

The Quantization of Gravity — an introduction

David Wallace*

Centre for Quantum Computation
The Clarendon Laboratory, Oxford University
Parks Road, Oxford OX1 3PU, UK

April 3, 2000

Abstract

This is an introduction to quantum gravity, aimed at a fairly general audience and concentrating on what have historically been the two main approaches to quantum gravity: the covariant and canonical programs (string theory is not covered). The quantization of gravity is discussed by analogy with the quantization of the electromagnetic field. The conceptual and technical problems of both approaches are discussed, and the paper concludes with a discussion of evidence for quantum gravity from the rest of physics.

The paper assumes some familiarity with non-relativistic quantum mechanics, special relativity, and the Lagrangian and Hamiltonian formulations of classical mechanics; some experience with classical field theory, quantum electrodynamics and the gauge principle in electromagnetism might be helpful but is not required. No knowledge of general relativity or of quantum field theory in general is assumed.

1 Introduction

Physicists — so the popular account goes — have reduced all interactions in the world to four fundamental forces: electromagnetism, the weak and strong nuclear interactions, and gravity. Forces and particles previously thought to be separate have been unified: electricity with magnetism, protons and neutrons (via the quark model) with pions. There is even a fairly good unification of electromagnetism with the weak interaction, and Grand Unified Theories adding the strong interaction (quantum chromodynamics) to this pair are constantly discussed. One often gets the impression that we are on the verge of a Theory of Everything, and that combining gravity with the others is just a matter of technical details.

And yet we do not have that unified theory, for a very good reason: we cannot yet construct a theory in which gravity is treated quantum-mechanically,

*email address: david.wallace@merton.ox.ac.uk

and there is little hope that such a theory is just around the corner. General relativity is far richer, conceptually and mathematically, than any of the classical field theories on which particle physics is based, and its quantization has eluded us for half a century, throwing up fascinating insights and open questions in the process.

This review is an attempt to explain what makes this quantum theory of gravity so unattainable, and to show what sort of progress has nonetheless been made. The strategy is as follows: to study, in parallel, electromagnetism and general relativity, and to show why the techniques that have served us so well in the quantization of the former theory should prove so difficult to apply to the latter. The mathematical complexities of gravity have been almost entirely omitted, and points relating to them are illustrated instead by analogy with the far simpler electromagnetic theory.

The review makes no pretence to be either complete or up-to-date. As far as the former is concerned, quantum gravity has become a vast subject and any attempt to survey the literature *in toto* would be doomed. As for the latter, while there have been exciting developments over the last two decades the mathematical complexity that would be required to discuss them would be beyond the scope of this paper. In particular, the huge and growing field of string theory is barely mentioned, despite its claim to be the most popular current route to a quantum theory of gravity. This omission is partly because string theory tries to construct a theory of all interactions, going beyond gravity. Primarily though, it is simply because the author is not qualified to discuss it.

As such, there is little in this review which would be news to a worker in the field in the early 1980s. However, the questions raised here were seen as relevant then and by-and-large are still seen as relevant today.

2 Classical Electromagnetism in Minkowski space-time

2.1 The action principle

For the purposes of comparison with gravity, it is easiest to begin with the covariant action form of electromagnetism. Recall that the theory is defined on the space of U(1) connections on spacetime [48] — that is, on the space of real vector fields \mathcal{A}_μ on spacetime, identifying any two vector fields which differ by a local gauge transformation (of which more later). An appropriate action for such a theory is

$$S_D = \int_D d^4x \mathcal{L} \tag{1}$$

where

$$\mathcal{L} = -\frac{1}{4} \mathcal{F}_{\mu\nu} \mathcal{F}^{\mu\nu},$$

$$\mathcal{F}_{\mu\nu} = \partial_\mu \mathcal{A}_\nu - \partial_\nu \mathcal{A}_\mu$$

with Greek indices ranging over spacetime indices 0,1,2,3, and with D any region of spacetime[42].

What this means is that for a classical field configuration within D to be dynamically acceptable, S_D for that configuration must remain unchanged under small variations of the field vanishing outside D . Introducing a small variation

$$\mathcal{A}_\mu \rightarrow \mathcal{A}_\mu + \delta\mathcal{A}_\mu \tag{2}$$

we get

$$0 = \delta D_S = - \int_D d^4x \mathcal{F}^{\mu\nu} \partial_\mu [\delta\mathcal{A}_\nu] \tag{3}$$

and integrating by parts gives us the familiar Maxwell equations in their covariant form:

$$\partial_\mu \mathcal{F}^{\mu\nu} = 0. \tag{4}$$

As is well known, this action is invariant under the ‘gauge’ variations

$$\delta\mathcal{A}_\mu = \partial_\mu f, \tag{5}$$

where f is any real function on spacetime. Physically speaking, this is the motivation for interpreting \mathcal{A} as a *connection* — that is, as a rule telling us the phase change for a (quantum) test charge when it is moved along a given path. Classically we can take \mathcal{F} to be fundamental and ignore \mathcal{A} ; quantum-mechanically phenomena such as the Aharanov-Bohm effect([1, 16]; [48, pages 365–369]) demonstrate the experimental consequences of \mathcal{A} .

The gauge freedom of the vector field \mathcal{A} is our motivation for viewing the space of possible field configurations not as the full space of vector fields on spacetime, but as the space of fields modulo gauge transformations. Viewing this space as a space of connections — that is, as a space of geometric entities — facilitates the comparison with gravity, but the reader will lose little by ignoring this aspect in the rest of this review.

2.2 Splitting spacetime into space and time: the 3+1 formalism

In this section, we break the manifest relativistic covariance of electromagnetism, essentially by choosing a preferred Lorentz reference frame. This is essential if we are to understand the dynamics of the theory, for dynamics is about the evolution of states through time and cannot be considered without picking out a preferred notion of time. Furthermore, this split is necessary to go to phase space and the Hamiltonian formalism, which we need to do in order to formulate a quantum theory.

The reader who feels that to make such a break is to do violence to relativity, may be comforted by the following observation. To specify a system’s state, we have to give its position in phase space: that is, the positions and momenta of all its constituent parts¹. However, once we know the positions and momenta

¹In the case of a field theory, this means the state of every point of the field at a given time.

we may calculate the state of the system at all times, and conversely if we know its state at all times we know where in phase space it is at any given time. The upshot of all this is that points in phase space are in 1:1 correspondence with solutions to the dynamical equations — and the space of solutions is not in any way dependent on our choice of reference frame. This ‘covariant phase space’ formalism will not be used further in this review, but is discussed in [72] and (in the context of quantum gravity) [13]. Further issues to do with the non-covariance of the theory will be discussed in section 4.1.

To make the split is very simple: choose a set of inertial coordinates on space-time and choose our region D to be all points with time coordinates between t_1 and t_2 . Then the action is

$$S = \int_{t_1}^{t_2} dt \int d^3\mathbf{x} \mathcal{L}. \quad (6)$$

Breaking \mathcal{A} into its components $\mathcal{A} = (V, \mathbf{A})$ we get the unlovely expression [45]

$$S = \int_{t_1}^{t_2} dt L \quad (7)$$

where

$$L = \frac{1}{2} \int d^3\mathbf{x} \left(\left(\nabla V - \frac{d\mathbf{A}}{dt} \right)^2 - (\nabla \times \mathbf{A})^2 \right). \quad (8)$$

Though this may look unfamiliar, it is now in the form we normally associate with nonrelativistic, classical-mechanics system: the dynamics are given by a Lagrangian which is a function of the configuration of the system at a given time and of the rate of change of that configuration. The only difference is that here, by ‘configuration’ we mean the \mathbf{A} and V fields on 3-dimensional space at a given time. To extract the dynamics we can use Hamilton’s variational principle

$$0 = \delta \int_{t_1}^{t_2} dt L \quad (9)$$

(note that this is just a thinly disguised version of the covariant action principle used above) from which follows the Euler-Lagrange equations. In fact these will turn out to be the familiar Maxwell equations, but we shall derive these later from the Hamiltonian viewpoint instead.

Given this format, it is natural to ask what is the ‘configuration space’, which must be analogous to the set of particle positions in classical mechanics. We might expect it to be the space of pairs (V, \mathbf{A}) but this is spoiled by gauge invariance. From elementary electromagnetism we know that V may be adjusted arbitrarily, so we do not treat it as a dynamical variable. But even the space of vector fields \mathbf{A} is too large, as we can always add the gradient of an arbitrary function to \mathbf{A} without changing the physical state. This issue will be discussed further in section 2.4.

2.3 The Hamiltonian formalism

To quantize a classical theory we usually begin with its description in the Hamiltonian formalism. To facilitate this we need the Hamiltonian form of electromagnetism.

Recall that in the case of finitely many point particles with coordinates q^1, \dots, q^n , we define the momenta p_1, \dots, p_n by [32]

$$p_j = \frac{\partial L}{\partial \dot{q}^j} \quad (10)$$

and rewrite Hamilton's principle as

$$0 = \delta S = \int_{t_1}^{t_2} dt (p_i \dot{q}^i - H(q, p)) \quad (11)$$

where

$$H(q, p) = p_i \dot{q}^i - L(q, \dot{q})$$

(in which \dot{q}_i is obtained as a function of p, q by inverting (10)). Then varying the path in Hamilton's principle gives Hamilton's equations

$$\dot{q}^j = \frac{\partial H}{\partial p_j} \quad (12)$$

and

$$\dot{p}_j = -\frac{\partial H}{\partial q^j}. \quad (13)$$

This procedure will also apply for electromagnetism, with two changes. The first is that electromagnetism has infinitely many degrees of freedom, with three coordinates per space point. This means that the momentum will be another vector field, which (with the benefit of hindsight) we will denote as \mathbf{E} . The other is due to the gauge freedom, which means that not all the coordinates are true degrees of freedom. The infinite-dimensional form of (10) will be

$$\mathbf{E} = \frac{\delta L}{\delta \dot{\mathbf{A}}} \quad (14)$$

and a quick calculation gives

$$\mathbf{E} = \dot{\mathbf{A}} - \nabla V; \quad (15)$$

of course this is just the expression for the electric field in terms of the potential, justifying our notation. The Hamiltonian turns out to be² ([45]; [67, pages 450–469])

$$H = \int d^3\mathbf{x} \left(\frac{1}{2} \mathbf{E}^2 + \frac{1}{2} \mathbf{B}^2 + V \nabla \cdot \mathbf{E} \right) \quad (16)$$

²To derive this we have had to integrate by parts, which necessitates our assuming the fields to drop off sufficiently quickly at infinity.

where we have defined $\mathbf{B} = \nabla \times \mathbf{A}$. Varying \mathbf{E} and \mathbf{A} in (11) gives analogies of Hamilton's equations:

$$\dot{\mathbf{E}} = \nabla \times \mathbf{B} \tag{17}$$

and

$$\dot{\mathbf{A}} = -\mathbf{E} - \nabla V \tag{18}$$

which on taking the curl of both sides becomes

$$\dot{\mathbf{B}} = -\nabla \times \mathbf{E}. \tag{19}$$

However, to get all of the dynamics we also have to vary the action with respect to V , which gives Gauss's Law³:

$$\nabla \cdot \mathbf{E} = 0. \tag{20}$$

This equation has a fundamentally different structure from the other dynamical equations above: while they tell us how \mathbf{E} and \mathbf{A} evolve from given initial conditions, Gauss's Law puts a constraint of what those initial conditions *are*. To be more specific, not all vector fields are admissible momenta: only those which are divergence-free can be used. An equation like this is called a constraint [45, 67], and will appear whenever a theory has gauge freedom.

2.4 Phase Space

Notionally, then, the phase space of electromagnetism consists of pairs (\mathbf{A}, \mathbf{E}) of vector fields on 3-space, implying that it has six dimensions per space point. However, this is complicated by the gauge freedom: in the space of \mathbf{A} -fields we should really treat any two fields differing by a gauge transformation as the same field (leaving two dimensions per space point), and in the space of \mathbf{E} -fields we have to impose the constraint $\nabla \cdot \mathbf{E} = 0$ at every space point, again reducing the dimensionality of the space of \mathbf{E} -fields from 3 to 2 per space point⁴.

This leaves electromagnetism having two degrees of freedom per space point, which accords with our quantum description in terms of massless spin-1 particles with two orthogonal spin states.

This restriction of the apparent phase space by gauge invariance and constraints is characteristic of gauge theories, and makes it hard to describe and classify the true dynamical states. Two normal ways of doing it are

1. To choose a particular gauge and describe the theory in that gauge, and
2. To find some way of describing the states that is itself gauge-invariant.

³The fourth Maxwell equation, $\nabla \cdot \mathbf{B} = 0$, follows trivially from the definition of \mathbf{B} .

⁴Here we see the reason to expect constraints in the theory: the 'position' \mathbf{A} and the 'momentum' \mathbf{E} should be dynamically conjugate, so if gauge invariance lowers the actual degrees of freedom of \mathbf{A} -fields then something else must restrict the degrees of freedom of \mathbf{E} -fields.

In general both of these (and especially the second) can be very difficult. In fact, in the case of electromagnetism the theory is simple enough that we can do either: there exist simple gauge choices like $\mathbf{A}^3 = 0$ (axial gauge) or $\nabla \cdot \mathbf{A} = 0$ (London gauge) which can be imposed at all times and for all solutions, and we can usefully describe the configuration space as

The space of equivalence classes of vector fields differing only by the gradient of a potential

(note that this is only a useful description because the mathematical description of a gauge transformation in electromagnetism is easy to understand).

3 Classical General Relativity

3.1 The dynamics of spacetime

Electromagnetism is a theory describing how certain fields evolve on a fixed spacetime background. So is Dirac field theory; so is QCD; so is every field theory used in modern physics, with one exception.

General Relativity differs from every other field theory in that spacetime itself is the dynamical entity under consideration: ‘gravitation’ occurs because curved space affects the path of matter. In Wheeler’s famous words [47],

Space acts on matter, telling it how to move. In turn, matter reacts back on space, telling it how to curve.

Unlike in the case of Newtonian gravity, the ‘gravitational field’ (i. e. curvature of spacetime) can propagate and have degrees of freedom even without matter — so vacuum gravity contains interesting objects, such as black holes and gravity waves.

The field equations for general relativity are highly complex, consisting of ten coupled non-linear partial differential equations, and I shall not write them here (see e. g. [47, 67]). This makes finding exact solutions extremely difficult and only a few are known, usually corresponding to situations with a high degree of symmetry (such as black holes and homogenous cosmologies). However, in the limit of a very weak field, we can approximate spacetime as nearly flat and look at the dynamics of the small variations from flatness. In this limit the equations become linear and can be solved approximately (this is the so-called ‘weak field’ approximation [47, 67]).

The fact that general relativity does not have a background spacetime makes it difficult to treat conceptually as well. In electromagnetism, for instance, we are entitled to use the background spacetime for reference: it is meaningful to say ‘at spacetime point x the value of the \mathcal{F} -field was such-and-such’. This does not work for gravity: to say ‘at spacetime point x the structure of spacetime was such-and-such’ implies that we have a meaningful way of referring to spacetime points in the absence of spacetime, which clearly is nonsense⁵.

⁵Some of the philosophical consequences of this are discussed in [24], Chapter 9, and references found therein.

3.2 Gauge invariance in general relativity

The ‘gauge’ invariance of general relativity actually corresponds to the difficulties noted at the end of the last section, in giving coordinates without reference to spacetime. However, for our purposes it can more simply be understood as follows.

To describe the dynamics of gravity, we need — as with electromagnetism — to split spacetime into space and time. However, in a theory where spacetime itself is dynamic there can be no preferred way of making this split. We have to cut 4-dimensional spacetime up into 3-dimensional slices, but in general we can do this in any way we like, freely choosing curved, bumpy slices. Any such cutting defines a concept of time, because we move forward in time by moving from one slice to the next. Gauge transformations correspond to a change in the definition of time, by changing our choice of how to slice. (This arbitrariness is often referred to as ‘many-fingered time’; see [47] for a discussion restricted to the classical context).

The mechanics of making a split were developed in the work of Arnowitt, Deser and Misner [2] on canonical general relativity; see, e.g., [47] for a brief treatment, and [39] for a full survey.

To understand better how much choice we have in the definition of time, we contrast general relativity with two previous theories of spacetime. In Newton’s conception of spacetime⁶ there is a unique way of making the division: time is universal, and this concept makes instrumental sense since there is no upper limit on speed of signal propagation in Newton’s physics⁷.

In special relativity we have more freedom. As is widely known there is no preferred reference frame in special relativity. However, there is still a preferred class of reference frames: inertial reference frames, i.e. those with Lorentz coordinates, related by Lorentz transformations. In the language used here, these preferred frames correspond to a preferred family of slices, and this family is in fact the family of *flat* slices. But the only reason that there exist such flat slices is because, in special relativity, spacetime itself is flat. If we take a highly non-flat spacetime — as is the generic case in general relativity — flat slices simply do not exist.

In fact, the difference between the freedom of choices of slicing in special relativity and in general relativity is analogous to the difference between local and global gauge transformations in quantum mechanics [43, 48]. Given a quantum wave-function, we can always make a global transformation

$$\psi(x) \rightarrow e^{i\phi}\psi(x) \tag{21}$$

but to make a local transformation

$$\psi(x) \rightarrow e^{i\phi(x)}\psi(x) \tag{22}$$

requires us to introduce the connection \mathcal{A} which in turn inspires us to make this connection dynamical. Similarly the Lorentz transformations of special

⁶Originally published in [50] and discussed extensively in [24].

⁷I am grateful to Harvey Brown for this observation.

relativity are necessarily global, i. e. affect the definition of time across all of spacetime, whereas in general relativity we can deform the slicing surfaces in a small region without affecting them outside that region. Much effort (e. g. [31, 36, 37, 49, 65]) has been put into making this more than an analogy, and trying to understand general relativity as the gauge theory built around special relativity. This program has had a degree of success, but unfortunately it has also led to much confusion and has tended to obscure the significant structural and conceptual differences between general relativity and other gauge theories.

In fact, one of these conceptual differences will be very important for us here. In electromagnetism two configurations differing only by a gauge transformation are regarded as different descriptions of the same physical state. However, in general relativity one possible ‘gauge transformation’ will be to move all of the slices forward through spacetime. If we keep our eye on a given slice, that slice will then describe a ‘later’ section of spacetime — in other words, time evolution is just one form of gauge transformation, and if we identify gauge-equivalent configurations then we will freeze out the time evolution [13, 45]. Classically this is by and large a curiosity; in the quantum theory it is a ferocious problem, as we shall see.

3.3 Phase space

To quantize gravity, we would like a description on phase space, and — again by analogy with electromagnetism — this turns out to be nontrivial because of the gauge transformations. We once again make some splitting of spacetime into space and time; note that the degree of arbitrariness here is much greater than for electromagnetism as we may no longer select a Lorentz slicing⁸. The configuration space of the theory, if we ignore the gauge freedom, is then the space of 3-dimensional geometries; the phase space turns out to include, in the guise of momenta, information on how to embed that three-dimensional geometry into a 4-dimensional one. The Hamiltonian is fearsomely complicated (see [67] for a derivation, and [28] for a detailed discussion) but its derivation follows the same pattern as for electromagnetism — it just has more algebra.

Of course, we have no business ignoring the gauge freedom. As with electromagnetism, we will find that some configurations should be identified as being gauge transformations of one another, and that there will be some constraint (the so-called ‘Hamiltonian’ constraint’) on the admissible states. This constraint is extremely complicated, and in general it is a very hard problem to check if any given state satisfies it⁹.

So far all of this is in accordance with electromagnetism - but now remember that gauge transformations correspond to changes in the definition of time. If

⁸Of course, the only reason we were able to select such a slicing before is that we were considering electromagnetism in *Minkowski* spacetime. If we were to study the propagation of electromagnetic fields on some more arbitrary fixed spacetime, this selection would not be possible.

⁹This is a significant oversimplification. There are other, ‘momentum’, constraints corresponding to changes in the coordinates used on the spatial slice, but these are conceptually and technically much simpler to handle and will be omitted hereafter.

we are identifying two states as identical under gauge transformations, we are just saying that they describe the same system at different times. No wonder that the constraint is so hard to solve: it encodes the dynamics for the entire system!

In fact, this is why we call the constraint ‘Hamiltonian’. It is responsible for the time evolution, and turns out to be equal to the actual Hamiltonian function. A corollary of this is that — since by definition the constraint has value zero on admissible states — all physical states have zero energy¹⁰.

As should be apparent from the above, understanding the true degrees of freedom of gravity is fearsomely difficult — far harder than for electromagnetism. The second strategy in section 2.4 (finding a simple gauge-invariant description) seems all but impossible, as it would amount to solving the field equations in the general case. As for choosing a specific gauge, this is extremely difficult to do: how can we specify a way of slicing spacetime that will make sense irrespective of the contents of that spacetime? In fact one way is known [73, 74, 47] (the so-called ‘York time’) but it has proved difficult to establish whether this method really is well defined. In any case, being restricted to a specific gauge is very inconvenient in general relativity (far more so than in electromagnetism) since we have lost our freedom to define time as we wish — analogous to, but more serious than, being forced to do all of special relativity in a specific frame.

4 Quantizing the electromagnetic field

Clearly there is something unsatisfactory about the whole notion of quantization. Presumably, the quantum theory is the more fundamental and we should begin from a quantum theory and then ‘classicalize’ it. Nonetheless, quantization remains the only program we have to generate realistic quantum field theories, so we are forced to use it. But it is important to bear in mind that the quantization process is a series of rules-of-thumb rather than a well-defined algorithm, and contains many ambiguities. In fact, for electromagnetism we shall find that there are (at least) two different approaches to quantization, and that while they appear to give the same theory they may lead us to very different quantum theories of gravity.

Before these methods are described, however, it is necessary to say something about covariance. The reader who is unconcerned about violations — real or apparent — of covariance may skip the next section.

¹⁰Actually, this only holds for ‘compact’ spaces, such as a closed universe. It is possible, by judiciously keeping track of boundary terms, to define energy, momentum etc. for an open universe in which the matter density and spacetime curvature tend to zero with large distances: see [67] for an introductory discussion, and [3, 6] for a full treatment.

4.1 Canonical vs. Path Integral approaches

Both of the approaches to quantization to be detailed below rely, as a first step, upon casting the classical theory into Hamiltonian form and thus breaking its manifest covariance. For all that phase space may be regarded in some sense as covariant (see section 2.2), and for all that the final quantum theory may have some measure of covariance restored to it, nonetheless it would be nice to have a manifestly covariant way to proceed.

On the surface, there is such a way: the Feynman path-integral approach to quantization [27, 43, 70] begins with the Lagrangian, rather than the Hamiltonian, formulation, and calculates all observable quantities in terms of the path integral

$$\int \mathcal{D}\mathcal{A} \exp i \left(\int d^4x \mathcal{L}[\mathcal{A}] \right) \quad (23)$$

where the second integral is over all of spacetime and the first is over all gauge-inequivalent \mathcal{A} -fields; this integral is manifestly covariant.

Unfortunately, appearances can be deceptive. To extract physical content from this path integral we need rules to go from the vacuum expectation value above to the expectation values for scattering between particles; these rules (the LSZ reduction formulae [41, 43]) and indeed the existence of the particle spectrum itself require the canonical formalism to derive. They, and indeed the path integral itself, fade into the background once the Feynman rules are derived, but for all that they are required in the foundations.

It may be objected that Feynman's path integral approach is usually described as an alternative formalism for all of quantum mechanics, and not just a calculation tool. Indeed this is so: the fundamental rule of the path integral approach is that, for two configuration-space eigenstates $|q_1; t_1\rangle$ and $|q_2; t_2\rangle$, we have

$$\langle q_1; t_1 | q_2; t_2 \rangle = \int_{q=q_1, t=t_1}^{q=q_2, t=t_2} \mathcal{D}q \exp \left(i \int_{t_1}^{t_2} dt L[q(t)] \right) \quad (24)$$

and from this all of the dynamics can be deduced. The problem is that, for a field theory, a 'point' in configuration space is a field configuration at a given time [70] — and of course, to define such a thing we have to choose a preferred time.

This example shows why we should not be too surprised to find problems with covariance when we try to quantize a theory. Ultimately states in quantum mechanics are *non-local*: entanglement makes it clear that a description in terms of isolated spacetime events simply will not do. But then any attempt to describe nonlocal states will struggle to avoid the problem of choosing preferred reference frames¹¹.

¹¹This is not to say that this problem cannot be avoided. The Heisenberg picture, with its time-independent states and operators which vary across spacetime, is much closer to manifest covariance — but the theory cannot, as far as I know, be formulated in this picture using the path-integral approach. When we consider the measurement problem, covariance issues become still more fierce: it appears impossible to carry out a collapse of the wavefunction without violating at least the spirit of relativity, and one of the major strengths of

See [5] for a discussion of the relationship between the canonical and path-integral methods, in the context of quantum gravity.

4.2 Particle quantization

The most common algorithm for quantizing electromagnetism goes as follows (I follow the treatment of [11]). Firstly, we assume the system to be confined within periodic boundary conditions ('putting the field in a box'); this step is fairly easily avoided but the presentation is simplified if we leave it in. A choice of gauge is made (London gauge is convenient for conceptual purposes; a relativistically covariant gauge is often easier for calculations) and the 4-potential \mathcal{A} is expanded in terms of plane waves:

$$\mathcal{A}_\mu(\mathbf{r}, t) = \sum_{\mathbf{k}, \alpha} (c_{\mathbf{k}, \alpha} \mathcal{A}_{\mathbf{k}, \alpha} + c_{\mathbf{k}, \alpha}^* \mathcal{A}_{\mathbf{k}, \alpha}^*) \quad (25)$$

where

$$\mathcal{A}_{\mathbf{k}, \alpha} = \sqrt{4\pi} \frac{e^{i(\mathbf{k} \cdot \mathbf{r} - \omega t)}}{\sqrt{2\omega}} \mathbf{e}^\alpha. \quad (26)$$

Here the \mathbf{e}^α ($\alpha = 1, 2$) are two orthogonal polarisation 4-vectors, whose precise form is given by (and determines) the choice of gauge, and the $c_{\mathbf{k}, \alpha}$ are expansion coefficients. Writing the Hamiltonian in this plane-wave expansion we get

$$H[\mathcal{A}] = \sum_{\mathbf{k}, \alpha} \frac{1}{2} \omega (c_{\mathbf{k}, \alpha} c_{\mathbf{k}, \alpha}^* + c_{\mathbf{k}, \alpha}^* c_{\mathbf{k}, \alpha}) \quad (27)$$

Guided by a mixture of theoretical¹² and experimental reasoning, we now reinterpret the theory as describing a system of spin-1 particles (photons) for which the $c_{\mathbf{k}, \alpha}$ are annihilation operators. This consists of promoting the c, c^* to the status of operators c, c^\dagger and imposing the commutation relations

$$[c_{\mathbf{k}, \alpha}, c_{\mathbf{k}', \alpha'}^\dagger] = \delta_{\mathbf{k}, \mathbf{k}'} \delta_{\alpha, \alpha'}. \quad (28)$$

The free-particle Hamiltonian now describes a collection of independent, free particles. To add interactions, we replace terms like $\mathcal{A}_\mu J^\mu$ in the Hamiltonian with their operator equivalents. Note that this process is fraught with ambiguities because quantum operators do not commute: for instance, the quantum form of the Hamiltonian (27) changes by an infinite constant if we replace $1/2(cc^* + c^*c)$ by c^*c in the classical version. Operator-ordering problems like this are a major source of ambiguities in quantization.

We pause a moment to consider what this theory is describing: a universe made up of particles moving through spacetime. For electromagnetism in the

Everett-style interpretations [26, 20] is the promise they hold out for a quantum theory which is more naturally compatible with relativity [60, 69]. See [46] for a discussion of the tension between the measurement problem and Lorentz covariance.

¹²Mostly the similarity of the expression for the Hamiltonian to that of a system of harmonic oscillators.

absence of sources, this quantization can be treated as exact; for most other quantum field theories there is self-interaction and we have to treat the self-interaction terms as perturbative corrections. Doing dynamics in this framework is mostly restricted to considering scattering situations: a collection of free particles far away from one another collide and in due course some other collection of free particles are observed. Some non-perturbative calculations can be done, but the lack of general methods for non-perturbative quantum field theory remains a major lacuna of the theory (especially in quantum chromodynamics where it is required to explain phenomena like quark confinement and spontaneously broken axial symmetry [17, 61]). The formalism we have constructed here would seem to have difficulties even expressing such non-perturbative ideas, since its construction is perturbative in nature.

To do perturbative calculations, we have an algorithm *par excellence*: the Feynman diagrams, which can quickly be derived from the quantization procedure above [70] and lend themselves naturally to an interpretation in terms of scattering events. Further, the Feynman diagrams are manifestly covariant, even if the formalism we used to derive them was not.

Actually doing calculations with Feynman diagrams rapidly reveals the presence of infinitely large terms, which have to be removed by sleight-of-hand: this ‘renormalisation’ process is now mathematically fairly well understood even if it is still physically somewhat mysterious [12, 43]. Essentially whether or not a theory is renormalisable depends on the number of lines coming into any vertex in a Feynman diagram: if there are more than a certain number of lines then the number of infinities that have to be removed will increase with the complexity of the diagram, and in removing an infinite number of infinities all predictive power would be lost. Fortunately this does not occur with electromagnetism: each vertex on a Feynman diagram corresponds to an interaction term (such as $\mathcal{A}_\mu \psi^\dagger \gamma^\mu \psi$ in QED) and the number of lines entering the vertex is equal to the power of the fields in the interaction term (so, for instance, in the term above there is one line corresponding to the \mathcal{A} -field and two corresponding to the ψ -field, for a total of three. This is the maximum permissible for a theory of electrons and photons, but it *is* permissible: QED is renormalizable [43]). These renormalisation considerations show why we would have difficulty with a theory with non-polynomial interaction terms: to write down the Feynman rules for such a theory we would need to make a power-series expansion, which would contain terms of all powers. Most of these terms would be nonrenormalisable.

4.3 Field quantization

While the method of the previous section seems to correspond nicely to what we observe in particle-physics experiments, it is at least noteworthy that it seems to bear little resemblance to the quantization procedure used for ordinary mechanical systems (such as a point particle).

Recall that in this case, the quantum system is described by a wave-function on the configuration space, i. e. a function ψ from the configuration space to the complex numbers. The physical interpretation of this wave function is that, if U

is some subregion of configuration space, the probability of finding the system to be in U on measurement is

$$\Pr(\text{system in } U) = \int_U d\mathbf{x} |\psi(\mathbf{x})|^2. \quad (29)$$

The position functions q^1, q^2, \dots, q^n are replaced by multiplication operators:

$$(\widehat{q^j} \psi)(\mathbf{x}) = x^j \psi(\mathbf{x}) \quad (30)$$

and the momentum functions p_1, \dots, p_n by derivative operators:

$$\widehat{p_j} \psi = -i \frac{\partial \psi}{\partial x^j}. \quad (31)$$

Functions of p and q are defined similarly, although major problems *may* result from operator-ordering ambiguities: classically, $pq = qp$, but this does not hold at the operator level.

In field quantization, we try to duplicate this process for the electromagnetic field (here I follow the treatment of [45]). The analogies we need are (ignoring, for the moment, subtleties due to gauge invariance):

Configuration space	Space of vector fields \mathbf{A}
Position functions q^i	Potential functions $\mathbf{A}^i(x)$ at each point x
Momentum functions p_i	Field-strength functions $\mathbf{E}_i(x)$ at each point x

So our quantum description of a state will be some sort of wave-functional $\Psi[\mathbf{A}]$ which assigns a complex number to each vector field. The physical interpretation of this wave functional will be that, if U is a collection of \mathbf{A} -fields, the chance of the system having vector potential in U upon measurement will be

$$\Pr(\text{system in } U) = \int_U \mathcal{D}\mathbf{A} |\Psi[\mathbf{A}]|^2 \quad (32)$$

(albeit this equation is largely figurative because of the difficulties in making rigorous sense of functional integrals¹³). The ‘position’ operator will be

$$(\widehat{\mathbf{A}^i(x)} \Psi)[\mathbf{A}] = \mathbf{A}^i(x) \Psi[\mathbf{A}] \quad (33)$$

¹³Actually, for a gauge theory there is another problem: this integral, as written, integrates over all \mathbf{A} -fields and not just the gauge-inequivalent ones. Strictly speaking we should introduce some sort of Fadeev-Popov gauge-fixing term into the functional integral [17], but this will not have any particularly interesting consequences.

and the ‘momentum’ operator will be

$$\widehat{\mathbf{E}}_i(x)\Psi = -i\frac{\delta\Psi}{\delta\mathbf{A}^i(x)}. \quad (34)$$

The Hamiltonian becomes

$$\hat{H} = \int d^3x \frac{1}{2} \left(\widehat{\mathbf{E}}(x)^2 + \left(\nabla \times \widehat{\mathbf{A}}(x) \right)^2 \right) + V \nabla \cdot \widehat{\mathbf{E}}(x) \quad (35)$$

(note that V is not quantized, since it is not genuinely dynamical but just chooses the gauge). Factor-ordering problems will not occur here because the $\widehat{\mathbf{A}}$ and $\widehat{\mathbf{E}}$ operators are not mixed.

There is one important difference between this system and the ordinary mechanical systems: the constraint equation (20). It turns out [43] that we cannot impose it as an operator equation because this will be too restrictive (it would be analogous to requiring the classical constraint equation to hold for all vector fields, not just the dynamically acceptable ones!) Instead we impose it as a condition on states:

$$\nabla \cdot \widehat{\mathbf{E}}(x) |\psi\rangle = 0. \quad (36)$$

Clearly this will restrict us to a subspace of the original Hilbert space: the space of physical states.

In fact, it turns out that we have a bonus: the physical states are already gauge invariant. This is because, as it turns out, the constraint is the infinitesimal generator of gauge transformations! As such, any state which it annihilates will be gauge invariant (the technical details may be found in the appendix). This is a completely general phenomenon, holding for any gauge theory.

For electromagnetism, we can reach this directly: recall that the ‘true’ configuration space of the system is the space of equivalence classes of vector fields under gauge transformations. We could have avoided any mention of the constraint by starting off with wave-functionals on *this* space, i. e. with functionals Ψ for which

$$\Psi[\mathbf{A} + \nabla f] = \Psi[\mathbf{A}]. \quad (37)$$

Effectively, we have imposed the constraint classically and only then quantized. However, the only reason that this trick works is that electromagnetism is simple enough to allow this direct description of the reduced configuration space (see [67, pages 450–469]). In general — and for gravity in particular, as we shall see — we must impose the constraint at the quantum level.

To summarise: physical states are those annihilated by the constraint operator, and these states are gauge-invariant.

The duality between particle and field descriptions of a quantum field theory means that these two methods of quantization produce physically equivalent theories, differing only by a unitary transformation. However, when we attempt to quantize gravity it is not clear that the duality will continue to hold: indeed, the two main approaches to quantum gravity follow the two different methods, and seem to have very little in common!

5 Covariant quantum gravity

5.1 Parting of the ways

Quantum gravity is a field driven almost entirely by theoretical considerations, with a crippling shortage of experimental data. As such, it has — for decades — been split into different schools, following different methods to construct different theories. Only when (and if) any one of these schools produces a theory developed enough to make experimental predictions, will it be possible finally to adjudicate between them¹⁴.

The two main approaches to quantum gravity — ironically they both go back to the same set of papers by Bryce Dewitt [21, 22, 23] — are known in the literature as ‘covariant’ and ‘canonical’ quantum gravity. An informal survey of published papers in 1998 [40] estimated that about 70% of publications were on covariant gravity (or more accurately its descendants, primarily string theory) and 30% were on the canonical theory. In addition, a small minority of papers covered a myriad other approaches: the Euclidean path integral approach of Hawking and co-workers [33]; the twistor theory pioneered by Penrose [57, 54], etc. , none of which will be discussed here.

5.2 Formulating covariant quantum gravity

The idea of the covariant approach to quantum gravity is to develop the theory along the lines seen in particle physics: that is, through interacting particles, Feynman diagrams and the like. (The use of Feynman diagrams is the justification for calling the theory ‘covariant’.)

To do this, we need to follow as far as is possible the particle quantization method of section 4.2, so we need an expansion of a general solution to the general-relativistic field equations in terms of plane waves. In general this will be impossible because of the extreme nonlinearity of general relativity, so we begin with the weak-field approximation mentioned in section 3.1 and treat the nonlinear terms as perturbations [43]. The weak-field approximation *is* linear, and we can make a plane-wave expansion; quantization reveals that the particle interpretation of gravity is in terms of massless spin-2 bosons: ‘gravitons’. We now include interactions, including the self-interactions due to the nonlinearity of general relativity, using Feynman diagrams.

5.3 Conceptual problems

It should be clear that this approach is essentially perturbative: we have taken an intrinsically non-linear theory, treated it as linear, and added its nonlinear part as a correction. Now this is not automatically a mistake: it is essentially the approach used in the rest of quantum field theory, after all. However, it

¹⁴Of course, it would be nice to suppose that the eventual theory would be some sort of synthesis of the various approaches — but currently we have no way of knowing how likely this is.

means that the framework is likely to have difficulty accommodating the sort of intrinsically non-perturbative behaviour which we would expect in the vicinity of strong fields. In fact as was mentioned in section 4.2 this already causes problems in QFT: specifically in quantum chromodynamics it makes it very difficult to understand features like quark confinement. In gravity, we might expect similar difficulties in coming to grips with features such as black holes, and other regions where, classically, spacetime is highly non-flat.

And this touches on a more profound problem with covariant quantization. The theory we have produced is one which describes spin-2 particles propagating through spacetime. But surely the theory is supposed to be a theory *of* spacetime? Or put another way, if quantized spacetime ‘is’ a collection of spin-2 particles, in what space do those particles exist¹⁵?

The concrete problems that this causes become more apparent when we consider how this quantum framework could describe certain classical features of gravity. In particular, general relativity, in making spacetime dynamical, allows the causal structure of spacetime to be influenced by entities like gravity waves (the causal structure of spacetime is the structure that defines concepts of ‘past’ and ‘future’; it tells us which events may causally affect other events, and is the generalisation to curved spacetime of the light-cones of special relativity). But, as has been stressed by Penrose among others [18], in covariant quantization the gravitons live in Minkowski spacetime. The causal structure of this spacetime is *fixed* and no amount of graviton propagation will change it.

If canonical quantum gravity were to be empirically useful, no doubt these problems could be sidestepped and left to be understood later: this is a method which has worked very well for quantum mechanics in the past. Unfortunately — as the next section shows — this is not the case due to formidable technical difficulties.

5.4 Technical problems, and further developments

(This treatment follows [43]).

The technical problems in this approach to quantum gravity can be summed up in six words:

Covariantly quantized gravity is not renormalisable

It almost goes without saying that this is a disaster. Nonrenormalisability means that there is no way of proving the infinities in the Feynman diagrams to be controllable: that, on the contrary, that there is every reason to expect them to increase without limit as the Feynman diagrams become more complicated.

¹⁵It might be objected that classical general relativity does not *have* to describe a dynamical spacetime: that the evidence is perfectly compatible with a flat-space field theory that just *looks* as if it describes a dynamical spacetime. This is of course perfectly true, just as astronomical evidence is perfectly compatible with a geocentric theory that just *looks* as if it describes a solar system with the sun at the centre. See [20] for further discussion of this point.

Obviously, this makes it completely impossible to extract information from the theory.

To understand why nonrenormalisability should occur, recall the remarks made at the end of section 4.2. There, it was explained that for a theory to be renormalisable the number of lines entering each Feynman-diagram vertex — equal to the power of the corresponding term in the interaction Hamiltonian — had to be below a certain finite number. But the interaction term in general relativity is *non-polynomial* in the fields, so (also as explained in section 4.2) we have to make a power-series expansion, leading to an infinite sum of interaction terms each of a higher power than the last. Clearly, for all but a finite number of these terms renormalisability is impossible.

Even when a theory is formally ‘nonrenormalisable’ in this way there is always the possibility that, through a series of miracles, the various infinities will cancel out. The formidable mathematical difficulties involved in quantum gravity have made calculations to test this very difficult, but it is now generally recognised that no such fortuitous cancellation is likely to occur above the one-loop level. (Having said this, it is as well to be cautious: for many years it was thought that *3-dimensional* quantum gravity was nonrenormalisable; then it was shown to be exactly quantizable ([71]; discussed in [13]) — and of course later analysis showed it to be renormalisable all along [58].)

Most workers in this branch of quantum gravity have so far concentrated on solving these technical difficulties, rather than conceptual ones. The usual approach has been to modify the classical theory in the hope of producing one with more tractable behaviour. In the 1970s gravity was coupled with a spin-3/2 field and a new symmetry (supersymmetry) introduced: this was shown to cause ‘fortuitous’ cancellation at the 2-loop level but unfortunately remained divergent at the 3-loop level. Since the 1980s most workers in the field — driven partly by input from particle physics — have moved towards string theory, whose Feynman diagrams appear to be not just renormalisable but finite, and which appears to have general relativity as a low-energy limit. String theory lies beyond the scope of this review and so I shall not discuss it, except to say that (despite recent attention to the issue, e. g. [64]) there is as far as I know no consensus on how to solve the problems of section 5.3 within a string-theoretic framework.

6 Canonical Quantum Gravity

6.1 Formulating canonical quantum gravity

In the canonical quantization of gravity, we try to emulate the ‘field quantization’ of electromagnetism discussed in section 4.3. To be precise, we start with the phase space of general relativity — i. e. the space of 3-dimensional geometries — and take our quantum states to be complex-valued wave-functionals on this space (so each wave-functional associates an amplitude to each geometry). We then impose the Hamiltonian constraint of section 3.3 as an operator equa-

tion on the states: this is the famous Wheeler-DeWitt equation ([21]; discussed in [13, 58]).

The technical problems raised here are extremely formidable. The Wheeler-DeWitt equation is extremely difficult to treat mathematically: indeed the mathematical pathologies here may be worse even than in the covariant case. In fact, until very recently no solutions at all were known.

In an attempt to get around this, we might try quantizing on the reduced phase space rather than the full one (recall that the reduced phase space consists only of those states satisfying the constraints, and furthermore identifies states differing by a gauge transformation. In this case the Wheeler-DeWitt equation would automatically be satisfied (essentially we are choosing to solve the Hamiltonian constraint before quantization rather than after it.) The problem is that explicitly describing the phase space is a virtually hopeless task, equivalent to a general solution of the field equations¹⁶.

As well as these technical problems, there are formidable conceptual problems. One of the most acute is the problem of defining the observables of the theory. Remember that an observable must be gauge-invariant: but with the gauge group of gravity including time evolution, this means that all of the observables must be conserved. This is a strong restriction, and in fact no-one has yet found any observables at all. Some intuition as to why can be got when we consider what we mean by observables in (say) electromagnetism: the paradigm observables are the value of the fields at *points in spacetime*. But, as was stressed in section 3.1, when there is no background spacetime to serve as a source of reference points this option is not available to us. Even if we could find observables it is unclear how they could really describe the universe if they must be time-independent. This issue is discussed for toy models in [45] and for quantum gravity in two space dimensions in [13]; the conclusion seems to be that, again, we would need to solve the classical dynamics in all generality to find all of the observables.

6.2 The Problem of Time

Perhaps the most profound conceptual problem of the canonical approach — almost certainly the most extensively discussed — goes by the name of the problem of time. The nature of the problem is very simple:

1. When we quantize a field theory with gauge transformations, the physical states will turn out to be gauge-invariant.
2. In general relativity, time evolution is just a particular gauge transformation.
3. Therefore, physical states of quantum gravity are time-independent.

¹⁶To see why this is, note that to describe the reduced phase space we have to identify states differing by a gauge transformation. Since in general relativity the dynamical evolution is just a gauge transformation (recall section 3.3) this would require saying, for any two states, whether one is the time-evolved version of the other — i. e. solving the dynamics.

This can be seen in various other ways: the Hamiltonian of quantum gravity vanishes on physical states; the Wheeler-DeWitt equation contains no mention of time, etc. Clearly the problem is deeply intertwined with the problem of defining observables mentioned above, but it is simpler and starker: *what has happened to time?* It is hard to take seriously a theory which denies that systems change in time!

Nonetheless, various strategies for dealing with the problem have been developed. The general principle is to assume that time evolution is somehow encoded in the other dynamic variables. This was spelled out — in the context of ordinary quantum mechanics — by Page and Wootters [51], and goes as follows.

Suppose the subsystem of the universe in which we are interested has Hilbert space \mathcal{H}_s . Then to actually measure the time-evolution of this system we will need a real, physical clock. Let us suppose that this clock has Hilbert space \mathcal{H}_c . The clock will also have some set of eigenstates which might be labelled $|4 \text{ o'clock}\rangle, |5 \text{ o'clock}\rangle$, etc., and the clock Hamiltonian will carry it from one such state to the next.

So, suppose the system starts in some product state

$$|\text{system at 4 o'clock}\rangle \otimes |4 \text{ o'clock}\rangle. \quad (38)$$

Then time evolution will duly carry it, after an hour, into the state

$$|\text{system at 5 o'clock}\rangle \otimes |5 \text{ o'clock}\rangle, \quad (39)$$

and thereafter into

$$|\text{system at 6 o'clock}\rangle \otimes |6 \text{ o'clock}\rangle, \quad (40)$$

etc. But now suppose that we prepare the system (skating blithely over some formidable mathematical problems, including normalisation) in an entangled state,

$$|\psi\rangle = \sum_{\text{all times } n} |\text{system at } n \text{ o'clock}\rangle \otimes |n \text{ o'clock}\rangle. \quad (41)$$

Then time evolution will carry each element of this superposition into the next, leaving the overall state unchanged. There is a strong analogy here with plane waves in ordinary quantum mechanics: a plane wave is an equally weighted superposition of all position eigenstates, and it is invariant (up to a phase factor) under any translation. So this entangled state is an equally weighted superposition of all instants of time and is invariant under time translation. Nonetheless — and here is the essence of the Page-Wootters strategy — the information about time evolution is still present, encoded in the entanglement correlations between the system and the clock.

It is unclear whether this strategy can deal with the problem of time in quantum gravity. A sceptical note was sounded by Unruh and Wald [66], who (among other objections) pointed out that the notion of time here will at best be approximate. This is because of a theorem they proved (the ‘no-clock theorem’)

to the effect that any realistic quantum clock will have a nonzero probability of running backwards. Nonetheless, the strategy has been widely discussed: see [7] and chapter 11 of [20] for a non-technical discussion, and [8, 9] for a detailed and philosophically-inclined analysis. The implications for observables are discussed by [59], using the terminology ‘evolving constants of motion’.

6.3 Recent progress: the Ashtekar formulation and knot theory

For most of the 1970s and 1980s relatively little work was done on the canonical approach to quantum gravity: researchers found the Hamiltonian constraint intractable and disliked the non-covariance that seemed inherent in a Hamiltonian description. This latter problem began to be seen to decline in importance when it was recognised that phase-space can be seen as covariant (see section 2.2), but the main impetus for research in the field since the late 1980s was the discovery by Abhay Ashtekar [4] of a new formulation of classical General Relativity in which the theory developed very strong analogies with other gauge theories, and in which the constraints became much simpler to solve. Space does not permit a discussion of this formalism (an extremely readable review paper is [58]) but it is worth mentioning one discovery which has come out of this work: a strong connection between general relativity and the mathematical theory of knots.

Knot theory, a branch of pure mathematics, is concerned with attempting to find rules to establish when one knot is the same as another (i. e. , one can be continuously deformed into another without untying it). The usual mathematical strategy for problems like this is to find ‘knot invariants’, mathematical objects which can be associated to a knot and which will not change no matter how much the knot is deformed. With hindsight, such invariants make excellent candidates for gauge-invariant states in quantum gravity: after all, if the quantity is unchanged by being deformed continuously it will be unaffected by gauge transformations. At time of writing, this topic has yet to be fully explored, but there appear to be deep connections (and the history of modern theoretical physics is filled with examples of deep and useful connections between branches of physics and of pure mathematics: consider Hilbert space and quantum mechanics, or fibre bundles and gauge theory, or vector calculus and Maxwell’s equations . . .). In particular, knot invariants have indeed furnished us with (a few) genuine gauge-invariant states [29], and as far as I know are the only such states known.

In the last decade canonical quantum gravity has been enjoying a renaissance: the technical problems are far from solved but it does at least seem possible to make progress on them, and although the conceptual problems are as fierce as ever it is perhaps an advantage of the approach that it brings them so clearly into view.

7 Evidence — experimental, observational, theoretical

7.1 Introduction

As was mentioned in section 5.1, quantum gravity suffers from a crippling lack of experimental findings to give us any indication about the ultimate theory. We know some such theory must exist to resolve the incompatibility between quantum mechanics and general relativity, but beyond that we are guided mainly by theoretical extrapolation. In this sense the development of quantum gravity resembles that of classical general relativity, driven primarily by theoretical and aesthetic considerations, more than it does that of quantum mechanics. Of course — as with general relativity — once we have the theory it will doubtless make many further testable predictions.

Furthermore, we are not entirely without guidelines: some of the main ones are sketched below. Some (few) refer to experimental results; more are due to cosmological observations, or are hints from our existing theories.

7.2 The parameters of the standard model

In general, it is hard to get direct data on quantum gravity from experimental particle physics: even the classical gravitational field between two subatomic particles is virtually unobservable, let alone quantum corrections. Further, the renormalisation procedure which makes it possible to do quantum field theory in the absence of an understanding of deeper theories, also tends to screen out observable effects due to those theories [12, 43]. This is largely why the Standard Model of particle physics has remained static for twenty years: higher-energy theories tend not to affect the relatively low-energy phenomena which we are able to observe.

Nonetheless the standard model has major unexplained features: primarily the large number (about 20) of physical parameters (masses of particles, fine structure constant, etc.) which have to be inputted by hand [17]. An explanation of the values of these constants certainly has to lie beyond the Standard Model, and it is at least possible that quantum gravity may have a hand in providing them (dealing as it does with extremely high-energy phenomena).

Further, the explanation given may be of a novel kind. As is now widely recognised, the values of the physical constants lie within an extremely narrow band of values compatible with the evolution of life. Small changes in, say, the mass of the electron or the fine structure constant would prevent star formation, or the formation of complex molecules, or any number of other finely balanced procedures. This implies that the physical constants will not simply be calculable from a higher-level theory, but will somehow have to be selected for — otherwise the compatibility of the universe with life will be an implausibly large coincidence. (See [10] for the definitive account of these *Anthropic arguments*, and [44] and [62, 63] for discussions from two very different viewpoints.)

7.3 Unsolved mysteries of cosmology

There are a number of unsolved problems in cosmology which are briefly mentioned here (all are more thoroughly discussed in [53]). None have generally accepted explanations within the current framework of physics, so some information might be gained about any of them by a better understanding of the unknown quantum-gravitational physics in the extreme vicinity of the Big Bang.

The horizon problem The universe, as we observe it, appears strikingly homogeneous, with this homogeneity increasing as we look further back in time. (Fractional fluctuations in the microwave background radiation are only about one part in 10^4 .) This is the case despite the fact that parts of the current universe far away from one another would have been out of causal contact with one another in the early universe, and so would have had no way to reach equilibrium with one another.

The flatness problem Generically, cosmological solutions of Einstein's equations can be open (describing infinite universes, with negatively curved space, expanding forever) or closed (describing finite universes, with positively curved space, doomed to collapse again in finite time). Only a very precise choice of the density of the early universe would produce a virtually flat universe, right on the boundary between open and closed — and yet our universe is virtually flat. To be more precise [53] the density of the universe is currently within one part in 10^3 of the critical density for flatness. Since the expansion of the universe tends to increase deviations from the critical value, this implies that at the point when our current theories fail us (about 10^{-35} seconds after the Big Bang) the density had to have been within one part in 10^{57} of the critical value¹⁷.

The preponderance of matter over antimatter The current universe contains an overwhelming predominance of matter, but physical laws are virtually symmetric in their treatment of matter and antimatter. Explanations for this imbalance within the domain of current physics are at best tentative.

Cosmological structure formation Current cold dark matter simulations of the expansion of the universe [19] produce a good fit to the currently observed large-scale structure, but require as an input a small level of inhomogeneities in the very early universe. These inhomogeneities have been observed in the microwave background radiation by the COBE satellite, but are still not understood theoretically.

In fairness I should add that there is a possible non-quantum-gravity explanation for the first two of these: the inflation mechanism (discussed in [53]). However, this explanation is still controversial, and in any case needs extremely fine tuning of the cosmological constant (see below!)

¹⁷It may or may not be a mere curiosity that a virtually flat universe is another requirement for the existence of life (see [10]).

7.4 The cosmological constant

General relativity contains a free parameter: the cosmological constant, which simulates a constant distribution of perfect fluid across spacetime. Einstein called its introduction ‘the biggest blunder of my life’ [30] and to this day it is not known empirically whether it is needed in the classical theory: certainly it is empirically indistinguishable from zero.

The situation is complicated by quantum theory. Quantum field theory assigns an energy to the vacuum of order \hbar per field mode — even imposing a cutoff around the quantum-gravity scale to make this finite gives an incredible energy density, of order 10^{107} joules per cubic metre. In nongravitational physics this energy has few effects, but it should give rise to an enormous effective cosmological constant, dwarfing the contribution from ordinary matter and totally contradicting observation.

This is a real dilemma in current physics (albeit one we can usually ignore) and is one of the most explicit clashes between general relativity and quantum theory. Simply removing the vacuum energy by fiat will not do: the Casimir effect ([15], discussed in [41]) proves its reality. The most conservative solution is to postulate a large, negative ‘real’ cosmological constant which cancels the vacuum contribution. However, this cosmological constant would have to be fine-tuned¹⁸ to better than one part in 10^{120} .

A recent review of the cosmological constant is [14].

7.5 Singularities

Even on its own terms (without the challenge from quantum theory) there is ground to think that general relativity needs modification. This is because the theory frequently predicts the occurrence of ‘singularities’, i. e. breakdowns of the predictability of the equations: the big bang is one example, and any black hole contains another. It is unclear to what extent these singularities require the replacement of general relativity (see [25] for a discussion) but the general consensus is that the very high curvatures of space found around them means that quantum gravity must be applied. Wheeler [47] in particular has strongly advocated the view that the singularities in gravitational collapse are our strongest clue to the nature of quantum gravity. This view has been championed more recently by Penrose [55, 56], in connection with the arrow of time: the singularity structure of spacetime is strongly time-asymmetric and Penrose has argued that this must be explained within the framework of quantum gravity.

7.6 Black hole thermodynamics

Possibly the strongest clue we have to the nature of quantum gravity is the discovery by Hawking [35] that black holes, contrary to the predictions of classical physics, must emit black-body radiation. This work has led to some extremely suggestive analogies between thermodynamics and statistical mechanics on the

¹⁸This fine-tuning is again required for the universe to support life.

one hand, and gravity on the other. Since Hawking's original discovery the result has been rederived within the framework of many approaches to quantum gravity, seeming to suggest it is not a mathematical anomaly; see [68] for a review.

Hawking's calculations were, of course, carried out within the framework of a semi-classical approximation, not true quantum gravity; nonetheless there is good reason to believe that the true theory must reproduce his results to be compatible with its classical limit. There is a nice analogy with the analysis by Einstein [52] of the A and B coefficients in atomic transitions: these are derived without using any detailed knowledge of quantum electrodynamics, but nonetheless QED has to conform to them on pain of violating fundamental principles of physics.

8 Conclusion

The problems involved in formulating a quantum theory of gravity are ferocious. Technical and conceptual problems beset us on all sides and there is precious little information from experimental data (one paper which the author has seen claims that quantum gravity will change the spectrum of Hawking radiation slightly — and then claims that this will be an 'experimental test' of quantum gravity!)

Still, progress has been made, and is continuing to be made. In a sense, the expectations we have of the completed theory are grand enough that it has to be difficult: different authors seek a resolution of the renormalisation problem in quantum field theory, a taming of the singularities of general relativity, a solution to the measurement (Schödinger-cat) problem, a reason for the direction of time, the unification of fundamental forces, It is far from clear which, if any of these will actually result, but certainly the complete theory will bring rich rewards for our understanding.

Currently, though, the complete theory is not at hand and there is no real reason to expect a breakthrough in the immediate future. In the words of one worker in the field [63],

We have made great progress, but the fact that general relativity and the quantum are not yet united means that we have no single picture of what the world is that we can believe in. When a child asks, What is the world, we literally have nothing to tell her.

Acknowledgements

I am grateful to Julian Barbour, Harvey Brown, Joy Christian, David Deutsch and Oliver Pooley for discussions of quantum gravity.

This work is based on a talk given in the Atomic and Laser Physics Department, Oxford University, on June 1st, 1999, and was completed in partial

fulfilment of the requirements for transfer from Probationary Research Student to DPhil student at Oxford University.

Appendix: The Gauss-law constraint generates gauge transformations

Recall that, in the field-quantization framework of section 4.3, the states are wave-functionals $\Psi[\mathbf{A}]$ on the space of \mathbf{A} -fields and the electric field operator $\widehat{\mathbf{E}}(x)$ is defined by

$$\widehat{\mathbf{E}}_i(x)\Psi = -i\frac{\delta\Psi}{\delta\mathbf{A}^i(x)}. \quad (42)$$

Now, if we had a *finite*-dimensional symmetry group, such as spin, there would be finitely many generators $\hat{\sigma}_1, \dots, \hat{\sigma}_n$ and a symmetry transformation would be given by

$$|\psi\rangle \rightarrow \exp\left(i\sum_i \lambda^i \hat{\sigma}_i\right) |\psi\rangle \quad (43)$$

where the λ^i give the specific gauge transformation (in the case of spin, for instance, they are related to the Euler angles). The infinitesimal version of (43) is of course

$$|\psi\rangle \rightarrow |\psi\rangle + i\epsilon \sum_i \lambda^i \hat{\sigma}_i |\psi\rangle. \quad (44)$$

In the case of gauge transformations the gauge group is infinite-dimensional so we replace the λ^i with a continuous function f and supply a continuous infinity of generators $\nabla^i \widehat{\mathbf{E}}_i(x)$. The analogue of (44) is then

$$|\psi\rangle \rightarrow |\psi\rangle + i\epsilon \left(\int d^3x f \nabla^i \widehat{\mathbf{E}}_i(x) \right) |\psi\rangle, \text{ i. e.} \\ \Psi[\mathbf{A}] \rightarrow \Psi[\mathbf{A}] - \epsilon \int d^3x f \nabla^i \frac{\delta\Psi}{\delta\mathbf{A}^i(x)}. \quad (45)$$

Integrating by parts, this gives us

$$\Psi[\mathbf{A}] \rightarrow \Psi[\mathbf{A}] + \epsilon \int d^3x (\nabla^i f) \frac{\delta\Psi}{\delta\mathbf{A}^i(x)}. \quad (46)$$

But this is just the first two terms of an infinite-dimensional Taylor expansion, and we have

$$\Psi[\mathbf{A}] \rightarrow \Psi[\mathbf{A} + \epsilon \nabla f] \quad (47)$$

which is a gauge transformation. We have proved what we set out to establish: that the constraint operators are the generators of the gauge group. It follows that any physical state, i. e. any state satisfying

$$\nabla^i \widehat{\mathbf{E}}_i(x) |\psi\rangle = 0, \quad (48)$$

must be gauge invariant.

References

- [1] Y. Aharonov and D. Bohm. Significance of electromagnetic potentials in the quantum theory. *Physical Review*, 115:485, 1959.
- [2] R. Arnowitt, S. Deser, and C. W. Misner. The dynamics of general relativity. In L. Witten, editor, *Gravitation: an introduction to current research*, pages 227–265. Wiley, 1962.
- [3] A. Ashtekar. Asymptotic structure of the gravitational field at spatial infinity. In Held [38].
- [4] A. Ashtekar. *Lectures in Non-Perturbative Canonical Gravity*. World Scientific, 1991.
- [5] A. Ashtekar. Overview and outlook. In J. Ehlers and H. Friedrich, editors, *Canonical Gravity: from Classical to Quantum — proceedings of the 117th WE Heraeus Seminar held at Bad-Honnef, Germany, 13-17 September 1993*, Lecture Notes in Physics 434. Springer-Verlag, 1993.
- [6] A. Ashtekar and R. O. Hansen. A unified treatment of null and spatial infinity in general relativity, i. universal structure, asymptotic symmetries, and conserved quantities at spatial infinity. *Journal of Mathematical Physics*, 19:1542–1566, 1978.
- [7] J. Barbour. *The End of Time*. Weidenfeld and Nicholson, 1999.
- [8] J. B. Barbour. The timelessness of quantum gravity:i. the evidence from the classical theory. *Classical and Quantum Gravity*, 11:2853–2873, 1994.
- [9] J. B. Barbour. The timelessness of quantum gravity:ii. the appearance of dynamics in static configurations. *Classical and Quantum Gravity*, 11:2875–2897, 1994.
- [10] J. D. Barrow and F. J. Tipler. *The Anthropic Cosmological Principle*. Oxford University Press, 1986.
- [11] V. B. Berestetskii, E. M. Lifshitz, and P. Pitaevskii. *Quantum Electrodynamics*. Butterworth Heinemann, 1982.
- [12] J. J. Binney, N. J. Dowrick, A. J. Fisher, and M. E. J. Newman. *The Theory of Critical Phenomena : an introduction to the renormalisation group*. Oxford University Press, 1992.
- [13] S. Carlip. *Quantum Gravity in 2+1 Dimensions*. Cambridge University Press, 1998.
- [14] S. M. Carroll, W. H. Press, and E. L. Turner. The cosmological constant. *Annual Review of Astronomy and Astrophysics*, 30:499–542, 1992.
- [15] H. B. G. Casimir. *Proc. Kon. Ned. Akad. Wetenschap*, B51:793, 1948.

- [16] R. G. Chambers. *Physical Review Letters*, 53, 1960.
- [17] T.-P. Cheng and L.-F. Li. *Gauge theory of elementary particle physics*. Oxford University Press, 1984.
- [18] J. Christian. Private conversations, 1998.
- [19] M. Davis, G. Efstathiou, C. S. Frenk, and S. D. M. White. *Nature*, 292:371, 1985.
- [20] D. Deutsch. *The Fabric of Reality*. Penguin, 1997.
- [21] B. DeWitt. Quantum theory of gravity, i. *Physical Review*, 160:1113–1148, 1967.
- [22] B. DeWitt. Quantum theory of gravity, ii: the manifestly covariant theory. *Physical Review*, 162:1195–1239, 1967.
- [23] B. DeWitt. Quantum theory of gravity, iii: application of the covariant theory. *Physical Review*, 162:1239–1256, 1967.
- [24] J. Earman. *World Enough and Space-Time*. MIT Press, 1989.
- [25] J. Earman. *Bangs, Crunches, Whimpers and Shrieks: Singularities and Acausalities in Relativistic Spacetimes*. Oxford University Press, 1995.
- [26] H. Everett III. Relative state formulation of quantum mechanics. *Review of Modern Physics*, 29:454–462, 1957. Reprinted in DeWitt and Graham (eds.), *The Many-Worlds Interpretation of Quantum Mechanics* (Princeton, 1973).
- [27] R. P. Feynmann and A. R. Hibbs. *Quantum Mechanics and Path Integrals*. McGraw-Hill, 1965.
- [28] A. E. Fischer and J. E. Marsden. The Einstein equations of motion — a geometric approach. *Journal of Mathematical Physics*, 13:546–568, 1972.
- [29] R. Gambini, B. Bruügman, and J. Pullin. Knot invariants as degenerate states of four dimensional quantum gravity. In S. Catto and A. Rocha, editors, *Proceedings of the Twentieth International Meeting on Differential Geometric Methods in Theoretical Physics*. World Scientific, 1992.
- [30] G. Gamow. *My World Line*. Viking, 1970.
- [31] M. Göckeler and T. Schücker. *Differential Geometry, Gauge Theories, and Gravity*. Cambridge University Press, 1987.
- [32] H. Goldstein. *Classical Mechanics*. Addison-Wesley, 1980.
- [33] S. Hawking. The path integral approach to quantum gravity. In Hawking and Israel [34].

- [34] S. Hawking and W. Israel, editors. *General Relativity: an Einstein Centenary Survey*. Cambridge University Press, 1979.
- [35] S. W. Hawking. Particle creation by black holes. *Communications in Mathematical Physics*, 43, 1975.
- [36] F. W. Hehl, J. Nitsch, and P. von der Heyde. Gravitation and the Poincaré gauge field theory with quadratic lagrangian. In Held [38].
- [37] F. W. Hehl, P. von der Heyde, and G. D. Kerlick. General relativity with spin and torsion: Foundations and prospects. *Reviews of Modern Physics*, 48:393–416, 1976.
- [38] A. Held, editor. *General Relativity and Gravitation: One Hundred years after the birth of Albert Einstein*. Plenum Press, 1980.
- [39] J. Isenberg and J. Nester. Canonical gravity. In Held [38].
- [40] C. Isham. Talk given to the Mathematics Institute, Oxford University, 1998.
- [41] C. Itzykson and J.-B. Zuber. *Quantum Field Theory*. McGraw-Hill, 1980.
- [42] J. D. Jackson. *Classical Electrodynamics*. Wiley, 1975.
- [43] M. Kaku. *Quantum Field Theory*. Oxford University Press, 1993.
- [44] J. Leslie. *Universes*. Routledge, 1989.
- [45] H.-J. Matschull. Dirac’s canonical quantization program. *arXiv*, quant-ph/9606031, 1996.
- [46] T. Maudlin. *Quantum non-locality and Relativity: Metaphysical Intimations of Modern Physics*. Blackwell, 1994.
- [47] C. W. Misner, K. S. Thorne, and J. A. Wheeler. *Gravitation*. W.H. Freeman and Company, 1973.
- [48] M. Nakahara. *Geometry, Topology and Physics*. IOP Publishing, 1990.
- [49] Y. Ne’eman. Gravity, groups and gauges. In Held [38].
- [50] I. Newton. *Philosophiae naturalis principia mathematica*. 1729. Translated by A. Motte and F. Cajori, University of California Press, 1962.
- [51] D. N. Page and W. K. Wootters. Evolution without evolution: Dynamics described by stationary observables. *Physical Review*, D27:2885–2892, 1983.
- [52] A. Pais. *Subtle is the Lord: the Science and the Life of Albert Einstein*. Oxford University Press, 1982.
- [53] P. J. E. Peebles. *Principles of Physical Cosmology*. Princeton University Press, 1993.

- [54] R. Penrose. Twistor theory, its aims and achievements. In C. J. Isham, R. Penrose, and D. Sciama, editors, *Quantum Gravity: an Oxford Symposium*. Oxford University Press, 1975.
- [55] R. Penrose. Singularities and time-asymmetry. In Hawking and Israel [34].
- [56] R. Penrose. *The Emperor's New Mind*. Vintage, 1990.
- [57] R. Penrose and W. Rindler. *Spinors and Space-time, volume 2: Spinor and Twistor Methods in Space-Time Geometry*. Cambridge University Press, 1986.
- [58] J. Pullin. Knot theory and quantum gravity in loop space: a primer. In J. L. Lucio, editor, *Proceedings of the V Mexican School of Particles and Fields*. World Scientific, 1993.
- [59] C. Rovelli. Time in quantum gravity: Physics beyond the Schrödinger regime. *Physical Review*, D43:442–456, 1991.
- [60] S. Saunders. Time, decoherence and quantum mechanics. *Synthese*, 102(2):235–266, 1995.
- [61] E. V. Shuryak. *The QCD vacuum, hadrons and the superdense matter*. World Scientific, 1988.
- [62] L. Smolin. Did the universe evolve? *Classical and Quantum Gravity*, 9:173–191, 1992.
- [63] L. Smolin. *The Life of the Cosmos*. Weidenfeld and Nicholson, 1997.
- [64] L. Smolin. Towards a background independent formulation of M theory. *arXiv*, hep-th/9898192, 1998.
- [65] A. Trautman. Fiber bundles, gauge fields, and gravitation. In Held [38].
- [66] W. G. Unruh and R. M. Wald. Time and the interpretation of quantum gravity. *Physical Review*, D40:2598–2614, 1989.
- [67] R. M. Wald. *General Relativity*. University of Chicago Press, 1984.
- [68] R. M. Wald. *Quantum Mechanics on Curved Space-Time and Black Hole Thermodynamics*. University of Chicago Press, 1994.
- [69] D. Wallace. Worlds in the Everett interpretation. Forthcoming, 2000.
- [70] S. Weinberg. *The Quantum Theory of Fields*. Cambridge University Press, 1995.
- [71] E. Witten. *Nuclear Physics*, B311:46, 1988.
- [72] N. M. J. Woodhouse. *Geometric Quantization*. Oxford University Press, 1991.

- [73] J. W. York, Jr. Gravitational degrees of freedom and the initial-value problem. *Physical Review Letters*, 26:1656–1658, 1971.
- [74] J. W. York, Jr. The role of conformal 3-geometry in the dynamics of gravitation. *Physical Review Letters*, 28:1082–1085, 1972.