The dissertation committee for Thomas Judson McClain II certifies that this is the approved version of the following dissertation:

Some Considerations in the Quantization of General Relativity

Committee:

Richard Matzner, Supervisor

Duane Dicus

Austin Gleeson

Richard Hazeltine

Paul Shapiro

Some Considerations in the Quantization of General Relativity

by

Thomas Judson McClain II

Dissertation

Presented to the Faculty of the Graduate School

of the University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

Doctor of Philosophy

The University of Texas at Austin

December 2018

Some Considerations in the Quantization of General Relativity

by

Thomas Judson McClain, Ph.D.
The University of Texas at Austin, 2018
Supervisor: Richard Matzner

In this dissertation, I explore a number of topics related to the quantization of the theory of general relativity. The first chapter presents a novel algorithm for determining the sky location of a gravitational wave source. The second chapter develops a new approach to polysymplectic covariant Hamiltonian field theory, then uses that approach to produce an original quantization procedure applicable to both particles and fields. The third and final chapter applies the quantization procedure of the second chapter to the particularly challenging case of general relativity. An appendix on noise in gravitational wave detectors and a Glossary of terms are included for the convenience of the reader.

Contents

Introduction

It is not an overstatement to say that the problem of quantizing general relativity is central to high energy theoretical physics. In essence, the problem is to formulate a quantum field theory counterpart to the classical field theory of general relativity. This theoretical problem has a number of potential solutions. String theory – in addition to purporting to be a theory of everything – is a possible solution to the problem of quantum gravity, in that every string theory builds in a spin 2 particle [1] indistinguishable from the graviton [2]. Loop quantum gravity is another relatively popular research program in quantum gravity, one in which the *only* goal is to define a quantum counterpart to general relativity (see, for example, [3] for a good overview).

The most natural approach to the problem is the first one, tried by Dewitt and others in the 1960s: to turn the classical field theory of general relativity into a quantum field theory of gravity in the same way – more or less – that we turn the classical field theory of electromagnetism into quantum electrodynamics [4]. This approach, however, runs into a number of technical challenges that do not exist for electrodynamics, the most damning of which is that the resulting theory is not perturbatively renormalizable, meaning that there are an infinite number of counter-terms in the perturbation expansion that must be defined by experiment if we hope to renormalize the theory (and thus render its predictions finite) [5]. Unless it is possible to find an analytical framework – thus avoiding the need to perturbatively renormalize the resulting theory – this is the end of the road: it is clearly impossible to do the infinite number of experiments needed to determine these counter-terms.

Indeed, the general feeling that has driven most high energy theorists to other, less conventional approaches is that the lack of perturbative renormalizability is a death-blow to the program, and the canonical quantization of general relativity is a dead end.

My research program begins by taking a different view. This is not to say that the canonical quantization of general relativity does not lead to serious mathematical difficulties (at least in the context of perturbation theory). It is instead to recognize that *every* quantum field theory arising from canonical quantization leads to serious mathematical difficulties; this is why all these theories need to be regularized and renormalized to begin with. But the fact that our ordinary tools for dealing with the mathematical pathologies of quantum field theory do not work with general relativity may interpreted another way: as reflecting poorly on the ordinary process of canonical quantization, rather than its application to general relativity. After all, the mathematical issues "solved" by the process of renormalization are solved in the physical sense, not the mathematical one: we can perturbatively renormalize some quantum field theories and therefore make experimental predictions, but the process is not one that most mathematicians would recognize as unambiguously well defined (see, for example, [6] for a general view or [7] for a more specific early example). So the question arises: are there other ways to formulate a quantum field theory starting from a classical one?

The answer here is an unequivocal yes: in addition to the path integral approach (which is plagued in most versions by similar and/or identical problems), we have also constructive [8], algebraic [9], and lattice gauge formulations [10] of quantum field theory. Each of these approaches has its own, different problems: either an extreme difficulty allowing useful calculations to proceed smoothly (as in algebraic and some constructive theories), or fundamental concerns about symmetries and arbitrary scales (as in lattice gauge theory). But do these alternatives exhaust the possible options?

To answer this question, it could be useful to look at the approaches to quanti-

zation that are used in classical particle physics. Here too there are issues (though not nearly as serious) with the canonical approach, and other approaches have been constructed to circumvent them: in addition to canonical quantization, there are alternatives like path integral quantization, geometric quantization, and deformation quantization, each of which solves one or more of the issues associated with the canonical quantization of classical particle systems. Could generalizations of these approaches be used to convert classical field theories into their quantum counterparts?

The reader familiar with quantum field theory will now object that the subject begins by considering the space of solutions to a set of classical field equations, which is an infinite-dimensional space to which most of the basic techniques of particle quantization seem to be inapplicable[1]. So the project seems to be stalled from the get-go. However, the notion of dimensionality is partially a matter of perspective: if we insist on dealing with the space of solutions "all-together", it is indeed infinite dimensional: change any solution at any point in space-time and it becomes a different solution, requiring a new dimension to add to the space. But if we characterize a solution by a section of a fiber bundle, the underlying bundle is almost invariably (for physical examples) finite dimensional. By working mostly with structures on the fiber bundle itself, rather than the space of all solutions directly, it is often possible to circumvent many of the hurdles that would otherwise prevent us from straightforwardly applying the techniques of particle quantization. In fact, although they are not very well-known there exist a number of different approaches to classical field theory that are essentially finite-dimensional, dealing primarily with structures on a finite-dimensional differentiable manifold and to which the quantization schemes of particles theory are therefore likely to be applicable. Perhaps the simplest of these is the formulation of Lagrangian field theory on the (finite-dimensional) jet bundle, as opposed to the infinite-dimensional space of solutions considered in the standard variational approach. However, since the process of canonical quantiza-

[1]Or at least to which we do not know how to apply them.

tion upon which we seek to improve relies on Hamiltonian rather than Lagrangian formulations, it may be more helpful to look for these kinds of approaches in finite dimensional classical field theory. Here again there is no shortage of candidates, going by names like polysymplectic and multisymplectic covariant Hamiltonian field theory.

My own review of these theories led to the conclusion that none was quite what I was looking for: a finite-dimensional approach to covariant Hamiltonian field theory that relies only on direct field-theoretic analogs to the symplectic structures used in most approaches to the quantization of classical particle theories. That being the case, chapter 2 of this dissertation is devoted to my own novel approach to polysymplectic covariant Hamiltonian field theory[2], as well as the ways the symplectic structures of this approach can be used to translate the method of canonical, geometric, and deformation quantization from classical particle physics to their field theoretic counterpart. Ultimately, I use the symplectic structures to reproduce manifestly finite canonical relations for a real-valued Klein-Gordon field equivalent to those of canonical quantization *after spatial integration* (see section 2.13 for details).

Despite its fundamentally more conservative approach to quantizing general relativity (as compared to loop quantum gravity or, especially, string theory), my research program is highly speculative. Indeed, this is true for all prospective quantum theories of gravity: even if we succeed in proving one formulation or another to be mathematically consistent and derive meaningful physical predictions (a monumental theoretical obstacle that has proven insurmountable for over five decades), we would still be in the painful position of having little or no experimental evidence to tell us if we had gotten it right!

There are many experimental efforts to probe strong-field gravitational phenomena that might eventually provide experimental input to the various approaches to quantum gravity. One such phenomenon is that of gravitational waves[3], which

[2]The field theory portion of this chapter was recently submitted to the *Journal of Mathematical Physics* [11] and is currently under review.

[3]See [12] for one example of how a quantum analysis of gravitational wave phenomena may be more straightforward than the full quantization of general relativity.

means that the relatively new field of gravitational wave astronomy could be of great interest to those attempting to quantizing general relativity. Now that the first handful of gravitational waves has been detected and the LIGO network has expanded to three detectors, the era of gravitational wave astronomy has really begun. As a small contribution to this new area of research, chapter 1 of this dissertation[4] details a novel approach to the problem of localizing gravitational waves sources in two detector networks. My long-term hope is that the field of gravitational wave astronomy will advance to the point where experimental results applicable to the problem of quantum gravity may be forthcoming.

[4]The first part of this chapter – dealing with the specialized case of circularly polarized gravitational waves – has been published by the *European Physical Journal Plus* [13], while the second part on elliptically polarized waves was recently submitted to the same journal [14].

Chapter 1

A Numerical Source Localization Algorithm for Two Detector Gravitational Wave Observatory Networks

1.1 Overview

This chapter presents a novel method for analyzing a signal in a LIGO-style interferometeric gravitational wave detector and extracting the most-likely source location. The method contains many departures from the standard protocols, some of the details of which will be explained in the next section. Due to the extraordinarily high computational costs of doing gravitational wave simulations, my analysis has proceeded in two stages. The first stage tested the basic ideas of the method in the limited situation in which the gravitational wave was assumed to be circularly polarized (an assumption that substantially reduces the computational cost of the simulations); see section 1.4 for details. Once the method was proven to work with circularly polarized gravitational waves, I conducted a more computationally expensive and time-consuming rounds of simulations removing the assumption of circular

polarization to test the method in a more realistic scenario; see section 1.5 for the results of this second analysis. Both sets of simulations showed good results, and I will go through each phase of testing in order.

1.2 Background

The use of modern computational methods to solve for the polarization amplitudes $h_+(t)$ and $h_\times(t)$ of a gravitational wave and to gain useful information about the direction to the source (θ, ϕ) goes back to the seminal paper by Gursel and Tinto [15]. That paper details a method for extracting all four of these parameters from the noisy output of a three detector system in a method that is all-but-optimal and that requires no assumptions about the nature of the source. However, even though the three detector problem has been essentially solved since 1989, it was already clear at that time that problems emerged in the case of a two-detector network, and later analysis [16] confirmed that the two-detector problem poses a number of additional challenges and does not admit unique solutions. The actual method implemented by the LIGO collaboration to determine source parameters in their gravitational wave burst search is implemented in a code package called LALInference [17]; an explication of this method can be found in [18]. Though the Bayesian inference framework of LALInference is the consensus of many teams and over a decade of research, the method as implemented in the first successful generation of LIGO science runs has at least three aspects upon which one might hope to improve: its reliance on model-specific waveforms, its high degree of final uncertainty, and its high computational cost. For example, the LALInference analysis of LIGO's early two detector data (VIRGO has since joined the network and participated in joint detections on August 14, 2017 and August 17, 2017) assumes that the source is one of only a few well-modeled astrophysical phenomena, reports sky-angles with uncertainties of $15 - 70\%$ and average errors of $1 - 10\%$, and takes on the order

of months of CPU time to run[1] [18]. Recent improvements to the basic methods of LALInference have enabled more rapid characterization of gravitational wave transient detection candidates of the type considered in this analysis [20] and much faster (on the order of minutes) determination of signal sky locations in some regions of the sky for some types of loosely-modeled sources [21]. Meanwhile, a substantially different basic method, though still very much in the mode of Bayesian analysis, represents another way to very rapidly (again, within minutes) determine the sky locations of gravitational wave transients at relatively high signal to noise levels, independent of assumptions regarding the signal source [19]. The results of this most recent generation of research show that rapid detection and sky localization is often (though not always) possible within the framework of Bayesian inference. However, as pointed out in [20], it is always worthwhile to have multiple algorithms available to cross-check and confirm results, especially if those algorithms differ enough in assumptions and implementation that they are unlikely to produce similar errors. It is in this spirit that I offer a rapid, non-Bayesian approach to sky location determination in gravitational wave detector networks.

The method to be discussed begins with a feature of the two detector system that was already known to Gursel and Tinto and is briefly discussed in [15]: if the source angles (θ, ϕ) are known, the two detector problem is analytically solvable, and gives a unique value for each of h_+ and $h_\times$ for each time at which the response of the two detectors is known. Unless the gravitational wave search is externally triggered (that is, induced by an optical astronomical observation), the sky angles are *not* known in advance. However, if an algorithm could be successfully implemented that could find the sky angles – without regard to any other model parameters – then the problem would be analytically solvable with this additional input.

The most obvious matching algorithm to try is very much like the one on which LALInference is based: simulate a great many possible signals and look for the one

[1]The possibility of parallelization means that CPU time is not the same as real time. However, many algorithms – including LALInference – make substantial use of techniques like MCMC that are notoriously resistant to parallelization. Nevertheless, the latest version of LALInference is actually run in about a day of real time [19].

with the smallest deviation from the actual detector responses. Even with a much-reduced model space (as compared to LALInference), this is method is workable, and can produce good results. However, it seems intuitively reasonable that, in the absence of any single excellent fit (as we must expect in the presence of substantial noise or if we do not wish to expend the extraordinary computational resources required to produce accurate results with LALInference), it should be the sky angles that produce the largest number of better-than-average fits that represent the true source location. This is the intuition underlying the new proposed fitting algorithms, and simulation results suggest that this method can not only reproduce the success of more standard fitting algorithms, but substantively improves upon the results of single best fit methods in almost every tested scenario.

1.3 Methods

To establish the validity of this modified approach, it was first necessary to write a simplified implementation of the single best fit method underlying most other approaches. I began with the starting assumption of a sine-Gaussian waveform, as used by the LIGO collaboration in their all-sky burst search event detection algorithms [22]. The success of a particular parameter set Θ is quantified by summing the absolute value of the difference between the algorithm's calculated responses R_{out} and the simulated detector responses R_{in} for each detector in the network N and each sampled time $t \in T$ in the lifetime of the signal:

$$Q(\Theta) := \sum_{i \in N\,,\, t \in T} \sqrt{\left(R_{i,out}(t,\Theta) - R_{i,in}(t)\right)^2} \tag{1.1}$$

The more common choice would be to normalize to the noise variance:

$$Q'(\Theta) := \sum_{i \in N\,,\, t \in T} \sqrt{\left(\frac{(R_{i,out}(t,\Theta) - R_{i,in}(t))^2}{\eta_i^2(t)}\right)} \tag{1.2}$$

where the term $\eta_i(t)$ represents the (estimated) noise in the i^{th} detector at time t. Since this analysis deals only with monochromatic signals, the noise is identical across all modeled parameter sets and there is no reason to normalize to the noise variance[2]. Since the algorithm is designed to sample the detector signal stream only very sparsely, one might expect poor results from allowing low-response pieces of the sample signal stream to be underweighted, as this effectively reduces the sampling rate even further. This intuition is not borne out by simulations over full signal lifespans, and eq. (1.1) is used throughout the analysis to determine the success of parameter sets.

The detector responses $R_i(t)$ are calculated in the standard way (see, for example, [23]) and in this analysis include both randomly generated noise and other, non-random but un-modeled contributions to the waveform (that is, signal that does not conform to the sine-Gaussian model and therefore cannot be readily fit by the algorithm). Following the conventions used by Schutz in [23], we compute the response functions

$$R_i(t) = h_+(t+\tau_i)F_i^+(\theta,\phi) + h_\times(t+\tau_i)F_i^\times(\theta,\phi) + \eta_i(t) \tag{1.3}$$

where h_+ and $h_\times$ represent the two independent polarization amplitudes of the incoming gravitational wave, $\tau_i = \frac{1}{c}(\vec{r}_0 - \vec{r}_i)\cdot\hat{e}_{gw}$ represents the time delay between the i^{th} detector and an arbitrarily chosen " 0^{th}" reference detector, F_i^+ and $F_i^\times$ represent the beam pattern response functions of the i^{th} detector (that is, the response of the i^{th} detector to a unit-amplitude, linearly polarized signal $h_+ = 1$ or $h_\times = 1$), and η_i represents the noise in the i^{th} detector. Because actual gravitational wave detectors have instrument noise that does not necessarily match the Gaussian noise model (see, for example, [24], as well the appendix to this work), this analysis assumes non-Gaussian noise. This noise is random-number generated, and is characterized throughout the analysis by its maximum allowed value within

[2]Of course, this would not be the case if non-monochromatic signals were being analyzed after Fourier decomposition. In this case, we would normalize to the noise variance as usual.

a given set of simulations, η_{max}, which is in turn set by the signal-to-noise ratio chosen for each simulation set: $\frac{\sqrt{h^2_{+,max}+h^2_{\times,max}}}{\eta_{max}} = \mathrm{SNR}$. The noise values in each detector are generated independently, and each is uniformly distributed within the range $[-\eta_{max}, \eta_{max}]$.

For the purposes of this algorithm, the beam pattern response functions are fully general (see, for example, [23] or [25], though I follow different angle conventions than the latter source). I modeled only monochromatic, sine-Gaussian signals of the form

$$s_{+,\times}(t, q, \omega, a_j) = \exp(-q^2t^2)(a_1 \cos\omega t + a_2 \sin\omega t) \tag{1.4}$$

though with the potential for the "real" (simulated) signal h_+ , $h_\times$ to be modified by an arbitrary, un-modeled function, which the algorithm includes up to fifth order in t :

$$h_{+,\times}(t, q, \omega, a_j, u_k) = \exp(-q^2t^2)(a_1 \cos\omega t + a_2 \sin\omega t) \times (1 + u_1t + u_2t^2 + u_3t^3 + u_4t^4 + u_5t^5) \tag{1.5}$$

As with the noise, these un-modeled signal amplitudes are random-number generated and are characterized throughout the analysis by their maximum allowed values within a given set of simulations. The fitting algorithms can easily be made to handle non-monochromatic signals after Fourier decomposition at the expense of greater computational cost; I have avoided these extra computational costs in my analysis.

Though my current modeling does not replicate much of the current state-of-the-art used, for example, in LALInference, the model does contain a few of the interesting features of LALInference, for example, quasi-random sampling of the parameter space [18]. It uses a more computationally intensive method of determining best fit [26]. More importantly, it is also designed to weigh many different fits with non-minimal Q values in the final determination of the "best fit" sky angles. Specifically, it allows searches other than the single best fit found by maximizing

the weighting function

$$W = \begin{cases} 1 & \text{if } Q = Q_{min} \\ 0 & \text{otherwise} \end{cases} \tag{1.6}$$

This weighting algorithm (the single best fit algorithm) simply looks for the parameter set with the minimum value of Q. My algorithm instead produces a weighted best fit after summing over all parameters with which we are not concerned by minimizing an arbitrary weighting function $W = f(Q/Q_{min})$

$$\min \left(\sum_{\Theta \setminus (\theta,\phi)} f\left(\frac{Q(\Theta)}{Q_{min}}\right) \right) \tag{1.7}$$

The output of the algorithm – which I call the "multiple close fit" algorithm – consists of the F^+, $F^\times$, and τ parameters that result from the (θ, ϕ) values of the parameter set that minimizes this weighted sum.

It is worth pointing out explicitly that we do not need the sky angles θ and ϕ if our only goal is to reconstruct h_+ and $h_\times$. Rather, it is $F_i^+(\theta,\phi)$, $F_i^\times(\theta,\phi)$, and $\tau_i(\theta,\phi)$ that determine whether a particular set of sky angles accurately determines the response functions measured in the detector network. Any sky angles that produce the same values of $F_i^+(\theta,\phi)$ and $F_i^\times(\theta,\phi)$ and the same time delays $\tau_i(\theta,\phi)$ will produce the same detector responses. Since the beam pattern response functions are multi-valued, there is typically not a single pair of sky angles determined by the algorithm. Consequently, for a two detector network what the algorithm most accurately determines are the beam pattern response functions of the two sites ($F_1^+, F_1^\times, F_2^+, F_2^\times$), as well as the time delay τ between the two sites. Again, it is this information that is needed to find $h_+, h_\times$ analytically (and again, the details can be found in [15]). The reason the algorithm computes all relevant parameters in terms of sky angles is to ensure that all modeled results are physically realizable. The actual values of these sky angles (θ, ϕ) are not necessary to solve the problem of determining h_+ and $h_\times$, but they can be recovered (non-uniquely) by finding

the intersection of the multi-valued inverses of the beam pattern and time delay functions.

1.4 Circularly polarized gravitational waves

In this analysis, I characterize the performance of the various algorithms by the RMS difference between the model's predicted values for F^+ and $F^\times$ at each site, the model's predicted τ between the two sites, and the actual values of those five functions calculated from the randomly generated values (θ, ϕ) that produce the simulated signal. Values with *in* subscripts denote the values that serve as the algorithm's (simulated) input values (in place of real signal data), and *out* subscripts denote the values predicted by the algorithm:

$$\delta F_{rms} := \frac{1}{4}\sqrt{(F^+_{1,out} - F^+_{1,in})^2 + (F^\times_{1,out} - F^\times_{1,in})^2 + (F^+_{2,out} - F^+_{2,in})^2 + (F^\times_{2,out} - F^\times_{2,in})^2} \tag{1.8}$$

$$\delta\tau_{rms} := \sqrt{(\tau_{out} - \tau_{in})^2} \tag{1.9}$$

For specificity, I continue to treat the two currently operating LIGO network detectors as the sites, with the Hanford detector set to be the reference site ($\tau = 0$).

To allow a modestly dense sampling of the angular parameter space, I first (and primarily) report the results of the highly specialized situation in which the frequency and q-factor of the incoming gravitational wave are known in advance and the wave is assumed to be circularly polarized (see the Glossary for the details of circular and elliptical polarization, [27] for more generic conditions, as well as section 1.5 below). This could, however, be a reasonable approximation to situations in which the inclination of the source system is known to be almost face-on, the signal is Fourier decomposed, and the analysis is performed on a handful of dominant frequencies. The parameter set used in these simulations is: $t \in [0, 0.19]$ s, $q = 4.3\,\mathrm{s}^{-1}$ (giving a half life of 0.5s), $f \in \{10, 100, 1000\}$ Hz, $\theta \in [0, \pi]$ rad, $\phi \in [0, 2\pi]$ rad, circular polarization, SNR $\in [3, 100]$, and the parameters u_1 through u_5 of eq. (1.5)

all capped at a single value $u_{\max} \in [1/100, 1/3]$. Time $t = 0$ is defined as the time of peak signal amplitude, and the small time interval used means that relatively little signal decay occurs over the times sampled in the analysis. Each simulation is run with different random values for the input parameters, as well as freshly randomized noise and un-modeled signal, all within the appropriate bounds. Few times are checked over the lifetime of the signal due to computational restrictions.

Naturally, the performance of a revised algorithm with weighting function $f(Q/Q_{min})$ depends entirely upon the function f. Extensive but non-exhaustive testing with weighting functions that are logarithmic, polynomial, and exponential in the argument Q/Q_{min} ultimately resulted in my choice to focus on weighting functions of the form

$$W = \exp\left[1 - \left(\frac{Q}{Q_{min}}\right)^n\right] \tag{1.10}$$

for values of $n \in \left[\frac{1}{2}, 64\right]$. These multiple close fit weighting functions peak at $Q = Q_{\min}$ with a value of 1, and reduce to exactly the result of eq. (1.6) in the limit $n \to \infty$. More intuitively, these functions look "almost" like the single best fit weighting function of eq. (1.6), but with some non-zero range over which "good" fits that are non-minimal can still contribute to the determination of the final angle values. For weighting value $n = 2$, the weighting function of eq. (1.10) is a Gaussian centered on the minimum Q value.

At every tested noise level there are one or more multiple close fit weighting functions of the form of eq. (1.10) that substantially improve upon the performance of the single best fit weighting function of eq. (1.6). As noise decreases, accuracy increases for the weighting functions of eq. (1.6) and eq. (1.10).

To give a sense of the raw accuracy of the method, figs. 1.1 and 1.2 show the cumulative probability distribution function of δF_{rms} over the frequencies and a few of the (maximum) SNRs and un-modeled signal fractions tested, respectively. As expected, higher frequency signals and higher maximum levels of noise and un-modeled signal substantially reduce the accuracy of the method, but we note that the method does not completely lose predictive power even at high frequencies and

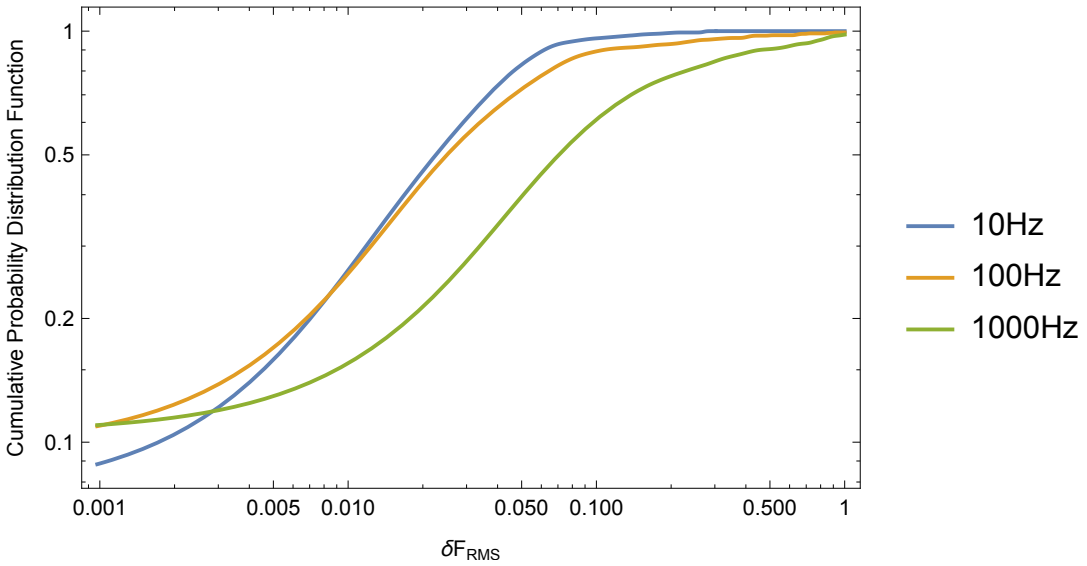

Figure 1.1: Cumulative probability distribution of δF_{rms} for the multiple close fit weighting function with weighting function value $n = 4$ at each tested frequency. The graph presents the specialized case of circularly polarized gravitational waves. The scaling of the axes is logarithmic in all graphs presented in this chapter.

high noise/un-modeled signal.

1.5 Elliptically polarized gravitational waves

Having established above that in the cases presented there are invariably a number of multiple close fit weighting algorithms that allow my revised source localization method to improve on the results of a single best fit analysis. Now I extend the analysis to the full parameter space of monochromatic all-sky burst searches.

As in my analysis of circularly polarized gravitational waves, I characterize the performance of the various algorithms by the RMS difference between the model's predicted values for F^+ and $F^\times$ at each site, the model's predicted τ between the two sites, and the actual values of those five functions calculated from the randomly generated values (θ, ϕ) that produce the simulated signal. I continue to report the results of the somewhat specialized situation in which the frequency and q-factor of the incoming gravitational wave are known in advance, but the waves are no longer assumed to be circularly polarized (see the Glossary for the details of circular and elliptical polarization). This is a reasonable approximation to situations in which the

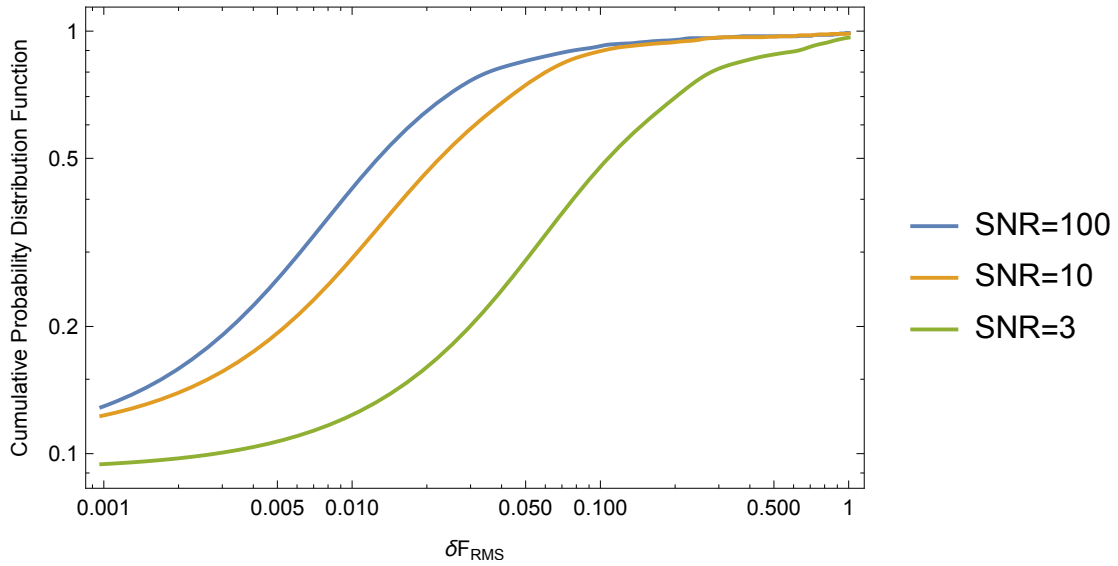

Figure 1.2: Cumulative probability distribution of δF_{rms} for the multiple close fit weight algorithm with weighting values $\{n = 2\,,\, n = 4\,,\, n = 2\}$ (respectively) for the given maximum values of noise and un-modeled signal. The graph presents the specialized case of circularly polarized gravitational waves.

signal is Fourier decomposed, and the analysis is performed on a handful of dominant frequencies within a handful of short time intervals. The parameter set used in generating the input signals for these simulations is: $q \in \{2.15, 4.29, 8.58\}\,\mathrm{s}^{-1}$ (see [3]) $f \in \{10, 100, 1000\}$ Hz (see [4]), $\theta \in [0, \pi]$ rad, $\phi \in [0, 2\pi]$ rad, SNR $\in [2, 100]$, and the parameters u_1 through u_5 of eq. (1.5) all capped at a single value $u_{\max} \in [1/100, 1/2]$. The same ranges of values are used to generate the output signals, with the exception that the noise and un-modeled signal values have no analog in the algorithm's output signal calculations. I have chosen the particular parameter values $q = 4.29\,\mathrm{s}^{-1}$, $f = 100\,\mathrm{Hz}$, SNR ≥ 10, and the parameters u_1 through u_5 of eq. (1.5) all capped at $u_{\max} = 1/10$ as the baseline. In figures 1.3 through 1.8 below, most parameters are set to their baseline values, while a single parameter is changed to a new (specified) value. A graph that shows, for example, $f = 1000$ Hz has that given frequency, with all other parameter set to their baseline values. Each simulation is run with different random values for the input parameters, as well as freshly randomized noise and un-

[3] These seemingly random q-values produce signal lifetimes – meaning times at which the signal drops to one-half its peak value – of 1s, $\frac{1}{2}$s, and $\frac{1}{4}$s, respectively

[4] The lower limit here pushes the state-of-the-art of what LIGO is capable of successfully detecting; we include it to show how the algorithm might behave in future GW detection scenarios

modeled signal, all within the appropriate bounds. Only a few times are checked over the lifetime of the signal (taken to be the time interval before it is suppressed by a factor of 2 due to its exponential decay) due to computational restrictions (see fig. 1.6 below for more information about the impact of this approach). All graphs in figures are the combined result of of 1000 simulations, each of which uses a single simulated signal (with noise as in eq. (1.3) and un-modeled contributions as in eq. eqrefeq:unmodeledsignal built in), but takes $N_t = 10$ random times, $N_{sd} = 100$ random source directions, and $N_{gwc} = 1000$ random gravitational wave amplitude combinations to generate model signals of the form of eq.
(1.4) to compare to the single simulated input signal.

Naturally, the performance of the revised algorithm with weighting function $f(Q/Q_{min})$ depends entirely upon the function f. I continue to focus on weighting functions of the form of eq. (1.10). However, in this larger parameter space, the specific values of n that I chose to focus on lie in the new range $n \in \left[\frac{1}{64}, 4\right]$. I began with the same, wider range reported in my analysis of circularly polarized gravitational waves above, but testing showed that the lower n-values here were generally more successful than the higher n-values I focused on above. All of these weighting functions still peak at $Q = Q_{\min}$ with a value of 1, and reduce to exactly the result of eq. (1.6) in the limit that the weighting value $n \to \infty$. More intuitively, these functions look "almost" like the weighting function of eq. (1.6), but with some non-zero range over which "good" fits that are non-minimal can still contribute to the determination of the final angle values. The baseline weighting value $n = 2$ was used to generate all the graphs that follow because it represents a well-understood Gaussian weighting function, and gives smaller (i.e., better) average values than the single best fit approach for the δF_{rms} and $\delta\tau_{rms}$ of eq. (1.8) and (1.9) that are used to characterize the accuracy of the algorithm's output in every scenario tested. However, in searching for the best possible weighting algorithms for specific parameter values, it is important to bear in mind that the weighting value $n = 2$ was often not the exponent with the smallest average values in eq. (1.8) and (1.9), and

therefore represents a sensible baseline choice rather than a universal best result.

As expected for this elliptically polarized scenario, at every tested noise level there is one or more weighting value for the multiple close fit algorithm of eq. (1.10) that substantially improves upon the performance of the single best fit weighting function of eq. (1.6); see fig. (1.4) for more details. Surprisingly, in this larger parameter space there are many scenarios in which *every* tested weighting value of the multiple close fit algorithm improves upon the performance of the single best fit weighting function of eq. (1.6), suggesting that the method offers a substantial improvement of the key performance indicators δF_{rms} and $\delta \tau_{rms}$ across a much wider range of weighting functions in this more realistic scenario. As noise decreases, the average value of both δF_{rms} and $\delta \tau_{rms}$ decreases for the weighting functions of eq. (1.6) and eq. (1.10). That the average values of the key performance indicators do not decrease more substantially at low noise levels is indicative of the fact that the parameter space is not being sampled densely enough to support many excellent fit results; see figs. 1.6, 1.7, and 1.8 to see the results of denser sampling.

In these graphs, the cumulative probability distribution function (CPDF) is plotted over the full possible range of δF_{RMS}; see eq. (1.8). These line graphs show the fraction of all fits produced by the algorithm in these simulations that have a δF_{RMS} below the value indicated on the horizontal axis. For example, the median value of δF_{RMS} can be found by finding the horizontal intercept of the value 0.5 on the CPDF; half of all fits had a δF_{RMS} value lower than this. Table 1.1 lists these median values for easy reference. More intuitively, these graphs give a sense of how accurately the algorithm is able to find the values of F^+ and $F^\times$ at each site.

Figs. 1.3, 1.4, and 1.5 show the cumulative probability distribution function of δF_{rms} over the frequencies, SNR/un-modeled signal fractions, and lifetimes indicated and show the method's capability to give good results over a range of parameter values within the larger parameter space. An ideal algorithm would give equivalent results (i.e., the same CPDF) across all possible frequencies, SNRs, and signal lifetimes, indicating that it could perform equally well throughout the entire parameter

Simulation	Median δF_{RMS}
f=10Hz	0.101
f=100Hz	0.195
f=1000Hz	0.219
SNR=100	0.218
SNR=10	0.195
SNR=2	0.211
q=8.58s^{-1}	0.184
q=4.29s^{-1}	0.195
q=2.15s^{-1}	0.204
N_t=1	0.358
N_t=10	0.195
N_t=100	0.122
N_{sd}=10	0.277
N_{sd}=100	0.195
N_{sd}=1000	0.161
N_{gwc}=100	0.189
N_{gwc}=1000	0.195
N_{gwc}=10000	0.174

Table 1.1: This table shows the median value of δF_{RMS} from eq. (1.8) in each set of simulations; elliptically polarized gravitational waves are used in these simulations. This is the same as the horizontal intercept of the value 0.5 on each of the CPDF graphs shown below. Each simulation set is referenced by its single, non-baseline value, except when the baseline value data itself is being presented. N_t stands for the number of time samples, N_{sd} for the number of source direction samples, and N_{gwc} for the number of gravitational wave amplitude combinations.

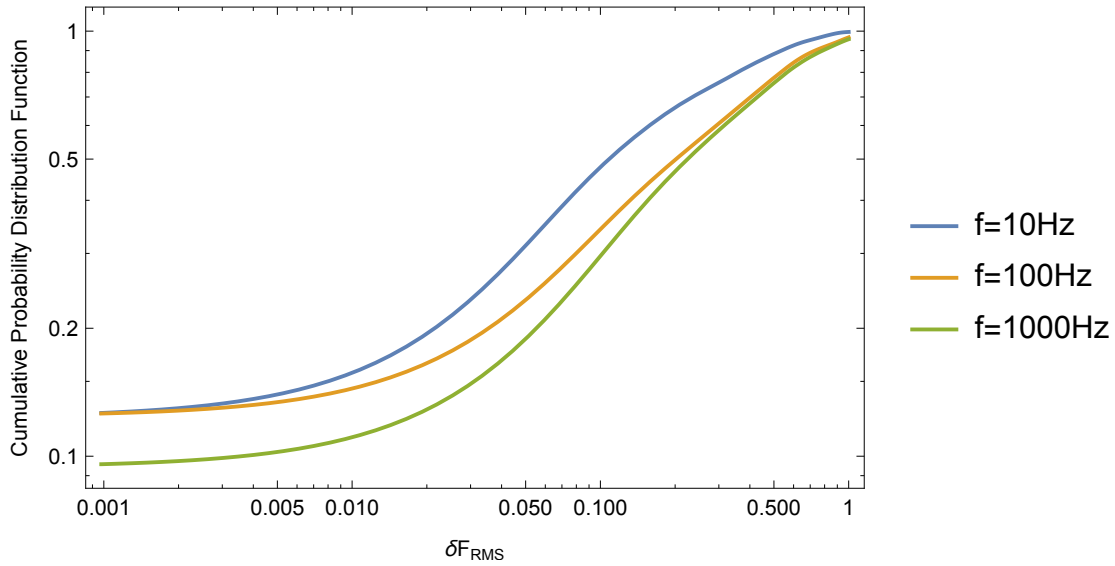

Figure 1.3: This figure shows the relatively small changes in RMS beam pattern response function error across three orders of magnitude in frequency of elliptically polarized gravitational waves. The lowest frequency data – near the limit of what LIGO can currently detect – shows substantially better performance when compared to the two more realistic scenarios. As noted in the text, the CPDF here allows us to see what fraction of simulations produced a δF_{RMS} (see eq. (1.8)) less than the value on the horizontal axis. For example, the median value of δF_{RMS} is the horizontal intercept of the value 0.5 on the line graph.

space. However, as seen in the circular polarization case above, higher frequency signals and higher maximum levels of noise and un-modeled signal reduce the accuracy of the method. Though the algorithm suffers at higher frequencies – an effect likely due to the fact that even small non-zero values $\delta\tau$ are amplified by the frequency when determining the mismatch of eq. (1.3) – fig. 1.3 shows that the method retains predictive power throughout the frequency range of these simulations. Fig. 1.4 shows that the algorithm proves surprisingly resilient to high levels of noise and un-modeled signal. Similar results were obtained above for circularly polarized GWs, but this unusually high tolerance is – as discussed above – likely due to the relatively coarse sampling of the parameter space that has been used in these simulations; I do not expect this level of resilience when the parameter space is sampled more densely. Only signal lifetimes show the ideal behavior: fig. 1.5 indicates that the impact of q-value/signal lifetime on the method's accuracy is negligible.

To give a sense of how the algorithm might fare when using higher sampling den-

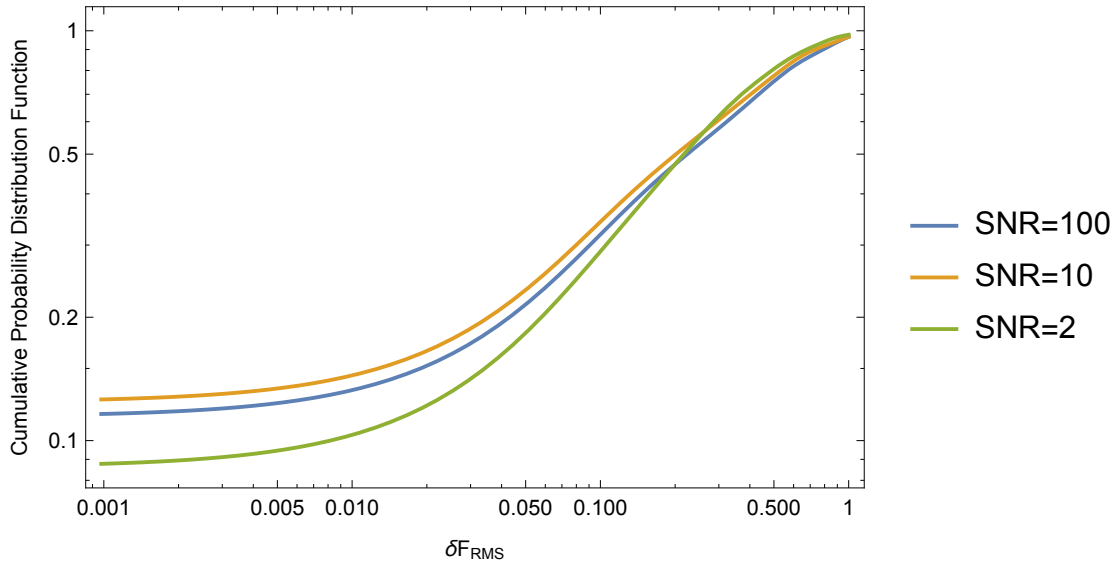

Figure 1.4: This figure shows very minimal changes in the RMS beam pattern response function error across a very wide range of signal-to-noise ratios (see eq. (1.3)) and fractional un-modeled signals (capped here at $\frac{1}{SNR}$; see eq. (1.5)) of elliptically polarized gravitational waves. Note, however, the clear decrease in the number of fits with $\delta F_{RMS} < 0.1$ at SNR 2 compared to the others. These minimal changes should not be expected to persist when the parameter space is sampled more densely.

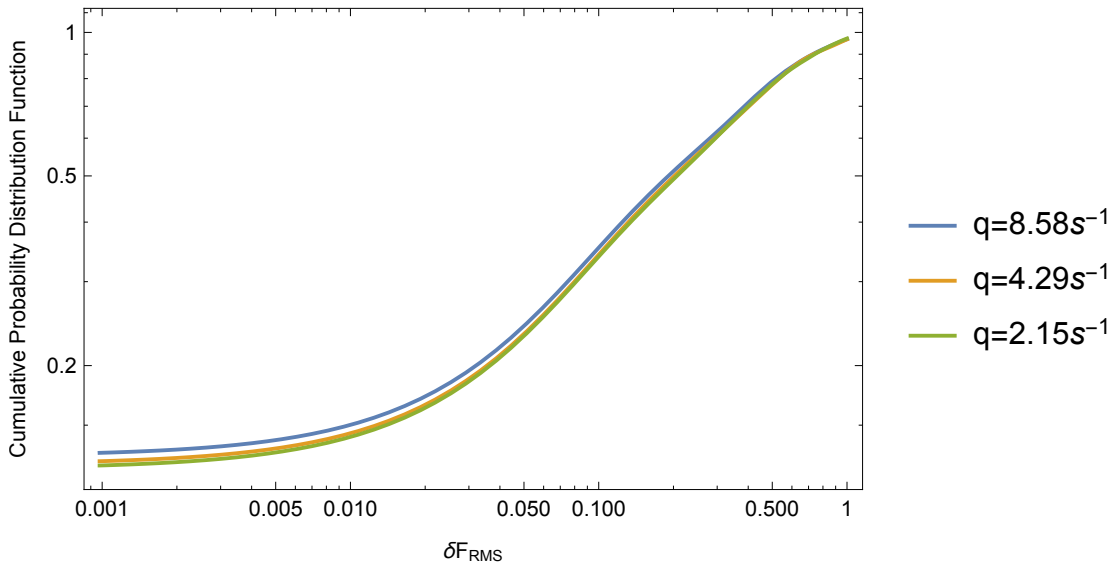

Figure 1.5: This figure shows that the algorithm handles signals with different q-values of elliptically polarized gravitational waves equally well. These q-values correspond to signal lifetimes – characterized by the time at which the signal amplitude drops to one-half its maximum value – of $\frac{1}{4}$s, $\frac{1}{2}$s, and 1s.

sities within this parameter space, I have done simulations that increase or decrease one of the three sampling densities while leaving the other two at their baseline levels; see figs. 1.6, 1.7, and 1.8. Here the ideal behavior is quite different from the three cases examined above: an ideal algorithm would show substantial improvement in the CPDF (an upward shift of the left-most portions of the graph) when any of its sampling densities is increased, indicating that the algorithm produces more accurate results at higher sampling densities. Similarly, the ideal algorithm would show a substantial worsening of the CPDF (a downward shift of the left-most portions of the graph) when any of its sampling densities is decreased, indicating that the ratio of sampling densities used in the algorithm is appropriate; i.e., there are no regions of the parameter space that are being oversampled relative to the others. Finally, an ideal algorithm would retain predictive power – and lower values of δF_{rms} and $\delta\tau_{rms}$ – even at lower sampling densities. My results show that the algorithm does retain predictive power – and continues to show smaller values of the key performance indicators δF_{rms} and $\delta\tau_{rms}$ as compared to the single best fit method – even at very low sampling densities. Also, the method's average δF_{rms} and $\delta\tau_{rms}$ values improve by significant amounts as any one of the sampling densities is increased. Increasing the time (N_t) and source direction (N_{sd}) sampling densities show a greater improvement in these key performance indicators than increasing the GW amplitude combination sampling density N_{gwc}. Similarly, decreasing the GW amplitude combination sampling density decreases the key performance indicators substantially less than decreasing either the time or the source direction sampling densities. Taken together, these results suggest that the ideal ratio of sampling densities is different from the baseline ratio used in these simulations *at the particular level of accuracy achieved in this analysis*. The fact that the far-left portion of the CPDF does rise considerably at very high gravitational wave amplitude combination sampling densities shows that this sampling density may actually be too low if our target accuracy is quite high. Determining the ideal ratio of sampling densities in these kinds of different scenarios will be an important consideration in future work.

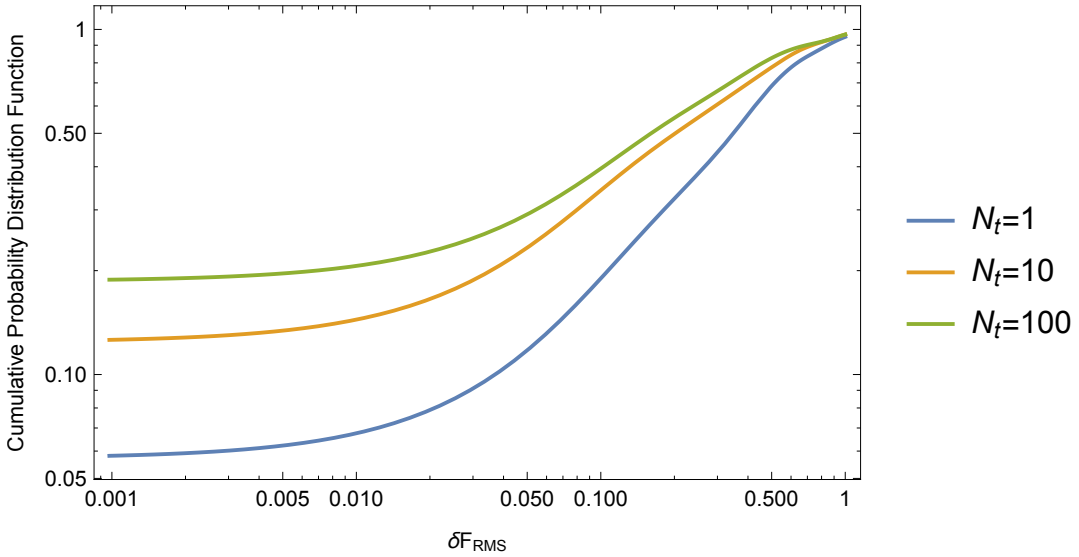

Figure 1.6: This figure shows the substantial differences in fit RMS beam pattern response function error when the time portion of the elliptically polarized gravitational wave parameter space is sampled more densely. The data indicates that the times are being appropriately sampled relative to source directions and GW amplitude combinations in the reference scenario. If times were being sampled much too densely then we would expect to see minimal change in the algorithm's ability to produce accurate predictions when the time sampling density was decreased. Conversely, if the times were being sampled much too sparsely then we would expect to see very large changes in the algorithm's ability to produce accurate predictions if they sampling density is increased. Instead, we see modest but significant changes in both scenarios. See section 1.5 for more details.

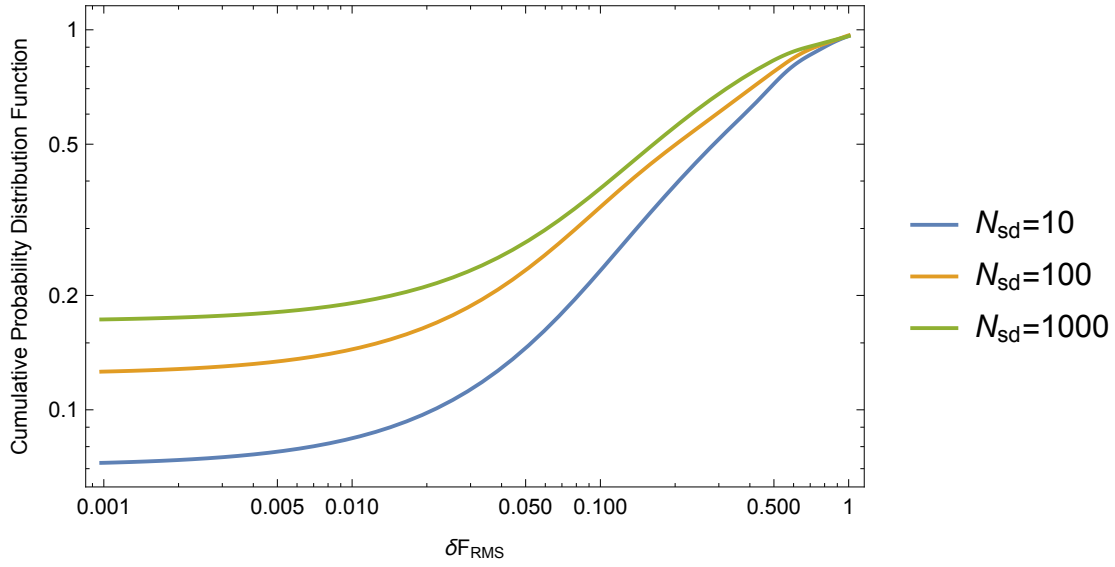

Figure 1.7: This figure shows the substantial differences in RMS beam pattern response function error across the different numbers of source direction samples of elliptically polarized gravitational waves. As with the time sampling data, this indicates that the source directions are being appropriately sampled relative to times and GW amplitude combinations in the reference scenario. See section 1.5 for more details.

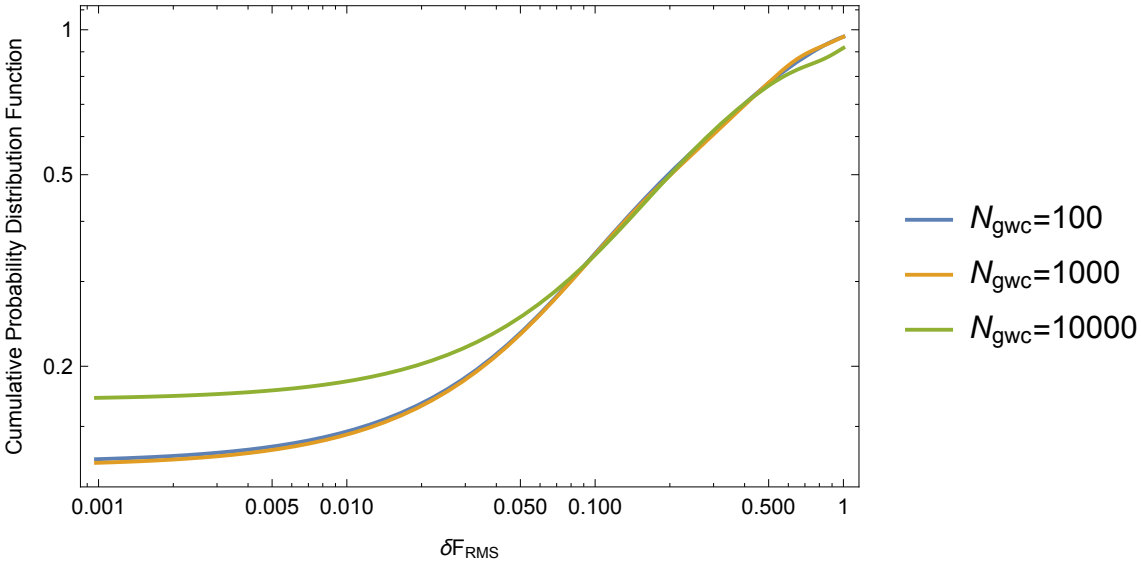

Figure 1.8: This figure shows the very small difference in RMS beam pattern response function error across the different numbers of GW amplitude combinations of elliptically polarized gravitational waves. This data is ambiguous, indicating both that GW amplitude combinations are being over-sampled relative to times and source directions in the reference scenario, but also that the sampling density would need to be significantly higher to produce high-accuracy fit results. See section 1.5 for more details.

1.6 Discussion

The a-priori circularly polarized simulations described in section 1.4 indicate that the multiple close fit weighting algorithm of eq. (1.10) outperforms the single best fit weighting of eq. (1.6) in many scenarios. Most importantly, my results indicate that the revised weighting is appreciably superior to the the single best fit weighting method in the presence of substantial noise and/or un-modeled signal; i.e., in precisely the scenario that most closely matches the actual all-sky burst search of the LIGO network. Indeed, we find that the weighting functions of eq. (1.10) with $n \in \{2, 4\}$ outperform the standard algorithm of eq. (1.6) at almost every noise level tested. This preliminary analysis indicates that use of a revised weighting algorithm of the form of eq. (1.10) is likely to produce a substantial improvement over the algorithm of eq. (1.6) even in situations in which the noise in the signal cannot be accurately quantified. Even better results can be obtained if the amount of noise can be quantified in advance of applying the weighting algorithm and the power n in the weighting function chosen appropriately; see figs. 1.1 and 1.2 for two circularly polarized examples. If the amount of noise cannot be quantified prior to the analysis, the weighting value $n = 4$ provides the best average result at the noise levels tested. To reiterate, the multiple close fit weighting algorithm can typically provide a substantial improvement in performance over the standard single best fit algorithm of eq. (1.6) regardless of whether or not the noise level is quantified before the analysis.

The elliptically polarized simulations described in section 1.5 show even greater promise. The percent increase in accuracy of the best-performing weighting function of eq. (1.10) over that of eq. (1.6) ranges from 14% to 26% in the scenarios tested, with the computational load increasing by only as much as 2.4% (and in some cases much less) due to the time it takes the algorithm to perform the additional weighting required by eq. (1.10) that is not required to compute eq. (1.6). Happily, both the higher fractional improvements in accuracy and the lower fractional increases in computational load occur at higher sampling densities – i.e., in the most realistic

scenarios tested. For reference, the baseline parameter values $q = 4.29\,\mathrm{s}^{-1}$, $f = 100\,\mathrm{Hz}$, SNR ≥ 10, the parameters u_1 through u_5 of eq. (1.5) all capped at $u_{\max} = 1/10$, $N_t = 10$ random times, $N_{sd} = 100$ random source directions, and $N_{gwc} = 1000$ random gravitational wave amplitude combinations yield a median accuracy increase of 17% for a median computational load increase of 0.1%.

There are two main results that suggest that the method remains promising in the larger parameter space used in this analysis. First, the weighting algorithm seems to perform almost equally well under small changes of the form of the weighting function, in contrast to the results of the circularly polarized GW analysis. Second, the weighting algorithm continues to show smaller average values of δF_{rms} and $\delta\tau_{rms}$ than the single best fit methods across a wide range of frequencies, q-values/signal lifetimes, and SNRs. In fact, by every calculated metric the multiple close fit weighting algorithm performs better in the larger parameter space than in the reduced parameter space of the circularly polarized GW analysis.

To re-iterate the results of section 1.5, the elliptically polarized GW simulations run for this analysis indicate that the sharply-decreasing exponential multiple close fit weighting of eq. (1.10) with weighting value $n = 2$ (i.e., a Gaussian weighting function) outperforms the single best fit weighting of eq. (1.6) in every scenario, including the presence of substantial noise and/or un-modeled signal. Though still not conclusive, the current analysis indicates that use of a revised weighting algorithm of the form of eq. (1.10) is likely to produce substantial improvement over the algorithm of eq. (1.6) in a realistic GW source localization scenario. Indeed, in the larger parameter space it is not even necessary to quantify the amount of noise in advance of applying the weighting algorithm to achieve near-optimal results (as was the case in the circularly polarized GW analysis).

To achieve the improved output of the multiple close fit algorithm in the context of an algorithm like LALInference, it would be necessary to use a finer mesh of parameter values, but the computational cost of doing so is vastly greater than that of implementing the revised weighting algorithm. Indeed, in the larger parameter

space of elliptically polarized GWs the savings in computational load is much greater than in the reduced parameter space of the circularly polarized GW analysis. To establish a baseline, first consider the cost of obtaining a particular result with the single best fit weighting function of eq. (1.6). Using this weighting function does not require us to calculate the large array of weighted Q values that we must have if we wish to use the multiple close fit weighting function of eq. (1.10); the computational cost of calculating this large array represents the additional computational load of the multiple close fit weighting algorithm. In the case of the monochromatic sine-Gaussian model (much simplified compared to LALInference, with only eight parameters), using the $n = 2$ weighting algorithm on a signal of frequency $f = 100$Hz with noise and un-modeled signal each less than 1/10 of the signal value is expected to produce about a 17% median improvement in accuracy based on the results of current simulations, while increasing the computational load of the model by less than 1% based on direct measurement of simulation runtimes. On the other hand, to achieve this same average improvement in accuracy using a finer mesh of parameters values requires (in a naive, best-case scenario) using a 17% finer mesh, which in turn increases the computational load of the model by $\approx 250\%$: every large array in the simulation grows by that same factor, and the computational cost of the simulation is dominated by a few computations involving those large arrays. In actual simulations, achieving a 17% increase in accuracy increases the computational load of the algorithm by 340% to 1300%, depending upon which parts of the parameter space are sampled more density to generate the improvement. The recommended strategy would therefore be to use the finest parameter mesh possible with a given set of computational resources, and then to implement the multiple close fit weighting algorithm to achieve a final 17% improvement in accuracy without much attendant increase in computational load.

The continued – indeed increasing – success of the multiple close fit weighting algorithm in the more realistic scenarios tested in this analysis suggests that it may be worthwhile to invest the time and resources necessary to test this algorithm in

real-world gravitational wave detection scenarios. To do this, two additional steps are necessary. First, the algorithm itself must be re-written in a more computationally efficient language to run in parallel on an arbitrary number of cores. Second, a very large amount of compute-time must be spent running the revised code on a highly parallelized network of CPUs to check the performance of the algorithm in many simulations of realistic GW search scenarios. It is my hope that this follow-up analysis will provide the necessary groundwork for these next steps.

The ideal algorithm would excel in all three areas where the current LALInference method could potentially be improved: its number of model-specific assumptions from computational general relativity, the uncertainty and error of the final fit parameters, and its computational cost. The method proposed in this analysis avoids any specific input from numerical general relativity, shows no sharp threshold below which it loses predictive power (as seems to be common with Bayesian methods; see, for example, [19] and [20]), and promises to be parallelizable to run in as little as 10 to 20 seconds on ordinary desktop computers. So the multiple close fit algorithm appears to make some progress in each of these areas. The present analysis does not prove that this algorithm can be made to reach the same degree of uncertainty or error attainable by LALInference. However, the apparent success of the multiple close fit algorithm in several areas of interest would seem to indicate that it may be worthwhile to spend the computational resources necessary to determine whether the new weighting algorithm can be made fully competitive in this regard.

The careful reader will have no doubt noticed that no new method has been suggested for using this algorithm's output (the beam pattern response functions F^+ and $F^\times$ at the two sites, as well as the time delay τ between the sites) to solve for the polarization amplitudes h_+ and $h_\times$. As alluded to in the introduction, this problem has already been solved [15]. The non-algorithmic novelty of this method consists in a change of perspective: we seek only to establish beam pattern response functions numerically, and then solve for the polarization amplitudes algebraically,

rather than trying to solve for these parameters (or, indeed, all relevant model parameters, as in the case of LALInference) numerically.

Even if this method is successfully implemented with uncertainty and error comparable to that of LALInference (while maintaining its relative stability under high noise and un-modeled signal and its relatively low computational cost), it is not designed to provide the other model parameters that can, in principle, be determined by LALInference. This inability is, of course, intimately related to one of the algorithm's major strengths, namely its lack of model-specific assumptions: without assuming a particular numerical model for the incoming signal, one cannot expect to extract model-specific parameter values for the source.

As a result, this method and algorithm should not be viewed as standing in competition with LALInference (or other similar modeling efforts). Its best possible use would be to supplement more model-specific and computationally expensive methods. For example, when a small enough range of true sky angle (θ, ϕ) pairs can be extracted from this algorithm's output, its relatively low computational cost (once parallelized) could allow it to be used as a cross-check for immediate (real-time) follow up to a detected signal using optical astronomy, something that has only recently become possible using Bayesian algorithms in the tradition of LALInference. Another possibility is that the algorithm can be used to rapidly characterize $(h_+ , h_\times)$ so that other (non-Bayesian) computational methods lying outside the scope of LALInference can be brought to bear on these intrinsic gravitational wave parameters as soon as possible after a signal is detected, then independently compared to the output of LALInference (and its rapid sky angle determination partners) once it is available. More detailed simulations will show whether this algorithm can achieve the low levels of uncertainty and error necessary to support these potential uses.

As VIRGO began participating in the LIGO network on August 1, 2017 and took part in a joint detections of GW170814 [28] and GW170817 [29] on August 14, 2017 and August 17, 2017, we are finally working with networks of three detectors, with more anticipated to come on-line in the future. Since the three-detector problem is

analytically solvable (as detailed in [15]), one may question the utility of an approach specific to two-detector networks. The first point to make is that there will be times in the future when only a subset of detectors within a larger network actually makes a detection. This may happen, for example, due to one or more detectors being offline, or only a subset of the detectors having the sensitivity to detect the signal. It is estimated that all three detectors will be simultaneously online only 50% of the time (see, for example, [19]), so this is by no means an unrealistic scenario. In this case, the two-detector situation may again become directly relevant. However, even in the case that three or more detectors all participate, the presence of noise makes analytical techniques potentially unreliable, and there may be benefit to a numerical approach like the one outlined in this analysis. This could be applied in several possible ways, either in a coherent approach (in the style of [30]) or by applying the method separately to each pair in the network and looking for coincidences (in the style of [31]). Future work will show whether this algorithm can support one or more of these potential applications.

Chapter 2

Quantization with Symplectic Structures

2.1 Overview

While the previous chapter was devoted to considerations related to potential experimental inputs to the program of quantizing general relativity (through the long term prospects of gravitational wave astronomy), this chapter is devoted to considerations related to the theoretical enterprise of quantizing general relativity. In the introduction, I claimed that there would be two phases to this program. The first is to generate a framework for classical field theory in which the tools of quantization familiar to theorists from the analysis of particle physics would more likely be applicable. The second is to generalize those quantization tools so that they in fact work for classical fields in the context of the new framework. In this chapter, I deliver on the goal of the first phase by constructing a novel approach to polysymplectic covariant Hamiltonian field theory (sections 2.3 – 2.8) and on the goal of the second phase by using the symplectic structures of that approach to quantize the simplest classical field: the Klein-Gordon field (section 2.13).

Much of the field theory portion of the chapter is devoted to showing that my approach reproduces standard results from covariant Hamiltonian field theory; the

genuinely new constructions are the polysymplectic structure of section 2.4 – which forms the foundation for the whole framework – and the generalization of the Poisson bivector in section 2.6 – which forms the foundation for my approach to quantization in the second half of the chapter. Since the Klein-Gordon field is the focus, I will use it as a running example throughout the field theory sections of the chapter, but the reader should bear in mind that all the results are completely general; indeed, the approach constitutes tremendous overkill for the simple example of a single, real-valued Klein-Gordon field on Minkowski space-time.

The second part of the chapter is more full of original material. Since my generalization of the Poisson bivector in section 2.6 is original and forms the foundation of my approach to quantization, all my examples are – in some important sense – new. However, since the ultimate goal of my quantization program is to quantize classical *fields*, it may seem like a detour when I first use the Poisson (and other) structures from my approach to recreate some of the key elements of classical particle quantization in section 2.12. This apparent detour will be justified by that fact that in my approach those elements are directly applicable to the field theory case in section 2.13, which constitutes the most important original result of the second half of the chapter.

2.2 Background

The study of Hamiltonian field theory from a geometric standpoint goes back at least to the pioneering works of Dedeker [32] and Goldschmidt and Sternberg [33] in the 1970s. This work was then taken up by Günther in the 1980s, and given a set of clear axiomatic foundations [34]. Though quite rigorous, Günther's original paper only dealt with the case in which the fields are sections of a trivial vector bundle (that is, the local version of the theory). One transition to a global theory for Günther's essential method was later provided by Carinena et al.. in [35], and the method was more recently revived and re-formulated in more modern mathematical language by Munteanu et al. in [36]. In addition to these more-or-less direct extensions,

Günther's work has been quite influential in the move toward a differential geometric foundation for covariant Hamiltonian field theory generally, as noted in many papers detailing different perspectives (see [37], [38], and [39], to name just a few key examples).

Though Günther's work was a major advance in the field, it was indeed only local in character, leaving an important gap that needed to be filled to provide a fully general treatment. In the extension of Carinena et al., the authors deviate substantially from the main path of Günther's paper, deeming the approach that too closely follows Günther's geometric foundation – the vector-valued polysymplectic structure – to be "artificial and unsatisfactory" [35] [1]. Later expositions and continuations of Günther's work also forgo its original geometric foundation in favor of other, more modern approaches (like the k-symplectic structures of [36]). To my knowledge, nowhere is Günther's original polysymplectic structure in its original form used to produce the full range of phenomena of covariant Hamiltonian field theory in the general (global) case. Many important physical phenomena in field theory are local in character, so a physicist might question the need for a global theory in the first place. However, the desire to apply field theories to cosmology and the importance of efforts to find correspondences like AdS/CFT that are fundamentally global in nature should convince us of this need. Thus, the first goal of this analysis is fill in this missing global link. I believe that I achieve that goal; see in particular section 2.5 below.

In doing so, we will find that the path from the starting fiber bundle through the polysymplectic structure to Hamilton's field equations is both feasible and well-defined, but it does require the substantial restriction of the domain of several key tensors to only vertical vectors[2]. The payoff of this path is a directly and naturally defined Poisson bracket on generic functions on the phase space of the

[1] A few comments made by those authors just before and after the quoted text shows that they understood several of the key points of the approach presented here. However, they seem never to have presented the details of that work.

[2] These restrictions may have been the grounds for the critique of [35], but those authors do not give the details that motivated their criticism.

theory: many of the common difficulties that surround the construction of natural, well-defined Poisson brackets in covariant Hamiltonian field theory (see, for example, [40]) are bypassed in this approach. Indeed, the outcome of the polysymplectic structure path seems to be the almost exact duplication of the natural results of the canonical transformation approach to covariant Hamiltonian field theory presented by Struckmeier and Redelbach in [41]. A second perspective on my approach to polysymplectic field theory is that it presents a natural geometric setting for [41] that directly reproduces most of the main results of [41].

One of the longstanding hopes for covariant Hamiltonian field theory is that it will provide a path for the geometric quantization of classical fields (see [42], [43], and [40], but note also [44]) and/or especially challenging classical particle systems (as in [45] and [46]). From this point of view a third, more aspirational perspective on the present analysis is that it presents another geometric framework for covariant Hamiltonian field theory, one which is at least superficially rather different from many of the other modern approaches and which may therefore prove itself to be more applicable to the as-yet-unsolved problem of the geometric quantization of classical fields. This topic is taken up in the second half of this chapter.

Given these three possible motivations for readers, I have tried to make my treatment as accessible as possible to as wide an audience as possible. I have avoided specialized or esoteric language and operations unless absolutely necessary, and I have explained such non-standard operations or language as I find it essential to employ. Though I give coordinate independent definitions of all structures and operations, I also provide local coordinate descriptions of all results, and the analysis can be followed by looking only at those local coordinate results (though in this case much of the motivation will be lost). As a concrete example, I apply my approach to the Klein-Gordon field at each step. It is my hope that a reader with a solid grasp of the fundamentals of differential geometry and elementary variational calculus will be able to follow the analysis with minimal difficulty.

In a field with so many different conventions and notations, a few remarks are in

order about which ones I will adopt throughout the analysis. Given a fiber bundle $\epsilon : E \to M$, I will denote the jet bundle simply as JE instead of J^1E (I will not use higher order jet bundles in this chapter), and the vector bundle upon which JE is modeled will be called $\vec{J}E$. Given a map $f : M \to N$, I will denote the tangent map $Tf : TM \to TN$ as Tf rather than f_*. I will use the notation $\Gamma(M, E)$ to denote sections of a fiber bundle E with projection map $\epsilon : E \to M$, and will often use notation like E_x to denote the part of one geometric structure (in this case the fiber of E) that lies over a specific point of another (in this case the point $x \in M$). I will use the notation $X \lrcorner \omega$ to denote the interior product of the vector-field X and the differential form ω. I make use of fibered local coordinate systems throughout the analysis, but nothing of the geometry is changed by using non-fibered coordinate systems (even though the local coordinate descriptions do change). The Einstein summation convention is employed throughout.

2.3 Geometric foundation

The theory begins with a fiber bundle $\epsilon : E \to M$ with total space E, base manifold M, projection map $\epsilon : E \to M$, and standard fiber Q. The base manifold M represents the space-time of our theory (typically represented by $\mathbb{R}^4$, but which we will assume is merely locally diffeomorphic to $\mathbb{R}^n$ for generality), while the space $Q \simeq E_x := \epsilon^{-1}(x)$ represents the space in which the physical field takes its values (typically $\mathbb{R}^N$). The total space E, therefore, represents possible pairings between space-time values and field values; in other words, all possible field configurations over space-time. E is the configuration space of our theory. Local fibered coordinates on E are given by

$$\{x^\alpha, \phi^I\} : E \to \mathbb{R}^{n+N} \mid e \mapsto x^\alpha e_\alpha + \phi^I e_I$$

In the case of the real-valued Klein-Gordon field, the base manifold $M = \mathbb{R}^4$ and the standard fiber $Q = \mathbb{R}$. The total space is therefore $\mathbb{R}^4 \times \mathbb{R}$; E is a trivial fiber

bundle. Coordinates on E are given by $\{x^\alpha, \phi\}$.

The next step is to construct the vector bundle

$$V := VE \otimes_E T^*M \simeq \vec{J}E$$

This vector bundle is isomorphic to the linearized first jet bundle over E; that is, the vector bundle upon which the ordinary jet bundle JE is modeled. The first jet bundle JE (an affine bundle over E) is the foundation for most discussions of covariant Hamiltonian field theory (see [40] for notation and details). More specifically, our bundle is defined to be the bundle over E with fibers $V_eE \otimes_E T^*_{\epsilon(e)}M$, where VE is the vertical bundle over E (see, for example, [40] and [38]) with its fibers defined by $V_eE := \{u \in T_eE \mid T\epsilon(u) = 0\}$, and $\epsilon(e)$ is the base space point over which $e \in E$ lies. We note, since it will be relevant later, that this bundle can be re-interpreted as a vector bundle over M rather than E. In the case of the real-valued Klein-Gordon field, coordinates on V are given by $\{x^\alpha, \phi, v_\alpha\}$ and a point $v \in V$ of this bundle is given by

$$v \in V = x^\alpha e_\alpha + \phi e_\phi + v_\alpha \frac{\partial}{\partial \phi} \otimes dx^\alpha$$

Intuitively, the final four coordinates v_α are related to the derivatives $\frac{\partial \phi}{\partial x^\alpha}$ [3].

The bundle that forms the foundation for the geometric structures of this analysis is really the dual of this bundle:

$$P := V^*E \otimes_E TM \tag{2.1}$$

This bundle has standard fiber $P_e := V^*_eE \otimes_E T_{\epsilon(e)}M$. Like the previous vector bundle, it can be interpreted either as a bundle over E with projection map $\pi : P \to E$ or as a bundle over M with projection map $\epsilon \circ \pi : P \to M$. We will make extensive use of both of these projections in defining the structures of our theory and their action on physical fields.

[3] I say "related to" rather than "coordinated by" because the bundle V is isomorphic to the linearized jet bundle $\vec{J}E$, not the true jet bundle JE. It is only in the true jet bundle JE (an affine bundle) that we can coordinate the bundle using the derivatives of the fields

P represents possible combinations of three things: 1) a point in space-time, 2) a point representing the values of all components of a physical field, and 3) a conjugate momentum value to each field value for each component of space-time. P is the phase space of our theory. Local coordinates on P that are compatible with a coordinate system $\{x^\alpha, \phi^I\}$ on E are given by

$$\{x^\alpha, \phi^I, \pi^\alpha_I\} : P \to \mathbb{R}^{n+N+nN} \mid p \mapsto x^\alpha e_a + \phi^I e_I + \pi^\alpha_I d\phi^I \otimes \frac{\partial}{\partial x^\alpha}$$

In the case of the real-valued Klein-Gordon field, coordinates on P are given by $\{x^\alpha, \phi, \pi^\alpha\}$, and a point $p \in P$ is given by

$$p \in P = x^\alpha e_\alpha + \phi e_\phi + \pi^\alpha d\phi \otimes \frac{\partial}{\partial x^\alpha}$$

The final four coordinates π^α are the conjugate momenta for the Klein-Gordon field. At this point, the major difference between the ordinary Hamiltonian approach to field theory and the covariant polysymplectic approach comes into sharp relief: there are four conjugate momenta to the single Klein-Gordon field, rather than a single one. This represents the fact that we have never broken space-time into space + time, and therefore cannot single out a single coordinate direction. Though these "extra" conjugate momenta look like they will overdetermine the physics of the system, we will find (section 2.5) that they are exactly what is needed. These additional conjugate momenta are a key feature of all covariant Hamiltonian field theories.

Finally, it is useful to consider the vector bundle

$$S := VP \otimes_P T^*M$$

where it is important to note that we are referring to the projection map $\epsilon \circ \pi$ in determining the vertical bundle VP. Local coordinates on S that are compatible

with a coordinate system $\{x^\alpha, \phi^I, \pi^\alpha_I\}$ on P are given by:

$$\{x^\alpha, \phi^I, \pi^\alpha_I, v^I_\alpha, \sigma^\alpha_{I\beta}\} : S \to \mathbb{R}^{n+N+2nNn^2N}$$
$$\mid s \mapsto x^\alpha e_a + \phi^I e_I + p^\alpha_I d\phi^I \otimes \frac{\partial}{\partial x^\alpha} + v^I_\alpha \frac{\partial}{\partial \phi^I} \otimes dx^\alpha + \sigma^\alpha_{I\beta} \frac{\partial}{\partial \pi^\alpha_I} \otimes dx^\beta$$

This bundle is useful only in relation to the bundle P; it has no obvious physical interpretation outside the mathematical structure of the theory. In the case of the real-valued Klein-Gordon field, coordinates on S are given by $\{x^\alpha, \phi, \pi^\alpha, v_\alpha, \sigma^\alpha_\beta\}$, and a point $s \in S$ is given by

$$s \in S = x^\alpha e_\alpha + \phi e_\phi + \pi^\alpha d\phi \otimes \frac{\partial}{\partial x^\alpha} + v_\alpha \frac{\partial}{\partial \phi} dx^\alpha + \sigma^\alpha_\beta \frac{\partial}{\partial \pi^\alpha} \otimes dx^\beta$$

There is nothing of any particular interest here except that the coordinates π^α and v_α are naturally dual to one another, a characteristic of the general case which I will exploit in the next section.

2.4 Polysymplectic structure

Looking at local coordinates on S makes it intuitively clear how we will define our canonical tensor θ (a vector-valued one-form): the third and fourth coordinates of a point $s \in S$ are naturally dual to one another. Define the local map:

$$\theta : S \to \mathbb{R} \mid (x^\alpha, \phi^I, \pi^\alpha_I, v^I_\alpha, \sigma^\alpha_{I\beta}) \mapsto p^a_I v^I_\alpha \tag{2.2}$$

In the case of the real-valued Klein-Gordon field, this specializes to:

$$\theta \mid (x^\alpha, \phi, \pi^\alpha, v_\alpha, \sigma^\alpha_\beta) \mapsto \pi^\alpha v_\alpha$$

This map can also be defined intrinsically. Since push-forwards by projection maps and identity maps on vector spaces are linear, the universal property of the tensor product [47] guarantees that the map $T\pi : TP \to TE$ can be uniquely extended to

$TP \otimes T^*M$, then restricted to $S := VP \otimes_P T^*M$. θ is therefore defined intrinsically by

$$\theta_p(u) := p \circ T\pi(u) \tag{2.3}$$

for any $u \in S_p$. In coordinates, the extension of $T\pi$ gives

$$T\pi(p) : S_p \to V_{\pi(p)}E \otimes T^*_{\epsilon\circ\pi(p)}M \mid v^I_\alpha \frac{\partial}{\partial \phi^I} \otimes dx^\alpha + \sigma^\alpha_{I\beta} \frac{\partial}{\partial \pi^\alpha_I} \otimes dx^\beta \mapsto v^I_\alpha \frac{\partial}{\partial \phi^I} \otimes dx^\alpha$$

so the intrinsic definition of θ leads to

$$\theta_p \mid u = u^I_\alpha \frac{\partial}{\partial \phi^I} \otimes dx^\alpha + u^\alpha_{I\beta} \frac{\partial}{\partial \pi^\alpha_I} \otimes dx^\beta \mapsto \pi^\alpha_I u^I_\alpha$$

meaning that the local coordinate representation of θ is

$$\theta_p = \pi^\alpha_I(p)\, d\phi^I \otimes \frac{\partial}{\partial x^\alpha} \tag{2.4}$$

In the Klein-Gordon example, this boils down to

$$\theta_p = \pi^\alpha d\phi \otimes \frac{\partial}{\partial x^\alpha}$$

As a linear map from S to $\mathbb{R}$, θ can be represented as a section of the dual bundle $S^* := V^*P \otimes_P TM$ with the local coordinate description given above.

Since θ is not simply a section of the exterior algebra of P, it does not have a canonical exterior derivative. Indeed, given its rather strange pedigree, defining any exterior derivative on θ is somewhat challenging (it is not even a vector-valued exterior form in the usual sense, since our vector fields are sections of TM, rather than P). Günther's original method in [34] was to limit the analysis to the local case (or, equivalently but perhaps more usefully, the global case in which the base manifold M is really just $\mathbb{R}^n$), then perform the exterior derivative operation on each coefficient one-form of the vector-valued one-form θ in canonical coordinates on M. For example, in the case of single real-valued Klein-Gordon field one can

take the exterior derivative in global coordinates on $P = \mathbb{R}^4 \times \mathbb{R} \times \mathbb{R}^4$ to get

$$\omega = d\theta = d\pi^\alpha \wedge d\phi \otimes \frac{\partial}{\partial x^\alpha} \tag{2.5}$$

Naturally, this approach does not work in the global case with a generic base manifold M. However, because of the vertical nature of the polysymplectic structure we are looking for, we can make this basic approach work with some modification.

Given any two vector fields u, $v \in \Gamma(P, TP)$ and any one-form $\beta \in \Gamma(M, T^*M)$, define the (non-tensorial) map ω via

$$\omega(u, v, \beta) := d(\theta \lrcorner \beta)(u, v) \tag{2.6}$$

where $\lrcorner$ denotes the contraction of the contravariant part of the tensor on the left with the covariant components of the form on the right. Though this is a well-defined map, it is not a tensor: the result of $d(\theta \lrcorner \beta)$ depends upon the particular one-form β, whereas a polysymplectic structure should be multi-linear. But looking at things in local coordinates makes it clear that defining ω only on vertical vector fields makes this structure unique:

$$d(\theta \lrcorner \beta) = d(\beta_\alpha \pi^\alpha_I d\phi^I) = \beta_\alpha d\pi^\alpha_I \wedge d\phi^I + \frac{\partial \beta_\alpha}{\partial x^\beta} \pi^\alpha_I dx^\beta \wedge d\phi^I$$

so inputting two vertical vector fields u, $v \in \Gamma(P, VP)$ gives

$$d(\theta \lrcorner \beta)(u, v) = \beta_\alpha (u^I v^\alpha_I - u^\alpha_I v^I)$$

This in turn means that we can write the resulting tensor $\omega_p : V_pP \times V_pP \times T^*_{\pi \circ \epsilon(p)}M \to \mathbb{R}$ as

$$\omega = d\pi^\alpha_I \wedge d\phi^I \otimes \frac{\partial}{\partial x^\alpha} \tag{2.7}$$

This is simultaneously a novel and comforting result. It is novel because it has been achieved without ever referencing anything like a k-symplectic (see Glossary)

or other over-arching geometric structure, resting only on the notion of vertical vectors intrinsic to the fiber bundle perspective. But it is comforting because it looks very much like the standard coordinate representation of Günther's original polysymplectic structure in [34] (essentially the generalization of eq. (2.5) above). Indeed, this coordinate form for ω defines local canonical coordinates in the theory, and it is the fact that this polysymplectic structure looks like a (vector-valued) collection of two forms that justifies the name poly-symplectic here as in its original context.

2.5 Hamilton's field equations

The vertical differential (see, for example, [48] for more details) $V\gamma : TP \to VP$ is defined by

$$V\gamma : TP \to VP \mid u \mapsto u - T(\gamma \circ \epsilon \circ \pi)$$

In local coordinates, this map looks like

$$V\gamma \mid u^\alpha_I \frac{\partial}{\partial \pi^\alpha_I} + u^I \frac{\partial}{\partial \phi^I} + u^\alpha \frac{\partial}{\partial x^\alpha} \mapsto (u^\alpha_I - u^\beta \frac{\partial \gamma^\alpha_I}{\partial x^\beta}) \frac{\partial}{\partial \pi^\alpha_I} + (u^I - u^\beta \frac{\partial \gamma^I}{\partial x^\beta}) \frac{\partial}{\partial \phi^I}$$

Intuitively, the vertical differential uses the information contained in a section γ to define – and then eliminate – the horizontal component of any vector tangent to the fiber bundle P, though it does so at the cost of altering the vertical part of the vectors. The vertical differential is a tensor, represented in local coordinates by

$$V\gamma = (d\pi^\alpha_I - \frac{\partial \gamma^\alpha_I}{\partial x^\beta} dx^\beta) \otimes \frac{\partial}{\partial \pi^\alpha_I} + (d\phi^I - \frac{\partial \gamma^I}{\partial x^\beta} dx^\beta) \otimes \frac{\partial}{\partial \phi^I}$$

Specializing to a single real-valued Klein-Gordon field $\gamma : M \to P \mid x^\alpha \mapsto \gamma(x^\alpha) e_\phi + \gamma^\alpha(x^\alpha) d\phi \otimes \frac{\partial}{\partial x^\alpha}$ gives

$$V\gamma = (d\pi^\alpha - \frac{\partial \gamma^\alpha}{\partial x^\beta} dx^\beta) \otimes \frac{\partial}{\partial \pi^\alpha} + (d\phi - \frac{\partial \gamma}{\partial x^\beta} dx^\beta) \otimes \frac{\partial}{\partial \phi}$$

Given the polysymplectic structure defined in eqs. (2.6) and (2.7), physical solution sections $\gamma : M \to P$ can be identified as those that satisfy

$$\omega(V\gamma, v) = dH(v) \tag{2.8}$$

at every point $p \in \mathrm{Im}\gamma$, for all vertical vectors fields $v \in \Gamma(P, VP)$. Here, $H : P \to \mathbb{R}$ is the covariant Hamiltonian function, which encodes the dynamical properties of the physical system under consideration; see section 2.8 for more details. Note that the restriction of eq. (2.8) to only vertical vector fields follows both from the restricted domain of ω *and* from the fact that Hamilton's field equations do not directly constrain the space-time dependence of H (so we must not end up with a piece like the carefully avoided $v^\alpha \frac{\partial H}{\partial x^\alpha}$ on the right hand side). In a local canonical coordinate system $\{x^\alpha, \phi^I, \pi^\alpha_I\}$, physical solutions $\gamma : M \to P$ must have components $\{x^\alpha, \phi^I(x^\alpha), \pi^\alpha_I(x^\alpha)\}$ that satisfy

$$\frac{\partial \phi^I}{\partial x^\alpha} = \frac{\partial H}{\partial \pi^\alpha_I} \tag{2.9}$$

$$\frac{\partial \pi^\alpha_I}{\partial x^\alpha} = -\frac{\partial H}{\partial \phi^I} \tag{2.10}$$

which are the standard form of Hamilton's field equations in covariant Hamiltonian field theory in canonical coordinates (see, for example, [41]).

Note that eqs. (2.9) and (2.10) do not constrain all the derivatives $\frac{\partial \pi^\alpha_I}{\partial x^\beta}$, only a particular sum of them (as is necessary to avoid over-determining the solution sections). Taken together, they serve to identify physically realizable field configurations (and their conjugate momenta) from unphysical configurations. Eq. (2.9) defines the relationship between the conjugate momentum coordinates and derivatives of the field-value parts of the solution section, while eq. (2.10) mimics the Euler-Lagrange equation of motion of the system when the relations imposed by eq. (2.9) are satisfied and the Lagrangian $\mathscr{L}$ is regular. The specific field configuration taken by a physical system is then determined by imposing appropriate initial

conditions and solving the system of partial differential equations. [41] provides the details of how these field equations apply a wide range of examples of physical interest.

To take one specific example, in the case of a real-valued Klein-Gordon field $\gamma : M \to P \mid x^\alpha \mapsto (x^\alpha, \phi(x^\alpha), \pi^\beta(x^\alpha))$ on Minkowski space-time, we have $H = \frac{1}{2}\eta_{\alpha\beta}\pi^\alpha\pi^\beta + \frac{1}{2}m^2\phi^2$ (where $\eta_{\alpha\beta}dx^\alpha dx^\beta$ is the Minkowski metric), giving

$$\frac{\partial\phi}{\partial x^\alpha} = \eta_{\alpha\beta}\pi^\beta$$

$$\frac{\partial\pi^\alpha}{\partial x^\alpha} = -m^2\phi$$

as Hamilton's field equations. Solving the first equation for $\pi^\alpha = \eta^{\alpha\beta}\frac{\partial\phi}{\partial x^\beta}$ and plugging it into the second equation gives us the more familiar Klein-Gordon equation of motion

$$\frac{\partial}{\partial x^\beta}\eta^{\alpha\beta}\frac{\partial\phi}{\partial x^\alpha} + m^2\phi = 0 = \frac{\partial^2\phi}{\partial t^2} - \nabla^2\phi + m^2\phi$$

2.6 Poisson brackets

In addition to the polysymplectic structure defined above, it will prove useful to define a second symplectic tensor that extends the notion of a Poisson bivector and is in some sense the inverse of the previous one. To do this, first note that it is possible to associate with any function f on P a *family* of sections $\mathscr{S}_f$, with each $s_f \in \mathscr{S}_f \subset \Gamma(P, S = VP \otimes_P T^*M)$, defined by the requirement:

$$\omega(s_f, v) = df(v) = v^\alpha_I s^I_\alpha - v^I s^\alpha_{I\alpha} \tag{2.11}$$

where the relation is expected to hold for all vertical vector fields $v \in \Gamma(P, VP)$.

In coordinates, eq. (2.11) means that the sections s_f all have components $\{\sigma^I_\alpha(p), \sigma^\alpha_{I\beta}(p)\}$ obeying the relations

$$\sigma^I_\alpha = -\frac{\partial f}{\partial \pi^\alpha_I}, \quad \sigma^\alpha_{I\alpha} = \frac{\partial f}{\partial \phi^I} \tag{2.12}$$

There is a family of sections rather than a single section associated with each function f because the second relation specifies only the trace of the second set of coordinate functions $\sigma^{\alpha}_{I\beta}$, rather than specifying every coordinate function uniquely. In local coordinates, the sections s_f look like:

$$s_f = -\frac{\partial f}{\partial \pi^{\alpha}_I}\frac{\partial}{\partial \phi^I} \otimes dx^{\alpha} + \frac{\partial f}{\partial \phi^I}\frac{\partial}{\partial \pi^{\alpha}_I} \otimes dx^{\alpha} + \sigma_{TF\, I\beta}^{\ \ \alpha}\frac{\partial}{\partial \pi^{\alpha}_I} \otimes dx^{\beta}$$

where the components σ_{TF} (TF stands for "trace-free") are arbitrary other than that they must obey the condition $\sigma_{TF\, I\alpha}^{\ \ \alpha} = 0$.

These families of sections are not especially important, except insofar as they allow us to define a Poisson tensor $\Pi \in \Gamma(TP \otimes TP \otimes T^*M)$. We do this by requiring that

$$\Pi(df) \in \mathscr{S}_f \tag{2.13}$$

for all functions $f \in C^{\infty}(P)$.

At first glance, this definition means that we must have

$$\Pi = -\delta^J_I \delta^{\alpha}_{\beta}\frac{\partial}{\partial \pi^{\alpha}_I} \otimes \frac{\partial}{\partial \phi^J} \otimes dx^{\beta} + (\Pi_{TF\, I\beta}^{\ \ \ J\alpha} + \delta^J_I \delta^{\alpha}_{\beta})\frac{\partial}{\partial \phi^J} \otimes \frac{\partial}{\partial \pi^{\alpha}_I} \otimes dx^{\beta}$$

where once again the notation "TF" in Π_{TF} refers to components that are arbitrary, except that this time they must obey $\frac{\partial f}{\partial \phi^J}\Pi_{TF\, I\alpha}^{\ \ \ J\alpha} = 0$. The fact that this same tensor is supposed to satisfy eq. (2.13) for all functions $f \in C^{\infty}(P)$, restricts the components of Π_{TF} significantly, but not enough to determine them all; as a tensor, Π is not yet well defined. If, however, we also demand that Π be anti-symmetric in its action on one-forms (so that, for example, $\Pi(df, df) = 0$), then all the components Π_{TF} are forced to 0. Perhaps somewhat surprisingly, the tensor Π – with the minimal additional requirement of anti-symmetry – is well-defined. Since we have made it anti-symmetric, we will typically write it in coordinates as:

$$\Pi = \frac{\partial}{\partial \phi^I} \wedge \frac{\partial}{\partial \pi^{\alpha}_I} \otimes dx^{\alpha} \tag{2.14}$$

Within the context of classical field theory, the primary significance of Π is that it allows us to define Poisson brackets in a natural and invariant manner: for any two functions $f,\, g \in C^\infty(P)$, we define their Poisson bracket $\{f,\, g\}$ by

$$\{f,\, g\} := \Pi(df, dg) \tag{2.15}$$

In coordinates, this reads

$$\{f,\, g\} = (\frac{\partial f}{\partial \phi^I}\frac{\partial g}{\partial \pi^\alpha_I} - \frac{\partial f}{\partial \pi^\alpha_I}\frac{\partial g}{\partial \phi^I})\, dx^\alpha \tag{2.16}$$

This means that the Poisson bracket of two functions is represented in this theory by a one-form over T^*M. It arises very naturally in this geometric setting, with all the natural transformation properties of the Poisson bracket of [41].

In the running Klein-Gordon example, this gives us

$$\Pi = \frac{\partial}{\partial \phi} \wedge \frac{\partial}{\partial \pi^\alpha} \otimes dx^\alpha \tag{2.17}$$

a tensor which can be interpreted as a one-form, each of whose components is an ordinary Poisson bivector.

The Poisson bivector and the Poisson bracket that it underlies are fundamental to many approaches to Hamiltonian field theory. In the context of this analysis, however, the real benefit of our Poisson tensor is that it allows us to perform quantization in myriad geometric ways. Indeed, we will find in section 2.12 below that the Poisson tensor – in addition to a few of the other structures from section 2.4 above – allow us to reproduce many of the results of canonical, geometric, and deformation quantization. That all of this is possible within the context of the unified approach to field theory laid out thus far is really quite remarkable.

2.7 Canonical transformations

The standard way to define the canonical transformations of a covariant Hamiltonian field theory like this one is to take them to be the symplectomorphisms of the appropriate symplectic structure (see, for example, [34] or [38]).

Following that standard, we say that a transformation $\Phi : P \to P$ is a canonical transformation if the pull-back by Φ leaves the polysymplectic structure unchanged, meaning

$$\omega(u, v, \beta) = \omega(T\Phi(u), T\Phi(v), \Phi^{-1\,*}(\beta))$$

for all vertical vector fields u, $v \in \Gamma(P, VP)$ and all one-forms $\beta \in \Gamma(M, T^*M)$. Since most authors take the transformation Φ to be a diffeomorphism anyway, the fact that the contravariant part of ω does not naturally pull-back under generic mappings (and therefore the map Φ must be a diffeomorphism) does not restrict the group of canonical transformations. Since ω is only tensorial for vertical vector fields, we must have $T_p\Phi(v) \in V_pP$ for all vertical vectors v at each point $p \in P$. This condition is genuinely restrictive. If the covariant part of ω is were simply an element of $T^*P \wedge T^*P$ (rather than having its domain restricted), the pull-back under canonical transformations would have been a much more generous condition, as only the vertical parts of the pushed-forward vector fields would have entered into the relation.

2.8 Legendre transformations and covariant Hamiltonian functions

To begin this section, it is important to note that it is a perfectly sensible approach to take the function $H : P \to \mathbb{R}$ as fundamental, simply a part – along with the fiber bundle E, its polysymplectic structure, and eq. (2.8) – of the data necessary to determine solution sections. However, it is far more common to choose the Lagrangian $L : JE \to \mathbb{R}$ or the Lagrangian density $\mathscr{L} : JE \to \Lambda^n(M)$ as fundamental,

and to derive the covariant Hamiltonian function H from L or $\mathscr{L}$. Following [41] (and almost everyone else), this is the approach I will follow.

To formulate H, begin by considering the Lagrangian L defined on the jet bundle JE with coordinates $(x^\alpha, \phi^I, v^I_\alpha)$. Remembering that the difference between any two elements of JE is an element of $\vec{J}E \simeq VE \otimes T^*M$, we define the vector Legendre transform[4] $\mathscr{V}L : \vec{J}E \to P$ by

$$\langle \mathscr{V}L(\gamma) \mid \gamma_1 - \gamma_2 \rangle := \lim_{\epsilon \to 0} \frac{L(\gamma + \epsilon(\gamma_1 - \gamma_2))}{\epsilon} \tag{2.18}$$

for all sections $\gamma \in \Gamma(E, JE)$. In terms of coordinates, this maps looks like

$$\mathscr{V}L : \vec{J}E \to P \mid (x^\alpha, \phi^I, v^I_\alpha) \mapsto (x^\alpha, \phi^\alpha, \frac{\partial L}{\partial v^I_\alpha}) \tag{2.19}$$

Define the covariant Hamiltonian on JE by

$$H : JE \to \mathbb{R} \mid \gamma \mapsto \frac{\partial L}{\partial v^I_\alpha}(\gamma)\gamma^I_\alpha - L(\gamma)$$

In the event that the mapping $\mathscr{V}L$ is bijective, the covariant Hamiltonian on P is then defined to be the pull-back by $\mathscr{V}L^{-1}$ of the Hamiltonian on JE:

$$H : P \to \mathbb{R} \mid p \mapsto H \circ \mathscr{V}L^{-1}(p) \tag{2.20}$$

which, in local coordinates, looks like

$$H = p^\alpha_I v^I_\alpha(p) - L(x^\alpha, \phi^I, v^I_\alpha(p)) \tag{2.21}$$

This is the same covariant Hamiltonian function defined in [41], and – in conjunction with the same Hamilton's field equations of eq. (2.8) – leads to same results in all the examples handled in that paper.

To go back to the now well-worn example of the single Klein-Gordon field γ :

[4] We use this non-standard notation to distinguish it from what most authors take to be the (ordinary) Legendre tranformation $\mathscr{F}L : JE \to J^*E$. For more details, see [38].

$M \to E \mid x^\alpha \mapsto (x^\alpha, \gamma(x^\alpha))$, the standard Lagrangian density $\mathscr{L} = \frac{1}{2}\eta^{\alpha\beta}\frac{\partial\gamma}{\partial x^\alpha}\frac{\partial\gamma}{\partial x^\beta} - \frac{1}{2}m^2\gamma^2$ gives $\mathscr{V}L \mid \frac{\partial\gamma}{\partial x^\alpha} \mapsto \eta^{\alpha\beta}\frac{\partial\gamma}{\partial x^\beta}$, which in turn yields the covariant Hamiltonian function $H = \frac{1}{2}\eta_{\alpha\beta}\frac{\partial\gamma}{\partial x^\alpha}\frac{\partial\gamma}{\partial x^\beta} + m^2\gamma$ on JE and, finally, the covariant Hamiltonian function $H = \frac{1}{2}\eta_{\alpha\beta}\pi^\alpha\pi^\beta + m^2\phi^2$ on P.

2.9 Concluding remarks on polysymplectic field theory

In this analysis, I have presented a novel approach to covariant Hamiltonian field theory that generates many of the key results of [41] – most importantly Hamilton's field equations, the form of the Poisson bracket, and the covariant Hamiltonian – in a more-or-less natural manner. In addition, the path I have taken generalizes Günther's approach in [34] in a manner that seems to me much more in keeping with Günther's original approach and spirit than most recent work. This accounts for two of the three motivations I outlined in the introduction to justify this enterprise.

The third motivation I offered for this particular treatment of covariant Hamiltonian field theory was that it might make possible a novel approach to geometric quantization. Though I have not yet successfully used this framework to reproduce all the results of (scalar) quantum field theory, it is a relatively simple matter to reproduce the canonical commutation relations for standard *particle* theories using the geometric structures of this analysis; indeed, the analysis can be successfully extended substantially farther, depending upon exactly which desirable properties one wishes the quantization procedure to reproduce[5]; see section 2.12 for details. On the field theory side, I do not produce the canonical commutation relations. However, it is possible to reproduce the results of the canonical commutation relations *after integration* (see section 2.13). This is a more reasonable expectation for a finite dimensional, differential geometric analysis of quantum field theory, as it is not clear how operator-valued distributions would ever arise from a well-defined differential

[5]That it will never produce all the properties one might reasonably expect for all smooth functions on the phase space is a result of Groenewald's theorem [49].

geometric structure (see, for example, [50] and [51] for similar efforts).

2.10 Introduction to quantization with symplectic structures

The main tools for taming the mathematical difficulties of mainstream quantum field theory – regularization and renormalization – though extremely successful in quantum electrodynamics are precisely where the canonical approach to quantization fails for general relativity: the result of applying the standard tools of canonical quantum field theory to general relativity is not perturbatively renormalizable, and therefore cannot be made to yield meaningful experimental predictions unless a way is found to avoid the need for perturbative analysis.

However, we often forget that, from a purely mathematical standpoint, pertubatively renormalizable quantum field theories like quantum electrodynamics are also problematic: many of the tools we regularly use to "tame the infinities" of renormalizable quantum field theories have thus far evaded rigorous mathematical analysis. The ordinary starting point for approaching the problem of quantum gravity is to say that general relativity is not quantizable in the same way as other classical field theories are, and then to ask: "well, what can we do now?" An alternative approach to the question of quantizing general relativity is to say "we don't really know how to quantize *any* classical field theories; how can we fix that?" It is in this spirit that I offer the following analysis on quantization with symplectic structures.

Since the quantization of particles is much better understood than that of fields, I will first look briefly at the way that symplectic structures from the first half of this chapter can be used to emulate the three main approaches to the quantization of classical particles systems: canonical, geometric, and deformation quantization. The point of this analysis is not to produce new results, but to show that many of the results of all three of these disparate approaches can be recreated under the single paradigm of quantization with symplectic structures. I will then straightforwardly

generalize a few of these tools from the particle case to produce a result equivalent to the canonical commutation relations of ordinary quantum field theory. This result is much more significant. The canonical commutation relations of momentum and position operators are fundamental to all quantum theories, and quantum field theory is no exception. This result suggests that it may be possible to base quantum field theory on the wholly new, fundamentally finite structures of the first half of the chapter. Though at least one other author has produced in [50] and [43] another result that seems to be equivalent to the canonical commutation relations of eq. (2.32), to my knowledge nowhere has an equivalent result been given from such a straight-forward procedure.

2.11 A brief review of the symplectic structures used in quantization

Before I begin talking about quantum particle theories from the various common perspectives, it may help to remind the reader about a few of the key pieces of symplectic structure upon which these quantization procedures will be based in my approach. The first and most important piece is the Poisson structure of eq. (2.13). It appears in my formulation of polysymplectic covariant Hamiltonian field theory (in canonical coordinates) as:

$$\Pi = \frac{\partial}{\partial \phi^I} \wedge \frac{\partial}{\partial \pi_I^\alpha} \otimes dx^\alpha$$

Intuitively, though it seems to have too many space-time indices on it, it looks like the kind of thing that will see if a function has any momentum dependence (by taking the appropriate momentum derivatives) and replacing it by an operator that takes derivatives with respect to the fundamental fields in the theory. Since this business lies at the heart of the process of turning classical observables into quantum operators (i.e., quantization), it is not hard to see intuitively why the Poisson structure will

form an essential part of the quantization process. The details, of course, are a bit more complicated.

Another key feature of quantization with symplectic structures is the canonical tensor of eq. (2.3), appearing in my field theory formulation as

$$\theta = \pi_I^\alpha d\phi^I \otimes \frac{\partial}{\partial x^\alpha}$$

where again we assume canonical coordinates. Above, I very briefly suggested that Π would be used to carry out the process of quantization, but the question naturally arises: with respect to what representation? The standard options are the position and the momentum representations, and the question of which we are going to use is answered (in favor of the position representation) by the canonical tensor, which acts to pick out a single representation from the various possible options.

Putting together these two structures gives us the last key symplectic structure we will use:

$$\Pi(\theta) = \pi_I^\alpha \frac{\partial}{\partial \pi_I^\alpha} \tag{2.22}$$

This structure was never used (or even defined) in the sections on classical field theory, but intuitively it will help to eliminate parts of functions that do not depend on the momentum variables, in effect enabling us to cancel them (or their complementary parts) out. How this works will be clearer once we begin dealing with the details of the process.

2.12 Quantization of classical particle systems

2.12.1 Quantization and the space of quantum states

The general consensus is that the primary goal of any quantization program should be to produce a map from the space of classical observables – represented by real-valued functions on the classical phase space – to the space of quantum operators – represented by linear operators on an appropriate space of quantum states.

The space of quantum states is itself not a trivial space, and in some cases there is considerable disagreement about the best way to explicitly construct this space. However, for the purposes of this analysis, I will usually consider quantum operators to be linear operators that act on complex-valued functions on the classical phase space. This is *not* a proper space of quantum states: it is much too big, being neither properly polarized (in the language of geometric quantization) – because the states can depend on both position and momentum – nor appropriately restricted to only normalizable states. Despite these problems, the space we will consider has the virtue of being easy to work with. Any fully realized quantization program would have to pare down this space of quantum states appropriately. In this analysis, I will only endeavor to illustrate how we can use symplectic structures to produce quantization maps that mimic some of the desirable properties of the current mainstream quantization programs, mostly leaving the task of paring down the space of quantum states for another work.

2.12.2 Canonical quantization

No approach to particle quantization can get off the ground without reproducing the canonical commutation relations of the basic coordinate functions of canonical coordinate systems. Given a quantization map Q, we must have

$$[Q(q^a), Q(p_b)] = \mathrm{i}\hbar\delta^a_b \tag{2.23}$$

This can be easily achieved with symplectic structures with the simple prescription

$$Q(f) = f - \Pi(\theta)f + \Pi(df)\llcorner \mathrm{i}\hbar\frac{\partial}{\partial t} \tag{2.24}$$

yielding the operators

$$\begin{aligned} Q(q^a) =& q^a + \mathrm{i}\hbar\frac{\partial}{\partial p_a} \\ Q(p_a) =& -\mathrm{i}\hbar\frac{\partial}{\partial q^a} \end{aligned} \tag{2.25}$$

Though the first operator doers not look right, it is in fact just fine provided that we deal with appropriately polarized wavefunctions in the position representation $\Psi \in C^\infty(P, \mathbb{C}) \mid d\Psi = \frac{\partial \Psi}{\partial q^a} dq^a + 0$ of the type used in geometric quantization: in this case, the quantum states have no momentum dependence and the extra derivative does no harm[6].

It might also appear that the extra step of needing to contract the symplectic structure Π with the vector field $\mathrm{i}\hbar\frac{\partial}{\partial t}$ is an ugly complication; after all, this would not happen in a particle theory where the parameter space $\mathbb{R}$ was suppressed (the standard approach). However, though it looks out of place here, it will turn out that it is in fact helpful in deriving the Schrodinger equation.

2.12.3 The Schrödinger equation

In each approach to quantization, we identify the Schrödinger equation from the relation:

$$\mathrm{i}\hbar\frac{\partial}{\partial t}\Psi = [Q(H), \Psi] = Q(H)\Psi \tag{2.26}$$

for any quantum state $\Psi \in C^\infty(P, \mathbb{C})$. Note that the left-hand-side of eq. 2.26 requires us to introduce the exact same vector field that we use to define the map Q above. In the absence of the parameter space, this vector field is not needed to define Q and instead comes out of nowhere at this stage. The process as a whole can in fact be viewed as more sensible in this approach, rather than less, since the same vector field is referenced in both critical phases of the quantization process – deriving the canonical commutation relations and deriving the Schrödinger equation – rather than only in deriving the Schrödinger equation. We will see a much more important example of this when we get to field theory, where the parameter space is much harder to suppress.

It is important to note that the simple quantization map of eq. 2.24 will *not* produce the correct Schrödinger equation, because it will not typically map H to

[6]In fact deformation quantization takes this approach of neglecting extraneous derivatives much farther; see section 2.12.4 for more details.

the appropriate quantum operator. For example, given a one-dimensional harmonic oscillator Hamiltonian of the form $H_{SHO} = \frac{1}{2m}p^2 + \frac{1}{2}m\omega^2 q^2$ we get the quantum operator

$$Q(H_{SHO}) = \frac{1}{2}m\omega^2 q^2 - \frac{1}{2m}p^2 + \mathrm{i}\hbar m\omega^2 q\frac{\partial}{\partial p} - \mathrm{i}\hbar\frac{p}{m}\frac{\partial}{\partial q}$$

which has a lovely antisymmetry to it, but is not even close to right. We will see that other approaches to quantization can do a better job of reproducing the correct Schrödinger equation.

2.12.4 Deformation quantization

Since our approach gives us direct access to the Poisson structure Π, it is not hard to reproduce many of the results of deformation quantization (see, for example, [52]). In the context of reproducing the standard basic results, it is helpful to choose local coordinates and rewrite the (contracted) Poisson structure as $\Pi \lrcorner \frac{\partial}{\partial t} = \alpha^{ij}(p)\partial_i \wedge \partial_j$. For example, standard local canonical coordinates in three dimensions gives $\Pi \lrcorner \frac{\partial}{\partial t} = \frac{\partial}{\partial q^1} \wedge \frac{\partial}{\partial p_1} + \frac{\partial}{\partial q^2} \wedge \frac{\partial}{\partial p_2} + \frac{\partial}{\partial q^3} \wedge \frac{\partial}{\partial p_3}$. Then locally the result of deformation quantization can be written (up to third order in $\hbar$) as:

$$\begin{aligned} f \star g = fg + \hbar \sum_{i,j} \alpha^{ij}\partial_i(f)\partial_j(g) + \frac{\hbar^2}{2}\sum_{i,j,k,l}\alpha^{ij}\alpha^{kl}\partial_i\partial_k(f)\partial_j\partial_l(g) \\ + \frac{\hbar^2}{3}\sum_{i,j,k,l}\alpha^{ij}\partial_j\left(\alpha^{kl}(\partial_i\partial_k(f)\partial_l(g) - \partial_k(f)\partial_i\partial_l(g)\right) + \mathcal{O}(\hbar^3) \end{aligned} \quad (2.27)$$

There does not appear to be any obvious way to get the various counting factors that are derived from graph theory in deformation quantization directly from the symplectic structures, but the key point is that the Poisson structure Π is the core of deformation quantization, and it appears naturally and essentially in symplectic structure quantization, providing a natural bridge between the two methods.

It is important to note here that deformation quantization has little or nothing to say about the space of quantum states on which the $\hbar$-deformed algebra of functions ultimately operates in quantum theory. Understanding how to reproduce the core

results of deformation quantization in a framework where the symplectic structures that underlie the method arise naturally and geometrically may allow us to supply this missing mathematical structure.

2.12.5 Geometric quantization

Actually, we have already seen how symplectic structures can be made to reproduce some of the key results of geometric prequantization: eq. (2.24) was in fact the precursor to the now standard method of prequantizing with connections over complex line bundles (see the Glossary or [53] for more details). So it is quite easy to reproduce the results of the historical precursor to what is now called geometric quantization. However, since connections do not play a role in my approach to polysymplectic covariant Hamiltonian field theory, it is *not* easy to reproduce the now-standard formulation of geometric prequantization[7].

A critical area in which geometric quantization differs from other approaches is in what it has to say about the space of quantum states. Indeed, where deformation quantization mostly ignores the issue of the space of quantum states, geometric quantization uses a careful construction of the space of quantum states as a key ingredient in its approach. Insisting on using only symplectic structures means that it is not possible to precisely follow the program of geometric quantization that allows us to choose an arbitrary polarization of the tangent bundle to define our space of quantum states. Luckily, however, there is a natural polarization to use.

Consider the space of complex functions over P that satisfy the condition

$$\Psi \in C^\infty(P, \mathbb{C}) \mid v(\Psi) = 0 \, \forall \, v \in VP \tag{2.28}$$

Let us call the space of these functions W. Intuitively, these are the functions that do not depend upon the momentum coordinates at all, only on the position coordinates

[7] In fact, a great deal of care is taken to avoid making the theory dependent upon any particular arbitrary connection, as this runs counter to the success of the enterprise when one is dealing with classical field theory. However, it is interesting to note that there could be an association in this regard between *non-standard* polysymplectic formulations of classical field/particle theory and geometric quantization. This possible link may be worth exploring in future work.

(and possibly time). That is, in local canonical coordinates we have

$$w \in W = w(t, q^a)$$

Now we have at least taken a first pass at paring down to a more reasonable space of possible wavefunctions[8].

2.13 Classical fields

The goal of all this preliminary work is to use symplectic structures to quantize classical *field* theories; this is the process that so many physicists have struggled with for so long. To do this, we will freely take inspiration from each of the approaches analyzed above: canonical, deformation, and geometric quantization. Ultimately, we will successfully recover the canonical commutation relations of a real-valued scalar field *after spatial integration*; see eq. (2.32) below.

2.13.1 Input from the canonical approach

Somewhat amazingly, I can write down the quantization map Q of eq. (2.24) without modification and it is sensible for use in the case of classical fields in the formulation of the first half of this chapter [9]:

$$Q(f) = f - \Pi(\theta) f + \Pi(df) \llcorner \mathrm{i}\hbar \frac{\partial}{\partial t}$$

This is a sensible map, but a complication arises from the fact that the base space of classical field theories is not time but space-time: there is no longer necessarily an unambiguous vector field $\frac{\partial}{\partial t}$ to use. We can get around this if – as is certainly the case in standard quantum field theory – there is a pseudo-Riemannian metric

[8] To make this more rigorous, we would need to make sure that admissible wavefunctions were also square integrable. However, this takes us into issues of whether or not the operators we will define will map square-integrable functions to other square-integrable functions, which is a bit outside the goals of the current, somewhat simplified, analysis.

[9] A similar method is applied in [54], but since our formulation of classical field theory is different, our results are also different

on space-time that allows us to identify a global unit-length time-like vector field. I assume that this is the case in what follows.

Taking again the case of a single, real-valued Klein-Gordon field, we compute the operators:

$$\begin{aligned} Q(\phi) =& \phi + \mathrm{i}\hbar\frac{\partial}{\partial\pi^0} \\ Q(\pi^0) =& -\mathrm{i}\hbar\frac{\partial}{\partial\phi} \end{aligned} \tag{2.29}$$

giving the commutation relation

$$\left[Q(\phi), Q(\pi^0)\right] = \mathrm{i}\hbar \tag{2.30}$$

This does not match the canonical commutation relation from standard quantum field theory:

$$\left[\hat{\phi}(\vec{x}), \hat{\pi}^0(\vec{y})\right] = \mathrm{i}\hbar\delta(\vec{x} - \vec{y}) \tag{2.31}$$

I label these operators with "hats" rather than making them arguments of quantization maps to clearly indicate the difference between these standard results and the operators computed from eq. (2.24). Indeed, the standard commutation relation is impossible to get in finite dimensional classical field theory, as we are not dealing with the full space of solutions of the Klein-Gordon equation but with ordinary coordinate functions on a finite-dimensional manifold; it is not clear how we would ever get a distribution (rather than an ordinary function) as an output.

This difference may be a blessing rather than a curse: many of the fundamental problems in quantum field theory arise from the fact that the canonical commutation relation (among many other important quantities) arises not from operators but from operator-valued distributions. These operator-valued distributions have no meaning apart from the integrals and test-functions that allow us to extract finite results from them. Nonetheless, there is a close relationship between the commutation relation

of eq. (2.30) and the standard one of eq. (2.31):

$$\big[Q(\phi), Q(\pi^0)\big] = \mathrm{i}\hbar = \int \star g(\frac{\partial}{\partial t})[\hat{\phi}(\vec{x}), \hat{\pi}^0(\vec{y})] = \mathrm{i}\hbar \int d^3x \delta(\vec{x} - \vec{y}) \tag{2.32}$$

where I have used the fact that the space-time metric allows us to define both a dual one-form $g(\frac{\partial}{\partial t})$ and a Hodge-dual to that one-form $\star g(\frac{\partial}{\partial t}) = d^3x$ [10]. Intuitively, the finite-dimensional classical field theory commutation relations from quantization with symplectic structures match not the raw commutation relations of infinite-dimensional quantum field theory, but the "tamed" commutation relations of ordinary QFT after spatial integration. Note that – as when we introduced the vector field $\frac{\partial}{\partial t}$ in deriving the map eq. 2.24 from the symplectic structures of the first half of the chapter – there is again the phenomenon adding a new structure that seems like it should not be necessary (in this case, a pseudo-Riemmanian metric g on space-time), only to find that it is essential not just where it was first introduced but to make other, perhaps deeper, connections to the canonical formulation.

The commutation relation of eq. (2.30) and eq. (2.32) is the key result of this section (and indeed, the whole chapter), as it underlies much of the structure of quantum field theory. However, it is possible to borrow more methods from the other approaches analyzed above to take this construction a bit further. For example, the idea from geometric quantization that the space of possible quantum states should consist of only polarized functions (as in section 2.12.5) can be used to reduce the space of quantum states in field theory, using the exact same form of eq. 2.28 [11]:

$$w \in W = w(t, \vec{x}, \phi) \tag{2.33}$$

meaning, intuitively, that wavefunctions can depend upon time, position, and field-value, but not on any of the conjugate momenta.

We could even borrow the construction of eq. (2.27) and modify it to better suit

[10] The far right hand side of this equation assumes that we are dealing with standard Cartesian coordinates on Minkowski space-time; everything else is completely general.

[11] The space P is, of course, different, which accounts for the different result

the less-formal results of geometric and canonical quantization, ultimately defining a quantization map that would produce reasonable results for classical observables beyond the basic coordinate functions. However, this process as I currently understand it involves a few seemingly arbitrary choices and would force the analysis even deeper into speculative territory. Such an analysis will be one focus of future work.

2.14 Discussion

One question we might ask of this analysis is: to what degree can the quantization methods successfully applied to particle theories be applied to field theories using only symplectic structures available in the formalism of the first half of this chapter? The answer would seem to be: very far indeed. I have found no clear mathematical obstacles to producing the key result of eq. (2.32), so things appear to be quite promising at the fundamental level considered here.

This by itself is a very interesting result, as a universal approach to quantization has thus far been very much lacking from our theoretical toolbox. However, we should not celebrate too early, as the immediate follow-up question must be: how well does this universal approach succeed in reproducing the key results of standard quantum field theory? Here, the answer is not so simple. We saw in section 2.13 that this universal approach does not reproduce the standard commutation relations of quantum field theory. Instead, it reproduces the integrated commutation relation of eq. (2.32) in place of the canonical result.

A final question is: what are we to make of this difference? The analysis is so far unclear. It may be that this result is tremendously beneficial, evidence that the universal approach will produce finite results in other aspects of quantum field theory as well, without ever requiring us to use the tools of regularization or renormalization that fail so badly for general relativity. I think this view is overly optimistic. Indeed, we will likely find that it is impossible to reproduce some of the results of quantum field theory – finite or otherwise – in a framework that avoids working with the space of classical solutions, as we do in this analysis by using the approach of the

first half of this chapter. A good deal more work will be required to put this final question to rest.

Chapter 3

Applications to General Relativity

3.1 Overview

As indicated at the end of the last chapter, my novel approach to the quantization of classical fields has by no means been proven successful. Though it reproduces a handful of critical results – most importantly the integrated canonical commutation relation of eq. (2.32) – and seems promising as a method, it has not yet reproduced most of the key results of standard quantum field theory, even for the simplest case of the real-valued Klein-Gordon field. It is therefore premature to apply this method to the much more challenging case of general relativity.

However, as the title of this dissertation promises that I will consider how to quantize general relativity, I will now do exactly that. The reader is strongly cautioned to take the analysis of this chapter with a very large grain of salt: not only are the quantization methods of chapter 2 not yet proven, there are also a number of critical choices that will have to be made to apply these methods to general relativity. The application of the basic procedure of section 2.13 to general relativity is the subject of section 3.2. Somewhat astonishingly, the analysis is relatively straightforward. But as I will show in section 3.3, the application of covariant Hamiltonian field theory to general relativity is quite the opposite[1]. Indeed, despite produc-

[1] This is by no means unique to the polysymplectic approach of chapter 2: applying *any* covariant Hamiltonian formalism to general relativity is fraught with difficulties.

ing perfectly reasonable quantum theories, all of the approaches I examine in these sections fail to reproduce the standard equations of motion of general relativity.

The dichotomy of a well-defined quantum theory arising from an ill-defined classical theory – so different from the standard situation in quantum field theory – is the main point of interest of the chapter. The fact that these quantum theories seem to be well-defined could be taken as evidence that my approach to quantization is extremely promising from the perspective of producing a quantum theory of gravity. Indeed, I will produce three different quantum theories of gravity in section 3.2, with many more within easy reach. But it is not at all clear that the choices I explore in this chapter – classical or quantum – are the right ones, and different choices lead to different results. *Caveat emptor.*

3.2 Quantum general relativity made easy

The first and most important question we must answer in order to apply the methods of chapter 2 to general relativity is: what fiber bundle E shall we begin with? It seems straightforward that we should start with the fiber bundle

$$E = T^*M \otimes T^*M \tag{3.1}$$

whose symmetric sections are possible metrics g on our space-time M. This is certainly something we can do; it leads to the fiber bundle $P = V^*E \otimes TM$ with local coordinates $\{x^\alpha, g_{\alpha\beta}, \pi^{\alpha\beta\gamma}\}$ in which a point $p \in P$ is given by

$$p \in P = x^\alpha e_\alpha + g_{\alpha\beta} dx^\alpha \otimes dx^\beta + \pi^{\alpha\beta\gamma} dg_{\alpha\beta} \otimes \frac{\partial}{\partial x^\gamma}$$

which in turn implies that the canonical tensor of eq. (2.4) has the local coordinate representation

$$\theta = \pi^{\alpha\beta\gamma} dg_{\alpha\beta} \otimes \frac{\partial}{\partial x^\gamma} \tag{3.2}$$

so that the polysymplectic structure is given locally by

$$\omega = d\pi^{\alpha\beta\gamma} \wedge dg_{\alpha\beta} \otimes \frac{\partial}{\partial x^\gamma} \tag{3.3}$$

and, more importantly, the Poisson structure is given locally by

$$\Pi = \frac{\partial}{\partial g_{\alpha\beta}} \wedge \frac{\partial}{\partial \pi^{\alpha\beta\gamma}} \otimes dx^\gamma \tag{3.4}$$

With this Poisson structure in hand, it is no more technically challenging to find the local representation of the quantum operators $Q(g_{\alpha\beta})$ and $Q(\pi^{\alpha\beta 0})$ than it was for the standard Klein-Gordon theory: we apply eq. (2.24) just as before. As a reminder, our quantization map is

$$Q(f) = f - \Pi(\theta) f + \Pi(df) \llcorner \mathrm{i}\hbar \frac{\partial}{\partial t} \tag{3.5}$$

Its application to the case at hand gives

$$Q(g_{\alpha\beta}) = g_{\alpha\beta} + \mathrm{i}\hbar \frac{\partial}{\partial \pi^{\alpha\beta 0}} \tag{3.6}$$

$$Q(\pi^{\alpha\beta 0}) = -\mathrm{i}\hbar \frac{\partial}{\partial g_{\alpha\beta}} \tag{3.7}$$

If as in chapter 2 our candidate wavefunctions are taken to be polarized functions

$$\Psi : P \to \mathbb{R} \mid v\Psi = 0 \, \forall \, v \in VP \tag{3.8}$$

then the extra derivative in eq. (3.6) does no harm and we really have

$$Q(g_{\alpha\beta}) = g_{\alpha\beta} \tag{3.9}$$

$$Q(\pi^{\alpha\beta 0}) = -\mathrm{i}\hbar \frac{\partial}{\partial g_{\alpha\beta}} \tag{3.10}$$

These operators obey the same fundamental commutation relations as the real-

valued Klein-Gordon field operators of section 2.13:

$$\left[Q(g_{\alpha\beta}), Q(\pi^{\gamma\epsilon 0})\right] = \mathrm{i}\hbar\delta^{\gamma}_{\alpha}\delta^{\epsilon}_{\beta} \tag{3.11}$$

As one would expect from the analysis of chapter 2, we obtain an integrated canonical commutation relation rather than a "raw" functional analytic one:

$$\left[Q(g_{\alpha\beta}), Q(\pi^{\gamma\epsilon 0})\right] = \int \star g(\frac{\partial}{\partial t}) \left[\hat{g}_{\alpha\beta}(\vec{x}), \hat{\pi}^{\gamma\epsilon}(\vec{y})\right] \tag{3.12}$$

where I have assumed in writing this as an equality that the operator-valued functions $\hat{g}_{\alpha\beta}(\vec{x})$ and $\hat{\pi}^{\gamma\epsilon}(\vec{y})$ obey the standard quantum field theory canonical commutation relations[2]

$$\left[\hat{g}_{\alpha\beta}(\vec{x}), \hat{\pi}^{\gamma\epsilon}(\vec{y})\right] = \delta^{\gamma}_{\alpha}\delta^{\epsilon}_{\beta}\delta(\vec{x} - \vec{y}) \tag{3.13}$$

Shockingly, we have encountered no major technical obstacle to the straightforward application of the quantization program of chapter 2 to the case of general relativity! However, the fact that we are able to do something without any major technical challenges does not mean that everything is okay; things are actually much more complicated than they seem.

Before turning to these complications, it is worth mentioning that the approach to general relativity presented above is by no means unique. Two other approaches are quite common: the Palatini approach and the ADM approach. Each of these begins with a different starting fiber bundle than that of eq. (3.1) and therefore ends with a different quantum theory. Ultimately, almost everything in my approach to polysymplectic Hamiltonian field theory rests on the geometric foundation chosen as the setting for the classical fields. Therefore, to change to something more like the Palatini perspective[3] we must change our geometric foundation. Rather than take our starting fiber bundle to be $E = T^*M \otimes T^*M$ as in eq. (3.1), let us instead

[2]This is actually a huge stretch: even at this basic level there is no such thing as a standard result in quantum gravity, so this commutation relation is by no means necessarily correct.

[3]I say "to something more like the Palatini perspective" rather than simply "to the Palatini perspective" because it will become apparent that this starting fiber bundle does not end up reproducing some of the key features of the Palatini approach. See section 3.3 for more details.

take as our starting fiber bundle

$$E = (T^*M \otimes T^*M) \oplus (TM \otimes T^*M \otimes T^*M) \tag{3.14}$$

where the $\oplus$ symbol denotes the Whitney sum of two fiber bundles (each of which is a bundle with base space M representing space-time). Intuitively, this sum means that we are taking independent combinations of sets of metric components and connection coefficients as our fundamental fields; see the Glossary for more details. A point $e \in E$ is represented by

$$e \in E = x^\alpha e_\alpha + g_{\alpha\beta} dx^\alpha \otimes dx^\beta + \Gamma^\alpha_{\ \beta\gamma} \frac{\partial}{\partial x^\alpha} \otimes dx^\beta \otimes dx^\gamma$$

This starting fiber bundle then gives $P = V^*E \otimes TM$, with a point $p \in P$ represented by

$$p \in P = x^\alpha e_\alpha + g_{\alpha\beta} dx^\alpha \otimes dx^\beta + \Gamma^\alpha_{\ \beta\gamma} \frac{\partial}{\partial x^\alpha} \otimes dx^\beta \otimes dx^\gamma + \pi^{\alpha\beta\gamma} dg_{\alpha\beta} \otimes \frac{\partial}{\partial x^\gamma} + \rho_\alpha^{\ \beta\gamma\epsilon} d\Gamma^\alpha_{\ \beta\gamma} \otimes \frac{\partial}{\partial x^\epsilon}$$

Note the two sets of canonically conjugate momenta: the $\pi^{\alpha\beta\gamma}$ for the metric components $g_{\alpha\beta}$ and the $\rho_\alpha^{\ \beta\gamma\epsilon}$ for the connection coefficients $\Gamma^\alpha_{\ \beta\gamma}$.

This phase space P in turn yields a polysymplectic structure with local canonical coordinate representation

$$\omega = \left(d\pi^{\alpha\beta\gamma} \wedge dg_{\alpha\beta} \otimes \frac{\partial}{\partial x^\gamma} \right) \oplus \left(d\rho_\alpha^{\ \beta\gamma\epsilon} \wedge d\Gamma^\alpha_{\ \beta\gamma} \otimes \frac{\partial}{\partial x^\epsilon} \right) \tag{3.15}$$

and a Poisson tensor with local canonical coordinate representation

$$\Pi = \left(\frac{\partial}{\partial g_{\alpha\beta}} \wedge \frac{\partial}{\partial \pi^{\alpha\beta\gamma}} \otimes dx^\gamma \right) \oplus \left(\frac{\partial}{\partial \Gamma^\alpha_{\ \beta\gamma}} \wedge \frac{\partial}{\partial \rho_\alpha^{\ \beta\gamma\epsilon}} \otimes dx^\epsilon \right) \tag{3.16}$$

The naive application of eq. (2.24) and the restriction to polarized wave-functions

of the form $\Psi = \Psi(x^\alpha, g_{\alpha\beta}, \Gamma^\alpha_{\ \beta\gamma})$ as required by eq. (2.28) leads directly to

$$Q(g_{\alpha\beta}) = g_{\alpha\beta} \tag{3.17}$$

$$Q(\Gamma^\alpha_{\ \beta\gamma}) = \Gamma^\alpha_{\ \beta\gamma} \tag{3.18}$$

$$Q(\pi^{\alpha\beta 0}) = -\mathrm{i}\hbar \frac{\partial}{\partial g_{\alpha\beta}} \tag{3.19}$$

$$Q(\rho_\alpha^{\ \beta\gamma 0}) = -\mathrm{i}\hbar \frac{\partial}{\partial \Gamma^\alpha_{\ \beta\gamma}} \tag{3.20}$$

giving the expected commutation relations (from the perspective of chapter 2) for all the field operators and their associated time-valued momentum counterparts:

$$\big[Q(g_{\alpha\beta}), Q(\pi^{\alpha\beta 0})\big] = \mathrm{i}\hbar = \big[Q(\Gamma^\alpha_{\ \beta\gamma}), Q(\rho_\alpha^{\ \beta\gamma 0})\big] \tag{3.21}$$

As expected, this different starting fiber bundle yields a whole extra set of field operators and conjugate momentum operators; i.e., a different quantum theory.

The ADM approach yields the opposite: fewer field operators and fewer conjugate momentum operators. To switch to something like this perspective, I will start with a rather odd looking fiber bundle. To mimic the fact that the ADM formalism singles out the spatial components of the metric as the main dynamical degrees of freedom [55], we will take our starting fiber bundle E to have base manifold M (space-time) and standard fiber $T^*_s S \otimes T^*_s S$, where the space S is a differentiable manifold that represents the spatial part of the space-time M. Symmetric sections of E then represent possible 3-metrics (spatial metrics) over M.[4]

Using the standard notation that Greek letters will denote space-time (four dimensional) indices and Latin letters will denote purely spatial (three dimensional)

[4]Just as the starting fiber of eq. (3.14) will not reproduce all the key results of the the Palatini approach, this setup will not capture all the features of the ADM formalism; for one thing, it completely ignores the lapse and shift functions $\{N, N_i\}$ that are essential to the ADM approach. However, since the point is to reduce the degrees of freedom to only those that are dynamical, this is a necessary step if we wish to use the machinery of chapter 2. The setup also assumes that the spatial portion of space-time can be represented by the same differentiable manifold S at every point in time, which need not be true in general relativity. See section 3.3 for more details.

indices, a point $e \in E$ is then given by

$$e \in E = x^{\alpha} e_{\alpha} + g_{ij} dx^{i} \otimes dx^{j} \tag{3.22}$$

which looks exactly the same as the starting setup of section 3.2 except for the Latin letter indices to indicate that we are dealing only with the spatial components of the metric. Note that the bundle is still taken over space-time, meaning that the three-metric components g_{ij} will be able to change over time as well as over space.

This starting fiber bundle leads to a phase space P in which a point $p \in P$ is represented by

$$p \in P = x^{\alpha} e_{\alpha} + g_{ij} dx^{i} \otimes dx^{j} + \pi^{ij\alpha} dg_{ij} \otimes \frac{\partial}{\partial x^{\alpha}} \tag{3.23}$$

Note that the approach of chapter 2 demands that we have four conjugate momenta for each component of g_{ij}, because the base manifold M remains space-time (rather than just the spatial manifold S).

This phase space in turn leads to a polysymplectic structure represented in local canonical coordinates by

$$\omega = d\pi^{ij\alpha} \wedge dg_{ij} \otimes \frac{\partial}{\partial x^{\alpha}} \tag{3.24}$$

and a Poisson tensor represented in local canonical coordinates by

$$\Pi = \frac{\partial}{\partial g_{ij}} \wedge \frac{\partial}{\partial \pi^{ij\alpha}} \otimes dx^{\alpha} \tag{3.25}$$

The naive application of eq. (2.24) and the restriction to polarized wave-functions of the form $\Psi = \Psi(x^{\alpha}, g_{ij})$ as required by eq. (2.28) leads directly to

$$Q(g_{ij}) = g_{ij} \tag{3.26}$$

$$Q(\pi^{ij0}) = -\mathrm{i}\hbar \frac{\partial}{\partial g_{ij}} \tag{3.27}$$

with the commutation relation

$$\left[Q(g_{ij}), Q(\pi^{kl0})\right] = \mathrm{i}\hbar\delta_i^k\delta_j^l \tag{3.28}$$

which in turn gives the integrated commutation relation

$$\left[Q(g_{ij}), Q(\pi^{ij0})\right] = \int \star g(\frac{\partial}{\partial t}) \left[\hat{g}_{ij}(\vec{x}), \hat{\pi}^{ij}(\vec{y})\right] \tag{3.29}$$

where I again assume that the operator-valued functions $\hat{g}_{ij}(\vec{x})$ and $\hat{\pi}^{ij}(\vec{y})$ obey the standard quantum field theory canonical commutation relations[5].

Now that we have three different quantum analogs of general relativity in hand, let us look at a first example of some of the problems inherent in applying this quantization method. First, let us examine equation eq. (3.5) more closely. In writing down this equation, I have naively used the vector field $\frac{\partial}{\partial t}$. This caused no great trouble in section 2.13, as we are used to working with real-valued Klein-Gordon fields on flat space-times. In these situations, there is always a global, unit-length, time-like vector field $\frac{\partial}{\partial t}$ to appeal to. But in the case of general relativity, there generally is not. So just by using eq. (3.5) we have restricted the scope of the analysis: we cannot accommodate relativistic space-times that lack a global, unit-length, time-like vector field. Let us call this *challenge one* (or **C1** for short). Moreover, we have violated one of the basic premises of classical general relativity: that we should not single out any particular coordinate direction for special treatment. Let us call this **C2**. While **C1** about a global, unit-length, time-like vector field is fundamentally technical in nature, **C2** is quite a bit deeper: the quantization procedure of eq. (2.24) *demands* that we single out the time coordinate direction for special treatment, while the fundamental principles of general relativity tell us that we should not. Note that both of these challenges are fundamentally classical in nature: the lack of a global, unit-length, time-like vector field is a feature of classical – not quan-

[5]This assumption is more likely to be justified in the present context. Indeed, if I allowed these operators to have space-time (rather than just spatial) dependence, eq. (3.29) would closely resemble the fundamental commutation relation of early quantum general relativity [4].

tum – general relativity. The admonition to preserve manifest general covariance is also a principle of classical – not quantum – general relativity.

I will return to these challenges in section 3.4. But there are actually more complicated technical challenges that I have swept under the rug in the foregoing analysis. Interestingly, they once again concern not the quantum but the classical theory of gravity.

3.3 Complications from classical general relativity

The forgoing analysis makes it look like the quantization of general relativity is actually quite easy in my approach. From a technical standpoint, this is basically true: though we made a number of choices in deriving the commutation relations of eqs. (3.12), (3.21), and (3.29), making other choices within the same basic framework would produce different results with no greater technical complication. This is not because quantizing general relativity is easy. Rather, it is because the greatest challenges in quantizing general relativity using the procedure of chapter 2 are not quantum, but classical. They pertain not to section 2.13, but to section 2.8.

Indeed, almost all of the preceding analysis is, if not wrong, then at least perilously ambiguous. Consider this seemingly simple question: what classical observable is the operator-valued field $\pi^{\alpha\beta 0}$ the quantum counterpart to? To try to answer this question, we need to define a Lagrangian on the jet bundle JE over the bundle $E = T^*M \otimes T^*M$ taken as the geometric foundation of the theory. We then need to consider the covariant Legendre transformation of section 2.8 to determine the relationship between the classical conjugate momenta $\pi^{\alpha\beta\gamma}$ and various derivatives of the metric $g_{\alpha\beta}$. To be fully confident in our approach, we should use our Legendre transformation to define a covariant Hamiltonian function and recover the Einstein field equations from eq. (2.9) and (2.10).

The problem is that this whole – purely classical – enterprise is fraught with difficulties. If we choose the standard Einstein-Hilbert Lagrangian density for general

relativity

$$\mathscr{L} = \sqrt{-\det g}R \tag{3.30}$$

we find that in the geometric framework of eq. (3.1) this Lagrangian density is not a function on the first jet bundle JE, but on the second jet bundle J^2E. Strictly speaking, this is the end of the story for us: the approach of chapter 2 works within the context of the (linearized) first jet bundle JE, not the second jet bundle J^2E. If we persist by trying to isolate the "first-order" part of eq. (3.30) we encounter another problem: the Legendre transformation of eqs. (2.18) and (2.19) is not invertible. This presents us with an insurmountable problem when attempting to define the covariant Hamiltonian function H, as there is no way to rewrite the Lagrangian density $\mathscr{L}(x^\alpha, g_{\alpha\beta}, \frac{\partial g_{\alpha\beta}}{\partial x^\gamma}, \frac{\partial^2 g_{\alpha\beta}}{\partial x^\gamma \partial x^\epsilon})$ in terms of the conjugate momenta [6]

$$\pi^{\chi\psi\omega} := \frac{\partial \mathscr{L}}{\partial \frac{\partial g_{\chi\psi}}{\partial x^\omega}} = \sqrt{-\det g}\,\frac{\partial g_{\alpha\beta}}{\partial x^\gamma}\Big(g^{\chi\psi}(g^{\omega\alpha}g^{\beta\gamma} - \frac{1}{2}g^{\omega\gamma}g^{\alpha\beta}) + g^{\chi\omega}(\frac{1}{2}g^{\psi\gamma}g^{\alpha\beta} - g^{\psi\alpha}g^{\beta\gamma}) + g^{\psi\omega}(\frac{1}{2}g^{\chi\gamma}g^{\alpha\beta} - g^{\chi\beta}g^{\alpha\gamma}) - g^{\psi\alpha}g^{\omega\beta}g^{\chi\gamma} + \frac{3}{2}g^{\chi\alpha}g^{\psi\beta}g^{\omega\gamma}\Big) \tag{3.31}$$

because the $\pi^{\chi\psi\omega}$ give no access to the second derivatives $\frac{\partial^2 g_{\alpha\beta}}{\partial x^\gamma \partial x^\epsilon}$. So it is not possible to use this Legendre transformation to define a consistent covariant Hamiltonian function with this geometric setup[7]. In addition to meaning that we will not be able to define any consistent energy operator in the quantum theory – already a serious problem for any standard quantum field theory – it also means that this procedure does not reproduce the equations of motion of general relativity: no covariant Hamiltonian means no Hamilton's equations of motion, let alone the correct ones. There is therefore no classical guarantee that the momenta $\pi^{\alpha\beta\gamma}$ I so readily

[6] This result is unlikely to be familiar even to the reader who knows the standard variational approaches to general relativity. In most treatments, it is the inverse metric that is varied – not the metric itself – and $g_{\alpha\beta}$ and $g^{\alpha\beta}$ taken to be independent fields. The partial derivative I calculated is not part of the standard variational approach. Indeed, there are even several ways to calculate this partial derivative depending upon how one enforces the relationships between $g_{\alpha\beta}$, $g^{\alpha\beta}$, and their derivatives. I present the one that seems to me the most reasonable. None of them gives access to the second derivatives of the metric.

[7] Actually, instead of proceeding with the standard Legendre transformation we could try to develop a generalized version as others have done (see, for example, [56]). Since this method is not necessarily consistent with the approach of chapter 2, I do not try to implement it here.

quantized in section 3.2 have anything to do with the physics of general relativity!

Let us call the fact that the Legendre transformation either does not exist or is not invertible **C3**, and the fact that we have no well-defined covariant Hamiltonian function that gives us the proper equations of motion of general relativity **C4** [8]. Though these are not necessarily the only problems associated with the naive process I used in section 3.2, they illustrate a very general situation. The quantization procedure of section 2.13 is no more challenging to apply to general relativity than to a real-valued Klein-Gordon field. Instead, all the new challenges I have discussed – **C1** through **C4** – come from the fact that general relativity does not fit the mold of covariant Hamiltonian field theory on the classical side. The framework of chapter 2 is very conducive to quantization, but it has serious difficulty handling general relativity. Indeed, there is no technical obstacle to quantizing the classical field theories I naively set up in the foregoing discussion to emulate general relativity. The problem is that these set ups do not emulate *all* of general relativity, so we should be skeptical that the quantum theories they lead to are physically meaningful.

Even this brief analysis has revealed two more technical challenges. These challenges seem like they might be solvable with a different geometric setup. Unfortunately, the most common alternative formulations of general relativity I explored in section 3.2 do not solve them.

Since **C3** first arose from the fact that the Einstein-Hilbert Lagrangian density of eq. (3.30) contains second-order derivatives of the metric g (in addition to first order derivatives), one natural way to try to address this challenge – and therefore, perhaps, **C4** as well – is to switch to a starting fiber bundle that emulates the jumping off point of the Palatini approach to the variational problem, in which the metric and connection coefficients are taken to be independent variables [57]. We will see that though the alternative starting fiber bundle of eq. (3.14) emulates this key feature of the Palatini approach, it does not reproduce the classical results of that approach. It still proves impossible to define a covariant Hamiltonian from

[8] Though these problems naturally arise together in the analysis of this section, they are in principle distinct issues.

the Lagrangian density of eq. (3.30), and we do not ever uncover the relationship between the metric and the connection coefficients. In the true Palatini approach this relationship is recovered from some of the Euler-Lagrange equations.

The apparent benefit of switching to the Palatini-like starting fiber bundle of eq. (3.14) is that the Einstein-Hilbert Lagrangian of eq. (3.30) is now based only on the first derivatives of these variables. The covariant Legendre transformation of eq. (2.19) goes through in this case, giving

$$\pi^{\alpha\beta\gamma} = \frac{\partial \mathscr{L}}{\partial \frac{\partial g_{\alpha\beta}}{\partial x^{\gamma}}} = 0 \tag{3.32}$$

and

$$\rho_{\alpha}^{\beta\gamma\epsilon} = \frac{\partial \mathscr{L}}{\partial \frac{\partial \Gamma^{\alpha}_{\beta\gamma}}{\partial x^{\epsilon}}} = \sqrt{-\det g}\left(g^{\beta\gamma}\delta^{\epsilon}_{\alpha} - g^{\beta\epsilon}\delta^{\gamma}_{\alpha}\right) \tag{3.33}$$

Here we must remember that the metric coefficients $g_{\alpha\beta}$ and connection coefficients $\Gamma^{\alpha}_{\beta\gamma}$ are taken to be independent in this geometric setup, and therefore also in evaluating these partial derivatives. The result is therefore quite different from eq. (3.31) because the scalar curvature R is most directly a function of the connection coefficients, rather than the space-time derivatives of the metric.

But just like eq. (3.31) the Legendre transformation of eqs. (3.32) and (3.33) is not invertible. Indeed, it fails to allow us to define a consistent covariant Hamiltonian function for the exact same reason as the Legendre transformation of eq. (3.31): it gives no access to the derivatives of either set of variables[9].

So just as before it is not possible to define a consistent covariant Hamiltonian function with this Legendre transformation, nor to define any consistent energy operator in the quantum theory. This procedure does not reproduce the equations of motion of general relativity. There is once again no obvious classical guarantee

[9]The reasons that the two Legendre transformations fail are different, though. The Legendre transformation discussed in reference to the standard Einstein-Hilbert approach fails because it contains no information about the second derivatives of the metric, which are needed to access the first derivatives of the connection coefficients that appear in eq. (3.30). The present Legendre transformation fails because the Einstein-Hilbert Lagrangian is linear (rather than quadratic) in the first derivatives of the connection coefficients and – from the Palatini perspective – no longer directly contains any first derivatives of the metric.

that the momenta $\pi^{\alpha\beta\gamma}$ and $\rho_{\alpha}^{\beta\gamma\epsilon}$ I quantized in section 3.2 have anything to do with the physics of general relativity. In other words, **C3** and **C4** remain unsolved, and we are right back where we started.

As a final attempt, one could go the opposite direction: rather than start with a fiber bundle that emulates some of the features of the Palatini approach – with many more apparent degrees of freedom than the standard approach – I could start with a fiber bundle that emulates some of the key features of the ADM approach by starting with the fiber bundle whose points are represented by eq. (3.22). This alternative instead reduces the number of apparent degrees of freedom.

Unfortunately, this approach encounters all the same issues I discussed in the build up to the covariant Legendre transformation of eq. (3.31). In fact it is somewhat worse, as we cannot even write the Einstein-Hilbert Lagrangian of eq. (3.30) as a function on J^2E, since this geometric framework does not contain any information about the temporal components of the metric. If we naively attempt to plow through, eq. (2.19) gives

$$\pi^{ij\omega} := \frac{\partial \mathscr{L}}{\partial \frac{\partial g_{ij}}{\partial x^{\omega}}} = \sqrt{-\det g}\,\frac{\partial g_{\alpha\beta}}{\partial x^{\gamma}}\Bigg(g^{ij}(g^{\omega\alpha}g^{\beta\gamma} - \frac{1}{2}g^{\omega\gamma}g^{\alpha\beta}) + g^{i\omega}(\frac{1}{2}g^{j\gamma}g^{\alpha\beta} - g^{j\alpha}g^{\beta\gamma}) + g^{j\omega}(\frac{1}{2}g^{i\gamma}g^{\alpha\beta} - g^{i\beta}g^{\alpha\gamma}) - g^{j\alpha}g^{\omega\beta}g^{i\gamma} + \frac{3}{2}g^{i\alpha}g^{j\beta}g^{\omega\gamma}\Bigg) \quad (3.34)$$

which once again lacks any information about the second derivatives of the metric $\frac{\partial^2 g_{ij}}{\partial x^{\gamma}\partial x^{\delta}}$. So it is once again impossible to rewrite the Lagrangian density of eq. (3.30) in terms of the g_{ij} and $\pi^{ij\omega}$, even neglecting the previously mentioned loss of information about the temporal components of the metric.

This result may seem surprising to a reader familiar with the standard variational approach to the ADM formalism. But there is no reason to expect eq. (3.34) to allow us to reproduce the ADM Hamiltonian. First, we cannot appeal to the lapse and shift functions of the ADM formalism in defining a covariant Hamiltonian, as they are not part of the starting fiber bundle whose points are represented by eq. (3.22). Second, the covariant Hamiltonian approach does not allow us to

appeal to derivatives of the momenta in formulating H, and the covariant Legendre transformation of eq. (2.19) is only expected to work properly for hyper-regular Lagrangians. So although it once again produces a perfectly well-defined quantum theory – indeed, one closer to the expected results of the canonical quantization of general relativity than those examined previously – this ADM-like approach does not lead to a consistent covariant Hamiltonian function any more than the previous two approaches did.

3.4 Remarks

In the last section I outlined two alternatives to the standard starting fiber bundle of eq. (3.1) that might have been expected to solve **C3** and (possibly) **C4**. In the process, I showed that solving these challenges is not necessarily as simple one would hope. But there are at least two other approaches to these challenges.

The first is to question the idea that we must start with the Lagrangian density of eq. (3.30) in order to define our covariant Hamiltonian function. As pointed out in section 2.8, it is a perfectly sensible alternative to take the covariant Hamiltonian function as the starting point, finding the relation between the momenta and the spacetime derivatives of the fields from Hamilton's field equations rather than the Legendre transformation. I have tried to stick more closely to standard formulations in the foregoing analysis, as this alternative would – to take just one important consequence – make it impossible to recover the Einstein-Hilbert action, and is therefore likely to provoke considerable skepticism if adopted. Nonetheless, it has the potential to bypass **C3** and solve **C4**, and will be explored in future work.

But perhaps a more interesting alternative is to ask the question: do these challenges really need to be solved? After all, the procedure I use gives perfectly well-defined conjugate momenta; it just doesn't produce well-defined Legendre transformations or covariant Hamiltonians.

Perhaps asking for a consistent covariant Hamiltonian is the wrong thing to do. We do not necessarily require – and may never be able to define – a covariant

Hamiltonian in a consistent quantum theory of gravity. Admittedly, any quantum theory that lacks this feature will differ greatly from those of standard quantum fields. But given the many difficulties involved in constructing a reasonable energy function even in classical general relativity, I would be hesitant to demand that a quantum counterpart to a classical Hamiltonian function should be a requirement of a quantum theory of gravity.

A similar attitude can be taken with regard to **C1** and **C2**. Rather than say that the requirement of a global, unit-length, time-like vector field is too restrictive, we might say that it is a basic requirement of a sensible quantum theory. Space-times without this feature may be classically admissible, but perhaps they simply do not work in quantum theory. After all, the space-times considered in ordinary quantum field theory certainty admit global, unit-length, time-like vector fields. Finally, we might say the same thing about the philosophical challenge that we ought not single out the time coordinate direction for special treatment. The procedure is indeed inconsistent with the philosophical foundations of general relativity. But despite this philosophical point, there is a major distinction between the time coordinate direction and the spatial ones: the square of the metric distance associated with it is negative. This in turn *allows* us to single it out for special treatment. Perhaps this is another requirement of the quantum theory that does not have a clear counterpart in classical general relativity.

Conclusion

This dissertation has been a whirlwind tour of different elements related to the quantization of general relativity: experimental input, reformulation of the classical elements of field theory to better allow for their quantization, and novel approaches to the quantization of field theories. The only obvious tie between these disparate elements is their importance to the program of quantizing general relativity.

The title of the dissertation emphasizes that these are only a handful of *considerations* in the quantization of general relativity. Despite the highly speculative analysis of chapter 3, nowhere in this work have I come anywhere close to successfully quantizing general relativity. It may be that one or more of the elements I have analyzed will play a crucial role in whatever quantization program actually succeeds in producing a successful quantum theory of gravity. The novel source-localization algorithm – if it passes a few more experimental tests – could be implemented in the real-world localization protocols of whatever future detector network produces data bearing on the domain of quantum gravity. Or future work could show that quantization procedures utilizing symplectic structures from finite-dimensional covariant Hamiltonian field theory can be made to produce accurate experimental predictions, justifying not only the quantization programs themselves but the importance of the approach to classical field theory upon which they are based. Even if this effort comes to fruition, there is still the big question of whether it can handle general relativity any better than the standard approaches.

For the moment, each of these three novel approaches seems to hold promise. It is my hope that, with a great deal of future work, each piece can make its own

small contribution to the ultimate goal of producing a successful quantum theory of gravity.

Appendix: Noise in the LIGO Detector Network

From the very beginning, noise has been the Achilles heel of interferometric gravitational wave detectors, and the LIGO network in particular. In 1999 when Barish and Weiss were describing the basic operating principles of the LIGO network to a general audience in [58], they emphasized that "controlling the noise in the measurement of exceedingly small strains...has been the prime technological challenge in this field for the past several decades, and it is the central focus of our development of the technology for LIGO." Even now – almost 20 years later – characterizing and controlling sources of noise remains the primary obstacle in improving the sensitivity of the LIGO network. Clearly, noise is an important topic in any discussion of interferometric gravitational wave detectors.

Many important attributes of the noise present in the LIGO detectors have not been definitely characterized. For example, it is generally accepted that much of the noise in the LIGO detectors is *not* Gaussian and *not* stationary [24]; however, there is no widely accepted model that characterizes what the noise *is*. With such fundamental questions unanswered, there is a great deal more research to be done in this area.

The presence of noise in detectors that are typically working very close to their minimal detection threshold is naturally cause for concern: a great deal of care must be take to verify that any candidate signal is the result of a genuine gravitational wave, rather than simple noise. It is therefore worth briefly discussing the

general principles by which LIGO and other interferometric gravitational wave detectors eliminate noise-only triggers from their analyses. This process consists of four stages: coincidence analysis, time-delay analysis, GW response analysis, and source-specific analysis. Each of these analyses basically asks that the candidate signal meet increasingly complicated requirements that allow us to rule out the possibility that the signal is really just noise.

The first stage of the process is to perform a basic coincidence analysis. This means making sure that any candidate signal present in the data stream of one detector is also present in the other detectors in the network. In addition, since the data stream of the detectors is typically frequency separated by means of Fourier analysis, we can ask that the candidate signal is present at the same frequency in each detector. This first stage rules out the vast majority of candidate signals, as noise is typically not correlated between the various detectors in the network. Indeed, as the number of detectors in the network grows, so does the ability of this basic test to rule out candidate signals arising from non-coincident detector noise.

The second stage of the process is to check the time delay between the coincident signals in the detector network. If the time delay is between zero and 10 ms in the LIGO network, then the signal passes this test. A time delay of zero would indicate a gravitational wave propagating perpendicular to the line connecting the detectors, while a time delay of 10ms would indicate a gravitational wave propagating along the line connecting the detectors (gravitational waves propagate at the speed of light regardless of the material medium between the detectors). Any time delay between these extremes is physically realizable; any time delay outside these extremes is not, and represents a false signal. This analysis may exclude close coincidences between different signals.

The third stage of the process is to check the gravitational wave response of the detectors in more detail. In other words, we need to verify that there is an actual gravitational wave of the form of eq. (G.5) (see the Glossary) that would produce the detector response signal of eq. (1.3) seen in the data stream – up to instrumental

noise – given the known beam pattern response functions F_+ and $F_\times$ of the two LIGO detectors and the known time delay between the detectors. See section 1.3 for more details.

The fourth and final stage of the process is the most delicate and the most error prone. It amounts to checking the signal not against the basic requirements of general relativity (that is stage three), but against the very specific source-model possibilities known to numerical general relativity and our current understanding of gravitational wave sources. This basically means checking the signal against an enormous database of template signals based on source-specific modeling from numerical general relativity. A match gives us even more confidence that the signal is real. In this case, however, a non-match does not necessarily indicate that the signal is spurious; it may simply be that the source of the signal is not currently known to gravitational wave astronomers from numerical general relativity. That being the case, the final stage of this process is much more of an art than the other stages, as well as being a less reliable way of ruling out signals.

This noise-only elimination procedure is important to understand, but stages three and four clearly rely on a strong understanding of the noise present in the detectors in order to decide whether a candidate wave-form is "good enough" or not. Indeed, the whole analysis can be sharpened substantially if it is possible to characterize the noise in the detectors precisely.

As already discussed, there is no universally accepted theoretical model for the noise present in the LIGO detectors. However, there are a number of widely accepted noise sources that are known to be present in the detectors [27]. At frequencies below 10 Hz, the total noise consists mostly of seismic noise (the dominant contribution), as well as test mass and suspension thermal noise and gravity gradients. Despite the fact that these noise sources are largely counteracted by active damping protocols, the total noise in this region is at least two orders of magnitude greater than in the most sensitive region; see [59] for more details. In the most sensitive region between 100 Hz and 1000 Hz, quantum noise (radiation pressure due to fluctuations in photon

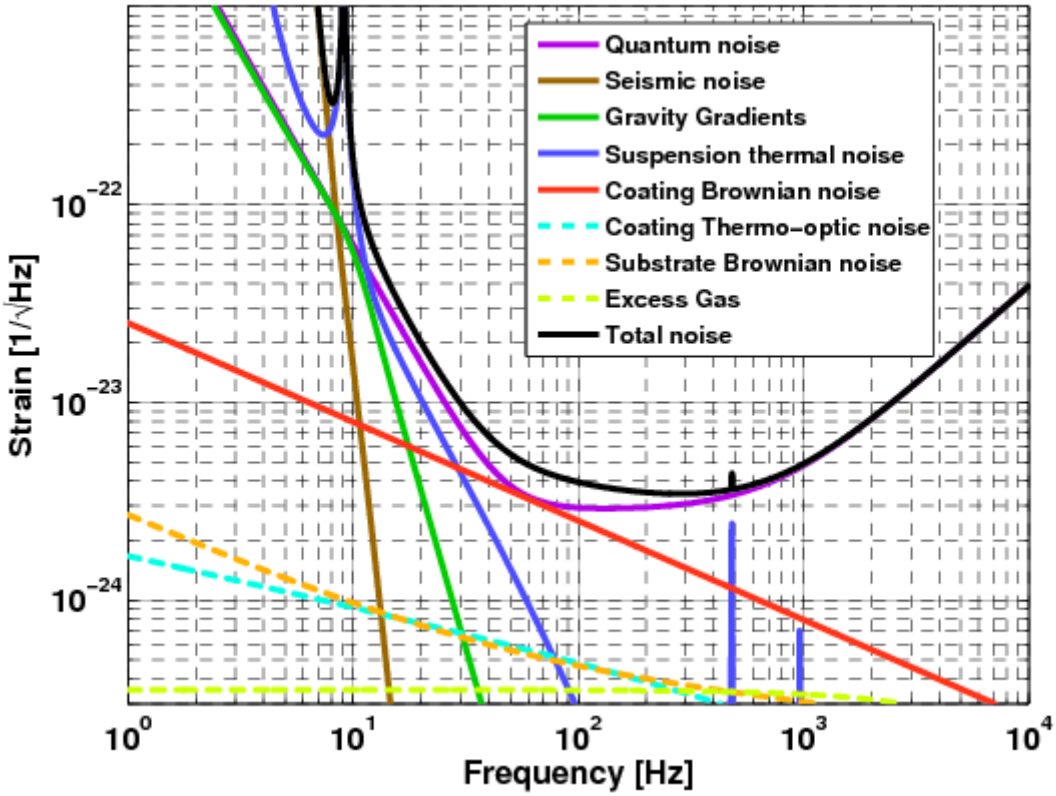

Figure A.1: This graph shows predicted values of the square root of the integrated noise present in each 1 Hz frequency bin in the LIGO detectors for the well-characterized noise sources listed in the key at around the time of the first observing run (O1). The figure is reproduced from [61].

number) and – to a lesser extent – thermal noise (mostly Brownian motion in the test mass coating) dominate the total noise. Nonetheless, the total noise in this region is at least an order of magnitude smaller than the total noise outside this region, and will further improve when finely tuned quantum squeezing is implemented (see [60]). Shot noise dominates the total noise at frequencies above 1000 Hz. Present throughout is a tiny contribution from variation in the dielectric constant of the matter in the beam path; this is all but eliminated by the almost complete evacuation of the beam tubes. See fig. A.1 for a graphical representation of these noise sources as estimated around the time of the first observing run (O1) of the LIGO network in its "Advanced" phase, September 12, 2015 to January 19, 2016.

In addition to these slowly varying noise sources, there are a few other known sources that contribute importantly to the overall noise. Noise from electrical os-

cillations (e.g., from power lines) occurs at integer multiples of 60 Hz, producing well characterized but significant noise spikes at regular intervals. Violin noise – the interaction of suspension thermal noise with the natural standing wave frequencies of the silicon fiber suspension rods for the interferometer mirrors – creates similarly sharp peaks at regular intervals; see [62] for more details. Both electrical oscillation noise and violin noise are well characterized and can be monitored and subtracted from the LIGO output, but glitches – uncharacterized noise spikes sometimes, but not always, caused by sudden shifts along material defects in the various mirror and test mass parts – are much more chaotic, producing sudden and unpredictable noise that is very challenging to eliminate from the data stream [63].

A great deal of work has been done to characterize the noise in the LIGO detectors empirically; some sample results – again characterizing the noise at around the time of the Advanced LIGO first observing run (O1) – are shown in fig. A.2. Comparison of figs. A.1 and A.2 shows that there is much to be gained from further analysis. Though it is not possible to characterize every potential noise source, work in this area is ongoing, with the result that more and more noise sources are characterized to higher degrees of precision and accuracy all the time.

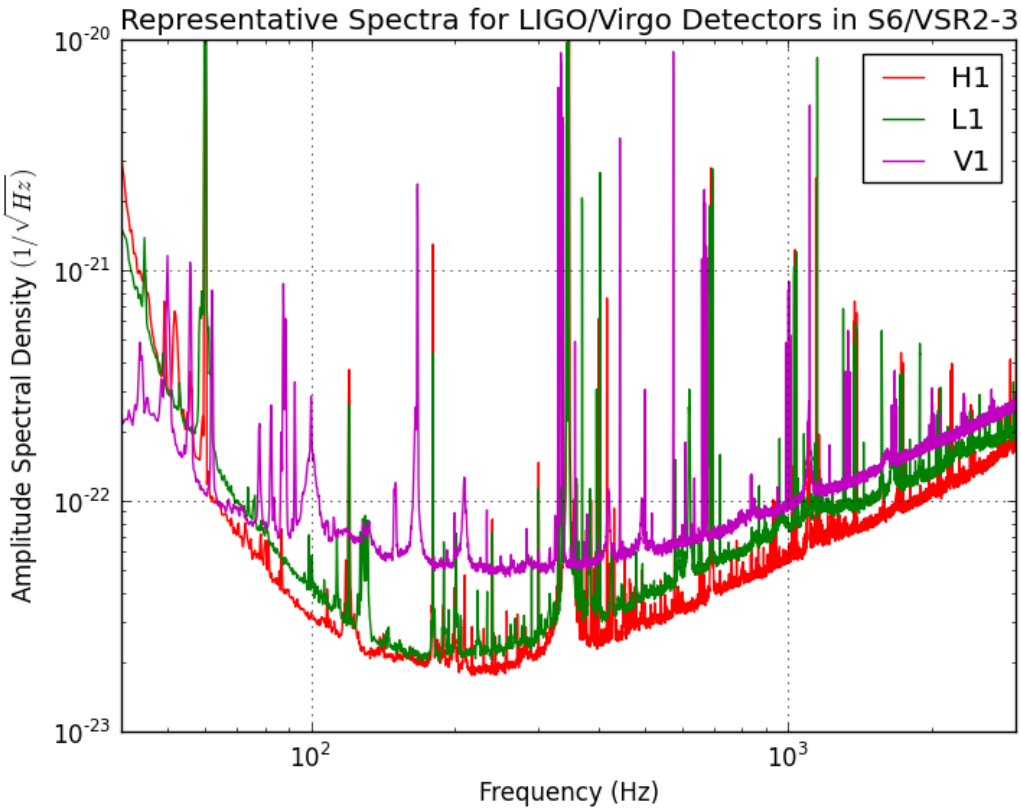

Figure A.2: This graph shows the square root of the integrated noise present in each 1 Hz frequency bin in the LIGO and Virgo detectors. Note the many sharp peaks present in addition to the smoother general trends of fig. A.1, as well as the difference in the scaling of both axes. The figure is reproduced from [64]. The figure shows the noise present in the detectors at representative times just prior to the first observing run (O1). Not included in this graph are the results of a few substantive upgrades made for the most recent observation runs.

Glossary

Axiomatic Quantum Field Theory

An axiomatic quantum field theory is one which attempts to describe quantum field theories in terms of the rigorous set of mathematical axioms that they should obey. The first example of these field theories is the Wightman axioms, which attempt to describe quantum field theories on Minkowski space-times as operator-valued distributions that act on a Hilbert space of quantum states. Moreover, the Haag-Kastler axioms underlie the first modern approach to algebraic quantum field theory, effectively making at least some algebraic quantum field theories examples of axiomatic quantum field theory.

Algebraic Quantum Field Theory

Algebraic quantum field theory is perhaps the most successful of the axiomatic quantum field theories, having produced a few important results on the expected properties of quantum field theories over space-times with large curvature (e.g., close to black holes). Technically, it consists of the set M of all open subsets of Minkowski space and a map (really a covariant functor) $\mathscr{F}$ from this set to a unital C^* algebra A such that every morphism (inclusion map) in M maps to an injective homomorphism in A. As the details of how this mathematical structure allows us to deal with quantum field theory algebraically (even on curved space-times) will not feature this work, I will not elaborate on them here.

Bayesian Inference

Generally, Bayesian inference is a statistical approach to reasoning quantitively about certainty and belief, in which Bayes theorem plays a central role:

$$P(H \mid E) = \frac{P(E \mid H)\, P(H)}{P(E)} \tag{G.1}$$

Here, $P(H \mid E)$ refers to the probability to find that some hypothesis is true given the evidence E, $P(E \mid H)$ refers to the probability that one would find all the evidence E if the hypothesis H were known to be true, and $P(H)$ and $P(E)$ refer to the raw probabilities of finding the evidence E and the hypothesis H, respectively. Bayesian inference relies greatly on having ideas about the raw probabilities of observing particular events (the evidence E and the hypothesis H); this means that the process is heavily reliant on external theoretical input. Bayesian inference is applied to gravitational wave astronomy by taking the hypothesis H to be that a particular detected gravitational wave matches a particular wave-form in a very large bank of template wave-forms, and the evidence E is the known response of the detector network. Typically, the Bayesian inference algorithms used in gravitational wave detection and analysis are adaptive, meaning that once they find a hypothesis H that has a higher-than-average probability of matching the evidence (that is, a large value of $P(H \mid E)$), the search is then narrowed to focus on template wave-forms that closely match the form of the template wave-form in H. These algorithms are both very successful and very prone to over-reliance on their wave-form templates and premature narrowing of their searches.

Canonical Coordinates

In Hamiltonian particle theory, a coordinate system on T^*M is called canonical if the symplectic form ω has the local coordinate representation

$$\omega = dp_i \wedge dq^i$$

In my approach to polysymplectic Hamiltonian field theory, a coordinate system on the phase space $P := V^*E \otimes TM$ is called canonical if the polysymplectic structure ω has the local coordinate representation

$$\omega = d\pi_I^\alpha \wedge \phi^I \otimes \frac{\partial}{\partial x^\alpha}$$

Circularly Polarized Gravitational Wave

In the tranverse, trace-free gauge of linearized general relativity (see below) a circularly polarized gravitational wave has the form

$$h = A\cos(\omega t + \phi)e_+ + A\sin(\omega t + \phi)e_\times \tag{G.2}$$

That is, the two (generically independent) polarizations e_+ and $e_\times$ are in this case exactly $\frac{\pi}{2}$ rad out of phase with one another; the only degrees of freedom are the amplitude A and the phase ϕ. In the context of gravitational wave searches, this implies that searching for circularly polarized gravitational waves means exploring a parameter space with two fewer degrees of freedom than the general (elliptical) case discussed below.

Cumulative Probability Distribution Function

In statistical analysis, a cumulative probability distribution function CPDF(x) of a random variable y gives the probability that the variable y has value less than or equal to x. That is,

$$\text{CPDF}(x) := \frac{\text{Number of samples of } y \text{ with } y \leq x}{\text{Total number of samples of } y} \tag{G.3}$$

In particular, we always have $\text{CPDF}^{-1}(0.5) = \text{Median}(y)$.

Deformation Quantization

Deformation quantization is a mathematical procedure for taking an ordinary algebra of functions F over a Poisson manifold M with Poisson structure Π and deforming it to a $\star$ algebra (operator algebra) over the same Poisson manifold with (first purely symbolic) deformation parameter $\hbar$. It proceeds by noting that the ordinary algebraic operator $*$ on F is simple given by

$$* : F \times F \to \mathbb{R} \mid (f, g) \mapsto fg$$

This operator is then deformed to a $\star$ operator by setting

$$\star : F \times F \to \mathbb{R} \mid (f, g) \mapsto fg + \sum_{k=1}^{\infty} \hbar^k B_k(f \otimes g) \tag{G.4}$$

where each of the B_k are bi-differential operators that depend on the Poisson structure Π of the Poisson manifold M. In the limit $\hbar \to 0$, the $\star$ operator returns to the ordinary product of functions; hence the name deformation quantization. The specific procedure for determining the bi-differential operators B_k from the Poisson structure Π (and, in particular, determining the appropriate weight to give each term) is the primary challenge of deformation quantization, and indeed there does not seem to be any well-known procedure that recovers the standard physical results in all cases of interest.

Elliptically Polarized Gravitational Wave

In the tranverse, trace-free gauge of linearized general relativity (see below) an elliptically polarized gravitational wave has the form

$$h = \{A\cos(\omega t) + B\sin(\omega t)\}e_+ + \{C\cos(\omega t) + D\sin(\omega t)\}e_\times \tag{G.5}$$

It therefore has four degrees of freedom, the four independent amplitudes A, B, C, and D. In the context of gravitational wave searches, this means that the generic

elliptical search must deal with a parameter space with two additional degrees of freedom beyond the circularly polarized case.

Fiber Bundle

A fiber bundle consists of three topological spaces, the base space M with dimension n, the fiber space F with dimension N, and the total space E with dimension $n + N$, along with a map $\pi : E \to M$ mapping points in the total space to points in the base space such that the map π is locally trivial. Local triviality means that for every point $e \in E$, there exists a neighborhood $U \subset M$ around $\pi(e) \in M$ such that there is a homeomorphism (a continuous map with a continuous inverse) $\phi : \pi^{-1}(U) \to U \times F$. Intuitively, this means that the total space E is locally just the cartesian product $M \times F$. A fiber bundle is called *trivial* if the local triviality condition is met globally; in this case, we just have $E = M \times F$ and the projection π is just projection onto the first factor in the cartesian product.

The notion of a fiber bundle specializes to many different cases in which the spaces M, F, and E have more structure than simple topological spaces. The primary case of interest to us will the case in which M is a differentiable manifold and F is a vector space; in this case, we call the fiber bundle a vector bundle to emphasize its special features.

Fibered Coordinates

Given a fiber bundle $\pi : E \to M$, we say a set of coordinates $\{x^\alpha, \phi^I\}$ are fibered coordinates on an open set in $U \subset E$ if there exist a set of coordinates $\{y^\alpha\}$ on an open set in $V \subset M$ such that $x^\alpha = y^\alpha \circ \pi$ on the intersection $\pi(U) \cap V$. Intuitively, fibered coordinates on a fiber bundle are those which do not mix the coordinates of the fiber with those of the base.

Geometric Quantization

Initially, the term geometric quantization was used to describe many different efforts to use the structures of differential (and, especially, symplectic) geometry to carry out the process of quantization: turning the observables of a classical system – typically real-valued functions over a phase space P with symplectic form ω – into a space of quantum operators – typically linear operators on some Hilbert space H. A physicist would naturally like to also know how to characterize this Hilbert space H on which the operators will act and make certain that it is physically appropriate. Since the first part of the quantization process is so challenging, several approaches to quantization (e.g., deformation quantization) omit this aspect of the analysis all together. The fact that geometric quantization grapples directly with the need to explicitly construct a space of quantum states is one of its great strengths in comparison to other quantization programs.

If one were to use the original definition of geometric quantization (i.e., quantization using symplectic structures), then the quantization program discussed in chapter 2 would certainly classify as a method of geometric quantization. However, the term today means something much more specific: typically, geometric quantization is taken to mean very specifically a quantization process that goes through three clearly defined steps:

1. Prequantization, in which the phase space P is replaced by a complex line bundle L over P and a connection ∇ on L is found such that $d_\nabla \propto \omega$.

2. Polarization, in which a specific representation for the operators is found by picking (arbitrarily) an appropriate sub-space of the tangent space TP (the"Polarization" entry for more details).

3. Metaplectic correction, in which the process is altered to sometimes reproduce standard results (like the energies of harmonic oscillators) by introducing the space of "half-forms" $\sqrt{\Lambda}$ and requiring the quantum operators produced by the procedure to act on the space $L \otimes \sqrt{\Lambda}$. This is the most technically of the

three stages, and the details will not enter into the analysis of this dissertation.

From this more specialized contemporary view, my quantization program is *not* geometric quantization.

Hyper-regular Lagrangian

In Lagrangian dynamics, a Lagrangian $L : JE \to \mathbb{R}$ on the jet bundle is called hyper-regular if its associated Legendre transformation

$$\mathscr{F}L : JE \to J^*E \tag{G.6}$$

is a diffeomorphism. Intuitively, since a diffeomorphism is an infinitely differentiable bijective map with an infinitely differentiable inverse, this means that the Legendre transformation carries all the information encoded in the Lagrangian over to covariant Hamiltonian formalism.

Jet Bundle

Given a fiber bundle with total space E, base space M, and projection map $\pi : E \to M$ with coordinates $\{x^\alpha, \phi^I\}$, its jet bundle (really its *first* jet bundle) J^1E is the space of all sections of E, modulo the equivalence relation $\sim_1$, where we say that two sections γ_1 and γ_2 are equivalent to first order (that is, under the equivalence relation $\sim_1$) at the point $x \in M$ if

$$\begin{aligned} \gamma_1^I(x) =& \gamma_2^I(x) \\ \frac{\partial \gamma_1^I}{\partial x^\alpha}(x) =& \frac{\partial \gamma_2^I}{\partial x^\alpha}(x) \end{aligned} \tag{G.7}$$

where we use the notation $\gamma^I(x) := \phi^I \circ \gamma(x)$. Intuitively, the jet bundle characterizes sections by looking at their values and derivatives over every point; coordinates on the jet bundle are given by $\{x^\alpha, \phi^I, \frac{\partial \phi^I}{\partial x^\alpha}\}$.

All of the ordinary results of Lagrangian field theory can be reformulated in terms of the jet bundle (usually, just the first jet bundle defined above). In these

formulations, the space of solution sections (infinite dimensional when considered as a single space in which two sections are different if they disagree at even a single point) is replaced by the finite-dimensional manifold J^1E. See, for example, [38] for a good overview.

K-Symplectic Manifold

Intuitively, a k-symplectic manifold is a differentiable manifold endowed with a k-tuple of symplectic forms, each of which is paired in a particular way with a subset of a tangent bundle so as to make the manifold "piece-by-piece" symplectic.

More precisely, a k-symplectic manifold consists of a manifold M with dimension $nk+n$, a set of k two-forms ω_i (where $0 \leq i \leq k$), and an nk dimensional distribution V (see the entry "Polarization" below) that satisfy the following properties:

1. $d\omega_i = 0$ for each i
2. $\bigcap_{i=1}^{k} \ker \omega_i = 0$
3. $\omega(u, v) = 0 \; \forall \, u, v \in V$

It is possible to build up the entire edifice of polysymplectic covariant Hamiltonian field theory using k-symplectic structures rather than following the approach of chapter 2 or other, more standard approaches. See [36] for more details.

Markov Chain Monte Carlo

Monte Carlo methods are a broad class of computational algorithms that rely heavily on random sampling to achieve their results. A Markov chain is a random process that possesses the Markov property, meaning that the probability of an event depends only upon the current step of the process, and not on any earlier or later states. Frequently, a Markov Chain Monte Carlo process is used to randomly sample the parameter space in the calculation of a high-dimension integral, evaluate how much that particular sample of the parameter space will contribute to the interval, and then sample nearby areas of the parameter space by calculating their approx-

imate contributions before settling on whether or not to jump to that point next. In this way, the method seeks out areas that will give large contributions to the integral and samples them quite densely in comparison to areas that will give small contributions, enhancing the overall accuracy of the calculation for a fixed number of sample points.

Markov Chain Monte Carlo processes are applied very frequently in the LALInference approach to determining the intrinsic parameters of gravitational wave sources: the parameter space of the wave-form template is sampled using precisely a Markov Chain Monte Carlo process, hopefully sampling densely those wave-forms that match the incoming signal relatively closely while mostly avoiding wave-forms that do not.

Poisson Bi-Vector

In ordinary Hamiltonian particle mechanics, a Poisson bi-vector is a tensor Π that yields the Poisson bracket:

$$\Pi(f,g) = \{f,g\} = \frac{\partial f}{\partial x^i}\frac{\partial g}{\partial p_i} - \frac{\partial f}{\partial p_i}\frac{\partial g}{\partial x^i} \tag{G.8}$$

where the final equality holds in canonical coordinates.

In some approaches to covariant Hamiltonian field theory (including mine), the generalization of the Poisson bi-vector is a Poisson tensor Π which generalizes the Poisson bracket to a one-form over space-time. For any two functions f and g on the phase space P, we have (in canonical coordinates):

$$\Pi(f,g) = (\frac{\partial f}{\partial \phi^I}\frac{\partial g}{\partial \pi^\alpha_I} - \frac{\partial f}{\partial \pi^\alpha_I}\frac{\partial g}{\partial \phi^I})dx^\alpha \tag{G.9}$$

The Poisson tensor Π plays the lead role in my approach to quantization with symplectic structures.

Polarization

In geometric quantization, a polarization (in mathematics more commonly called a distribution) on a manifold M is a choice of subspace P_x of the tangent bundle TM at every point $x \in M$. This subspace must have constant dimension

$$\dim P_x = \text{constant} \, \forall \, x \in M$$

and the subspace should form an integrable distribution, meaning that

$$[u, v] \in P \, \forall \, u, v \in P$$

A Hilbert space H (usually just the space $C^\infty(M)$ of all smooth functions over M) can then be pared down by requiring that the functions be covariantly constant in the direction of the polarization:

$$Q \subset H = \{f \in C^\infty(M) \mid v(f) = 0 \, \forall \, v \in P\} \tag{G.10}$$

Intuitively, the original Hilbert space is much too large, containing functions that depend simultaneously on the position and the momentum coordinates of the phase space M. A polarization is a geometrically invariant way of choosing a consistent representation, for example, the position representation or the momentum representation.

Polysymplectic Hamiltonian Field Theory

Within the context of covariant Hamiltonian field theory, a Hamiltonian field theory is called polysymplectic if the main structure used to derive Hamilton's field equations and to define canonical coordinate systems is a vector-valued analog of the ordinary symplectic structure of Hamiltonian particle theory (see below).

Symplectic Structure

In Hamiltonian particle physics, the symplectic structure (or just symplectic form) ω is the exterior derivative of the canonical one-form θ:

$$\omega = d\theta = dp_i \wedge dq^i \tag{G.11}$$

where the second equality holds in canonical coordinates, and there are differences of convention with respect to whether this form or its negative is called the symplectic form.

In covariant Hamiltonian field theory, some theories (called polysymplectic, see above) treat a multi-valued analog of this two form as the main structure used to derive the equations of motion and define canonical coordinate systems. In my approach, this polysymplectic structure can be written – in canonical coordinates, and with its domain restricted to vertical vectors (see below) – as:

$$\omega = d\pi^\alpha_I \wedge d\phi^I \otimes \frac{\partial}{\partial x^\alpha} \tag{G.12}$$

These similarities in role and coordinate representation justify the name polysymplectic.

Transverse, Trace-Free Gauge

Gravitational wave analysis is based on the weak gravitational field (linearized) equations of general relativity, as detected gravitational wave signals are always very weak. This process begins by assuming that our space-time is approximately flat, in the sense that

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu} \tag{G.13}$$

where the magnitude of the "non-flat" components $|h_{\mu\nu}| \ll 1$. For technical reasons

[10], we define a new scalar

$$h := \eta^{\mu\nu} h_{\mu\nu} \tag{G.14}$$

and a modified set of "non-flat" components

$$\bar{h}_{\mu\nu} := h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu} h \tag{G.15}$$

Now we begin to use our tremendous freedom to choose coordinates, to impose the condition [11]

$$\frac{\partial \bar{h}^{\mu\alpha}}{\partial x^{\alpha}} = 0 \tag{G.16}$$

so that the equations of motion reduce to

$$\eta^{\alpha\beta} \frac{\partial^2 \bar{h}_{\mu\nu}}{\partial x^{\alpha} \partial x^{\beta}} = 0 \tag{G.17}$$

The simplest solutions to these linearized equations is the monochromatic plane wave solution

$$\bar{h}_{\mu\nu} = \mathscr{R}\left[A_{\mu\nu} \exp(i k_{\alpha} x^{\alpha})\right] \tag{G.18}$$

where the amplitudes $A_{\mu\nu}$ and wave-vector k_{α} satisfy

$$\eta^{\nu\beta} A_{\mu\nu} k_{\beta} = 0 = \eta^{\alpha\beta} k_{\alpha} k_{\beta} \tag{G.19}$$

Amazingly, we have yet more gauge freedom to exploit (see, for example, [57]), so we choose a fixed four velocity u^{α} and also set

$$A_{\mu\nu} u^{\nu} = 0 = \eta^{\mu\nu} A_{\mu\nu} \tag{G.20}$$

Now, we choose a purely time-like u (that is, $u^0 = 1$, $u^i = 0$). In this gauge, there

[10] Mostly this is to keep the number of terms in the Einstein equations from blowing up too badly.

[11] Indices on $h_{\mu\nu}$ and $\bar{h}_{\mu\nu}$ are raised and lowered using the Minkowski metric $\eta_{\mu\nu}$

is no distinction between h and $\bar{h}$, and we have

$$h_{\mu 0} = 0 = \eta^{ij} h_{ij} \tag{G.21}$$

as well as

$$\eta^{kj} \frac{\partial h_{ij}}{\partial x^k} \tag{G.22}$$

In particular, this means that in this transverse, trace-free gauge we have

$$h_{xx} = -h_{yy} =: h_+ \tag{G.23}$$

and

$$h_{xy} = h_{yx} =: h_\times \tag{G.24}$$

(hence the name of the gauge). This makes clear that gravitational waves have, once fully gauge-fixed, only two degrees of freedom: h_+ and $h_\times$.

Vertical Vector

Given a fiber bundle $\pi : E \to M$, a tangent vector $v \in T_e E$ is called a vertical vector if $T\pi(v) = 0$. Here, $T\pi$ is the differential of the projection map π. Intuitively, a vertical tangent vector is tangent to the fiber, rather than the base space. The set of all vertical vectors is called the vertical bundle, and is usually denoted VE:

$$VE := \{v \in TE \mid T\pi(v) = 0\} \tag{G.25}$$

Whitney Sum

Given two vector bundles E and F over the same base space M, the Whitney sum $E \oplus F$ is another vector bundle over M with each fiber defined by

$$(E \oplus F)_m := E_m \times F_m \tag{G.26}$$

where the $\times$ in this equation is the ordinary Cartesian product. Intuitively, the

fibers of $E \oplus F$ are produced by taking all possible pairings between the elements of each fiber of E and the elements of the fiber of F over the same point in the base space.

Bibliography

[1] Joel Scherk and John H Schwarz. Dual models for non-hadrons. *Nuclear Physics B*, 81(1):118–144, 1974.

[2] Steven Weinberg. Photons and gravitons in perturbation theory: Derivation of maxwell's and einstein's equations. *Physical Review*, 138(4B):B988, 1965.

[3] Carlo Rovelli. Loop quantum gravity. *Living reviews in relativity*, 11(1):5, 2008.

[4] Bryce S DeWitt. Quantum theory of gravity. i. the canonical theory. *Physical Review*, 160(5):1113, 1967.

[5] Stanley Deser and Peter van Nieuwenhuizen. One-loop divergences of quantized einstein-maxwell fields. *Physical Review D*, 10(2):401, 1974.

[6] Gerald B Folland. *Quantum Field Theory: A tourist guide for mathematicians*. Number 149. American Mathematical Soc., 2008.

[7] Freeman J Dyson. Divergence of perturbation theory in quantum electrodynamics. *Physical Review*, 85(4):631, 1952.

[8] Arthur S Wightman. Quantum field theory in terms of vacuum expectation values. *Physical Review*, 101(2):860, 1956.

[9] Rudolf Haag and Daniel Kastler. An algebraic approach to quantum field theory. *Journal of Mathematical Physics*, 5(7):848–861, 1964.

[10] Kenneth G Wilson. Confinement of quarks. *Physical Review D*, 10(8):2445, 1974.

[11] Tom McClain. A global version of günther's polysymplectic formalism and the geometric analog of the canonical transformation approach to covariant hamiltonian field theory. *Journal of Mathematical Physics (under review)*, 2018.

[12] Karel Kuchař. Canonical quantization of cylindrical gravitational waves. *Physical Review D*, 4(4):955, 1971.

[13] Tom McClain. Using weighting algorithms to refine source direction determinations in all-sky gravitational wave burst searches with two-detector networks. *The European Physical Journal Plus*, 133(6):210, 2018.

[14] Tom McClain. Using weighting algorithms to refine source direction determinations in all-sky gravitational wave burst searches with two-detector networks ii: The case of elliptical polarization. *The European Physical Journal Plus (under review)*, 2018.

[15] Yekta Gürsel and Massimo Tinto. Near optimal solution to the inverse problem for gravitational-wave bursts. *Physical Review D*, 40(12):3884, 1989.

[16] Malik Rakhmanov. Rank deficiency and tikhonov regularization in the inverse problem for gravitational-wave bursts. *Classical and Quantum Gravity*, 23(19):S673, 2006.

[17] J Aasi, J Abadie, BP Abbott, R Abbott, TD Abbott, M Abernathy, T Accadia, F Acernese, C Adams, T Adams, et al. Prospects for observing and localizing gravitational-wave transients with advanced ligo and advanced virgo. *Living Reviews in Relativity*, 19, 2016.

[18] John Veitch, Vivien Raymond, Benjamin Farr, W Farr, Philip Graff, Salvatore Vitale, Ben Aylott, Kent Blackburn, Nelson Christensen, Michael Coughlin, et al. Parameter estimation for compact binaries with ground-based gravitational-wave observations using the lalinference software library. *Physical Review D*, 91(4):042003, 2015.

[19] Leo P Singer and Larry R Price. Rapid bayesian position reconstruction for gravitational-wave transients. *Physical Review D*, 93(2):024013, 2016.

[20] Ryan Lynch, Salvatore Vitale, Reed Essick, Erik Katsavounidis, and Florent Robinet. Information-theoretic approach to the gravitational-wave burst detection problem. *Physical Review D*, 95(10):104046, 2017.

[21] S Klimenko, G Vedovato, M Drago, F Salemi, V Tiwari, GA Prodi, C Lazzaro, K Ackley, S Tiwari, CF Da Silva, et al. Method for detection and reconstruction of gravitational wave transients with networks of advanced detectors. *Physical Review D*, 93(4):042004, 2016.

[22] BP Abbott, R Abbott, R Adhikari, P Ajith, B Allen, G Allen, RS Amin, SB Anderson, WG Anderson, MA Arain, et al. Search for gravitational-wave bursts in the first year of the fifth ligo science run. *Physical Review D*, 80(10):102001, 2009.

[23] Bernard F Schutz. Networks of gravitational wave detectors and three figures of merit. *Classical and Quantum Gravity*, 28(12):125023, 2011.

[24] Benjamin P Abbott, R Abbott, TD Abbott, MR Abernathy, F Acernese, K Ackley, M Adamo, C Adams, T Adams, P Addesso, et al. Characterization of transient noise in advanced ligo relevant to gravitational wave signal gw150914. *Classical and Quantum Gravity*, 33(13):134001, 2016.

[25] Warren G Anderson, John T Whelan, Patrick R Brady, Jolien DE Creighton, David Chin, and Keith Riles. Beam pattern response functions and times of arrival for earthbound interferometers. *LIGO Internal Document*, 10110, 2001.

[26] Bangalore Suryanarayana Sathyaprakash and Bernard F Schutz. Physics, astrophysics and cosmology with gravitational waves. *Living Reviews in Relativity*, 12(1):2, 2009.

[27] Junaid Aasi, BP Abbott, Richard Abbott, Thomas Abbott, MR Abernathy, Kendall Ackley, Carl Adams, Thomas Adams, Paolo Addesso, RX Adhikari, et al. Advanced ligo. *Classical and Quantum Gravity*, 32(7):074001, 2015.

[28] Benjamin P Abbott, R Abbott, TD Abbott, F Acernese, K Ackley, C Adams, T Adams, P Addesso, RX Adhikari, VB Adya, et al. Gw170814: A three-detector observation of gravitational waves from a binary black hole coalescence. *Physical review letters*, 119(14):141101, 2017.

[29] Benjamin P Abbott, R Abbott, TD Abbott, F Acernese, K Ackley, C Adams, T Adams, P Addesso, RX Adhikari, VB Adya, et al. Gw170817: observation of gravitational waves from a binary neutron star inspiral. *Physical Review Letters*, 119(16):161101, 2017.

[30] Sergey Klimenko, I Yakushin, A Mercer, and Guenakh Mitselmakher. A coherent method for detection of gravitational wave bursts. *Classical and Quantum Gravity*, 25(11):114029, 2008.

[31] S Klimenko, I Yakushin, M Rakhmanov, and G Mitselmakher. Performance of the waveburst algorithm on ligo data. *Classical and Quantum Gravity*, 21(20):S1685, 2004.

[32] Paul Dedecker. On the generalization of symplectic geometry to multiple integrals in the calculus of variations. In *Differential geometrical methods in mathematical physics*, pages 395–456. Springer, 1977.

[33] Hubert Goldschmidt and Shlomo Sternberg. The hamilton-cartan formalism in the calculus of variations. *Ann. Inst. Fourier*, 23(1):203–267, 1973.

[34] Christian Günther et al. The polysymplectic hamiltonian formalism in field theory and calculus of variations. i. the local case. *Journal of differential geometry*, 25(1):23–53, 1987.

[35] José F Cariñena, Mike Crampin, and Luis A Ibort. On the multisymplectic formalism for first order field theories. *Differential geometry and its Applications*, 1(4):345–374, 1991.

[36] Florian Munteanu, Angel M Rey, and Modesto Salgado. The günther?s formalism in classical field theory: momentum map and reduction. *Journal of Mathematical Physics*, 45(5):1730–1751, 2004.

[37] G Sardanashvily. *Generalized Hamiltonian formalism for field theory: constraint systems*. World Scientific, 1995.

[38] Mark J Gotay, James Isenberg, Jerrold E Marsden, and Richard Montgomery. Momentum maps and classical relativistic fields. part i: Covariant field theory. *arXiv preprint physics/9801019*, 1998.

[39] Igor V Kanatchikov. Canonical structure of classical field theory in the polymomentum phase space. *Reports on Mathematical Physics*, 41(1):49–90, 1998.

[40] Michael Forger, Cornelius Paufler, and Hartmann Römer. The poisson bracket for poisson forms in multisymplectic field theory. *Reviews in Mathematical Physics*, 15(07):705–743, 2003.

[41] Jürgen Struckmeier and Andreas Redelbach. Covariant hamiltonian field theory. *International Journal of Modern Physics E*, 17(03):435–491, 2008.

[42] Scott Axelrod, Steve Della Pietra, and Edward Witten. Geometric quantization of chern simons gauge theory. *representations*, 34:39, 1991.

[43] Igor V Kanatchikov. On quantization of field theories in polymomentum variables. In *AIP Conference Proceedings*, volume 453, pages 356–367. AIP, 1998.

[44] Mark J Gotay. Functorial geometric quantization and van hove's theorem. *International Journal of Theoretical Physics*, 19(2):139–161, 1980.

[45] Victor Guillemin and Shlomo Sternberg. Geometric quantization and multiplicities of group representations. *Inventiones mathematicae*, 67(3):515–538, 1982.

[46] G Giachetta, L Mangiarotti, and G Sardanashvily. Geometric quantization of mechanical systems with time-dependent parameters. *Journal of Mathematical Physics*, 43(6):2882–2894, 2002.

[47] Garrett Birkhoff and Saunders Mac Lane. *Modern Algebra*. Macmillan, 1965.

[48] Arturo Echeverria-Enrıquez, Miguel C Munoz-Lecanda, and Narciso Román-Roy. Geometry of multisymplectic hamiltonian first-order field theories. *Journal of Mathematical Physics*, 41(11):7402–7444, 2000.

[49] Hilbrand Johannes Groenewold. On the principles of elementary quantum mechanics. In *On the Principles of Elementary Quantum Mechanics*, pages 1–56. Springer, 1946.

[50] Igor V Kanatchikov. Toward the born-weyl quantization of fields. *International journal of theoretical physics*, 37(1):333–342, 1998.

[51] Igor V Kanatchikov. Precanonical quantum gravity: quantization without the space–time decomposition. *International Journal of Theoretical Physics*, 40(6):1121–1149, 2001.

[52] Maxim Kontsevich. Deformation quantization of poisson manifolds. *Letters in Mathematical Physics*, 66(3):157–216, 2003.

[53] Nicholas Michael John Woodhouse. *Geometric quantization*. Oxford University Press, 1997.

[54] Igor V Kanatchikov. Geometric (pre) quantization in the polysymplectic approach to field theory. *arXiv preprint hep-th/0112263*, 2001.

[55] Richard Arnowitt, Stanley Deser, and Charles W Misner. Dynamical structure and definition of energy in general relativity. *Physical Review*, 116(5):1322, 1959.

[56] G Magnano, M Ferraris, and M Francaviglia. Legendre transformation and dynamical structure of higher-derivative gravity. *Classical and Quantum Gravity*, 7(4):557, 1990.

[57] Charles W Misner, Kip S Thorne, and John Archibald Wheeler. *Gravitation*. W.H. Freeman and Company, 1973.

[58] Barry C Barish and Rainer Weiss. Ligo and the detection of gravitational waves. *Physics Today*, 52:44, 1999.

[59] F Matichard, B Lantz, R Mittleman, K Mason, J Kissel, B Abbott, S Biscans, J McIver, R Abbott, S Abbott, et al. Seismic isolation of advanced ligo: Review of strategy, instrumentation and performance. *Classical and Quantum Gravity*, 32(18):185003, 2015.

[60] Junaid Aasi, J Abadie, BP Abbott, Richard Abbott, TD Abbott, MR Abernathy, Carl Adams, Thomas Adams, Paolo Addesso, RX Adhikari, et al. Enhanced sensitivity of the ligo gravitational wave detector by using squeezed states of light. *Nature Photonics*, 7(8):613, 2013.

[61] Gregory M Harry, LIGO Scientific Collaboration, et al. Advanced ligo: the next generation of gravitational wave detectors. *Classical and Quantum Gravity*, 27(8):084006, 2010.

[62] DH Santamore and Yuri Levin. Eliminating thermal violin spikes from ligo noise. *Physical Review D*, 64(4):042002, 2001.

[63] L Blackburn, L Cadonati, S Caride, S Caudill, S Chatterji, N Christensen, J Dalrymple, S Desai, A Di Credico, G Ely, et al. The lsc glitch group: monitoring noise transients during the fifth ligo science run. *Classical and Quantum Gravity*, 25(18):184004, 2008.

[64] LIGO Scientific Collaboration, Virgo Collaboration, et al. Sensitivity achieved by the ligo and virgo gravitational wave detectors during ligo's sixth and virgo's second and third science runs. *arXiv preprint arXiv:1203.2674*, 2012.