Chapter 7
Loop Quantum Gravity

7.1 Introduction

LQG is, along with string theory, one of the best-established quantum gravity programs. Proponents of LQG hold that the most important lesson of GR is the diffeomorphism invariance of the gravitational field, and thus seek to preserve diffeomorphism invariance at the high-energy level of quantum gravity. Like the discrete approaches to quantum gravity discussed in the previous chapter, LQG is non-perturbative and researchers in LQG suggest that the problems with perturbative approaches (i.e. the problems associated with the non-renormalisability of gravity) may be a consequence of the failure of perturbation theory when applied at the scales being considered. And, like the discrete approaches, LQG describes the small-scale structure of spacetime as being discrete. The difference, though, is that some proponents of LQG claim that the discrete nature of spacetime is not postulated from the outset, as it is in the discrete approaches, but rather follows from the theory itself, as a prediction. However, it is not clear that this is indeed the case, since, as they stand, the discrete operators described by LQG are not physical observables.

This chapter is concerned with the conception of spacetime described by LQG. The most well-developed formulation of the theory is based on canonical quantum gravity and uses the Hamiltonian formalism, which (as described in Sect. 1.5) means that 4-dimensional spacetime is split into (3+1)-dimensions; the kinematics of the theory concerns primarily the microstructure of *space*, which is introduced in this section, and discussed in Sect. 7.2.1. The microstructure of spacetime will be discussed in Sect. 7.2.3. In Sect. 7.2.2, I consider the semiclassical limit and the recovery of (large-scale) spacetime. Finally, the idea of emergent spacetime in LQG is discussed in Sect. 7.3.

The birth of LQG is generally acknowledged as having occurred in 1987, and began when Ted Jacobson and Lee Smolin rewrote the Wheeler de-Witt equation (1.5) using Ashtekar variables, which Abhay Ashtekar (Ashtekar 1986, 1987) had used to construct a novel formulation of GR the year before, building upon the work of

K. Crowther, *Effective Spacetime*, DOI 10.1007/978-3-319-39508-1_7

Amitabha Sen.[1] Ashtekar variables are *connection* variables rather than metric ones and allow GR to be cast in a form similar to a Yang–Mills theory, and thus in a way that more closely resembles the standard model than it does otherwise.[2] Jacobson and Smolin (1988) discovered that, when rewritten using the Ashketar variables, the Wheeler–DeWitt equation has solutions that seem to describe loop excitations of the gravitational field.

As Carlo Rovelli (2004, pp. 15–16) points out, there is a natural old idea that a Yang–Mills theory is really a theory of loops: recalling Faraday's intuition that there are "lines of force" that connect two electric charges and which form closed loops in the absence of charges (the direction of the electric field at any point along such a line is given by the tangent vector at that point). More technically: the relevant mathematical quantity is the holonomy of the gauge potential along the line, and in LQG the holonomy is a quantum operator that creates "loop states". A loop state is one in which the field vanishes everywhere except along a single Faraday line.[3] In 1987, Rovelli and Smolin defined a theory of (canonical) quantum gravity in terms of loop variables. Doing so, they discovered that not only did the formerly intractable Wheeler–DeWitt equation become manageable and admit a large class of exact solutions, but that there were solutions to all the quantum constraint equations in terms of knot states (loop functionals that depend only on the knotting of the loops). In other words, knot states were proven to be exact physical states of quantum gravity (Rovelli and Smolin 1988, 1990).

Although the idea that loops are the appropriate variables for describing Yang–Mills fields may be a natural one, Rovelli (Rovelli 2004, 2008) explains that it was never able to be properly implemented except within lattice theories (QFT on a lattice).[4] One of the problems with using loops in a continuum theory is that loop states on a continuous background are over-abundant; a loop situated at one position on the background spacetime must be considered a different loop state from one that is positioned only an infinitesimal distance away, and so there are an infinite number of loop states on the continuum. Thus, the space spanned by the loop states is non-separable and therefore unsuitable for providing a basis of the Hilbert space of a QFT.[5]

Although it is not an obvious matter, Fairbairn and Rovelli (2004) argue that this problem does not arise for a background independent (diffeomorphism invariant) theory, such as GR. The argument is that, if we treat spacetime itself as made up of loops, then the position of a loop state of a QFT is relevant only with respect to other

[1]This discussion draws upon Ashtekar and Lewandowski (2004), Carlip (2001), Nicolai and Peeters (2007), Rovelli (2003, 2004, 2008, 2011), Wüthrich (Forthcoming).

[2]Since the standard model is a quantum Yang–Mills theory, meaning it has local (non-Abelian) gauge symmetry.

[3]The idea of loops and loop states will be discussed again below, and will hopefully become clearer by the end of this section.

[4]For example, Wilson loops, as a gauge-invariant observables obtained from the holonomy of the gauge connection around a given loop, were developed in the 1970s to study the strong interaction in QCD (after Wilson 1974). They now play an important role in lattice QCD.

[5]This will be discussed in more detail in Sect. 7.2.1.

loops, rather than the continuum background, and so there is no sense in saying that two loops are separated in spacetime. The idea is that an infinitesimal (coordinate) displacement will not produce a distinct physical state, but only a gauge equivalent representation of the same physical state. Therefore, the size of the state space is dramatically reduced by diffeomorphism invariance; only a finite displacement, which involves a loop being moved across another loop, will represent a physically different state. In the context of GR, the, loop states are thus (arguably) able to provide a basis of the Hilbert space, and the state space of LQG is a separable Hilbert space, $\mathcal{H}_K$, spanned by loop states (Rovelli 2004, p. 18).[6] Quantum states are represented in terms of their expansion on the loop basis, that is: as functions on a space of loops.

In a quantum theory, the discrete values of a physical quantity can be found by calculating the eigenvalues of its corresponding operator. In a theory of quantum gravity (where the gravitational field is identified with the geometry of spacetime), any quantity that depends on the metric becomes and operator, and it is by studying the spectral properties of these operators that we can learn about the quantum structure of spacetime. Most significant in LQG is the operator, $\hat{\mathbf{A}}$, associated with the area, $\mathbf{A}$, of a given surface, $\mathcal{S}$, and the operator, $\hat{\mathbf{V}}$, associated with the volume, $\mathbf{V}$, of a given spatial region, $\mathcal{R}$.

The area operator $\hat{\mathbf{A}}$ can be calculated by taking the standard expression for the area of a surface, replacing the metric with the appropriate function of the loop variables, and then promoting these loop variables to operators. An essentially similar procedure can be followed in order to construct $\hat{\mathbf{V}}$.[7] Both $\hat{\mathbf{A}}$ and $\hat{\mathbf{V}}$ are mathematically well defined self-adjoint operators in the kinematical Hilbert space $\mathcal{H}_K$; their spectra, first derived in 1994, were found to be discrete (Rovelli and Smolin 1995a). For instance, the spectrum for $\hat{\mathbf{A}}$ is given by:

$$\mathbf{A} = 8\pi\gamma\hbar G \sum_i \sqrt{j_i(j_i + 1)} \tag{7.1}$$

Where $i = 1, \ldots, n$, so that j is an n-tuplet of half-integers, labelling the eigenvalues, and γ is the Immirzi parameter, which is a free dimensionless constant (i.e. not determined by the theory).

The discrete spectra of the area and volume operators implies that the gravitational field is quantised. These quanta of space may be thought of as "chunks" of space, of definite volume given by the eigenvalues of $\hat{\mathbf{V}}$. Each chunk (or region, $\mathcal{R}$) can be thought of as bounded by a surface: if two chunks are adjacent to one another (i.e. direct neighbours), then the part of the surface that separates them (i.e. the fence that lies between the two neighbours) is $\mathcal{S}$, of area given by the eigenvalues of $\hat{\mathbf{A}}$. This idea is illustrated in Fig. 7.1, where the grey blobs represent the chunks of space.

---

[6] The subscript K is used in order to signal that this is the *kinematical* Hilbert space of the theory.

[7] More precisely: the construction of the area operator first requires the classical expression be regularised, then the limit of a sequence of operators, in a suitable operator topology, be taken. Both $\hat{\mathbf{A}}$ and $\hat{\mathbf{V}}$ have been derived several times using different regularisation techniques (e.g. Ashtekar and Lewandowski 1997a, b; Frittelli et al. 1996; Loll 1995, 1996; Rovelli and Smolin 1995a).

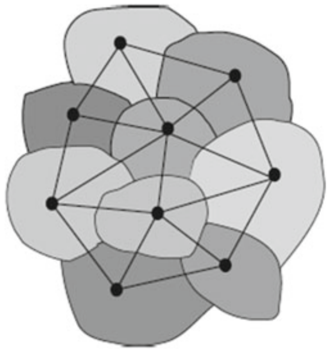

Fig. 7.1 Quanta of volume (*grey blobs*). Adjacent chunks are separated by a surface $\mathcal{S}$ of quantised area. The corresponding spin network graph is overlaid. Each link "cuts" one quantised surface $\mathcal{S}$

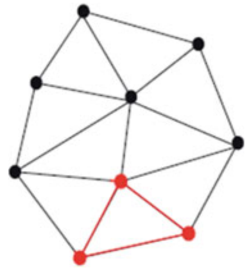

Fig. 7.2 Spin network: Nodes represent quanta of volume, which are adjacent if there is a link between them. Connected links form loops, like the one highlighted in *red/grey*

In LQG, this intuitive picture takes the form of abstract graphs called *spin networks*, in which each volume chunk is represented by a node (the black dots in each grey blob in Fig. 7.1), and each $\mathcal{S}$ separating two adjacent chunks is represented by a link (the lines joining the nodes). The spin network without the heuristic background illustration is shown in Fig. 7.2. This diagram also enables us to visualise the loops of LQG: they are the links that meet up to enclose white space, for instance the red loop highlighted.

An spin network graph, Γ with N nodes represents a quantum state of space, $|s\rangle$ formed by N quanta of space. The graph has each node n labelled i_n, which is the quantum number of the volume (i.e. the volume of the corresponding quanta, or chunk, of space), and each link l labelled j_l, which is the quantum number of the area (i.e. the quantised value of the area of $\mathcal{S}$ separating the two adjacent chunks of space being represented by those nodes being linked). The choice of labels is called the *colouring* of the graph. The area of a surface cutting n links of the spin network with labels j_i $(i = 1, \ldots, n)$ is given by the spectrum (7.1). The spin network s may thus be designated $s = (\Gamma, i_n, j_l)$ as shown in Fig. 7.3: these quantum numbers completely characterise and uniquely identify an spin network state. It is worth briefly mentioning (because it will be important in Sect. 7.2.1) that the labels j_l attached to the links are called *spins*, while the labels i_n are *intertwiners* associated to the nodes.[8]

[8]The meaning of these terms will not be discussed here, except to say that these quantum numbers are determined by the representation theory of the local gauge group, $SU(2)$. See Rovelli (2004, pp. 234–236).

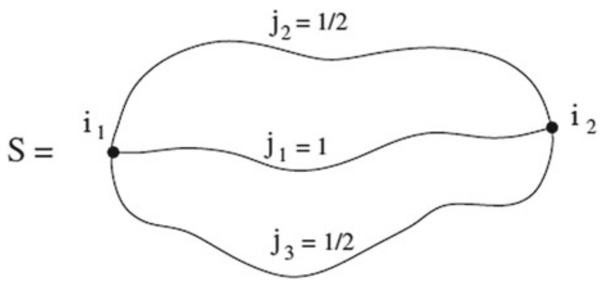

Fig. 7.3 A simple spin network with colouring. Labels i_n indicate the quantised volume of the corresponding node, and j_l give the quantised area represented by the corresponding link (Rovelli 2004, p. 19)

The spin network states $|s\rangle$ provide a basis for $\mathcal{H}$, and represent the general (unmeasured) quantum states of the gravitational field. Since (a region of) physical space is a state in $\mathcal{H}_K$, it is a quantum superposition of abstract spin network states. A loop state is a spin network state in which the graph Γ has no nodes, i.e. it is a single loop, and in such a state, the gravitational field has support only on the loop itself, with the direction of the field at any point along the loop given by the tangent vector at that point (again, recalling Faraday's intuition mentioned above). It is important to emphasise that spin networks are *abstract* graphs: spin network states are not quantum states of a physical system *in* space, rather they are the quantum states of physical space itself; only abstract combinatorial relations defining the graph are (physically) significant, not its shape or position in space.[9]

The significance of the spin network states as providing a basis for the Hilbert space of LQG was only realised in 1995 (Rovelli and Smolin 1995b). However, it turns out that spin networks were only "rediscovered" rather than invented in the context of LQG: the spin networks themselves had been created independently many years earlier, by Roger Penrose, based simply on what he imagined quantum space could look like (e.g. Penrose 1971). It was Penrose who named these graphs "spin networks", since their quantum numbers and their algebra resembled the spin angular momentum quantum numbers of elementary particles.

7.2 Spacetime in LQG

Recall that the canonical quantisation program (on which LQG is based) begins with canonical GR, which casts GR as a Hamiltonian system with constraints. The goal of the quantisation procedure is to find the Hilbert space corresponding to the physical

[9] When referring to "spin networks" I mean only the *abstract* graphs. Embedded (i.e. non-abstract) spin networks are of significance, and will be discussed in Sect. 7.2.1, where they will be explicitly referred to as embedded spin networks. The abstract spin network states $|s\rangle$ are equivalence classes under diffeomorphism invariance of the embedded spin networks, and are also known as s-knots.

state space of theory, and to define operators on the Hilbert space that represent the relevant physical quantities. In LQG, the procedure begins with a classical phase space coordinatised by the "holonomy" and its conjugate "flux" variable, which are constructed from the Ashtekar connection A_a^i and its conjugate, E_i^a a densitised triad "electric field".[10]

The geometrical structure of the classical phase space is encoded by the canonical algebra given by the Poisson brackets among these basic variables; in the quantisation, an initial functional Hilbert space of quantum states is defined, and the basic canonical variables are turned into operators whose algebra is determined by their commutation relations, which come from the classical Poisson brackets. These are then used in the construction of the constraints, which, in turn, serve to select a subset of states that correspond to the physical states of the theory.

In LQG, there are three types of constraint: the $SU(2)$ Gauss gauge constraints, which come about from expressing GR as an $SU(2)$ Yang–Mills theory, and are comparatively easy to solve; the spatial diffeomorphism constraints, which stem from diffeomorphism invariance and are hard to solve; and the Hamiltonian constraint (the general form of which is the Wheeler–DeWitt equation (1.5)), which has not yet been solved. In fact, there are many different forms of the Wheeler–DeWitt equation in LQG, and it is not clear which—if any—is correct. Solving the Gauss constraints and the diffeomorphism constraints gives us the kinematical Hilbert space $\mathcal{H}_K$ (i.e. this is the Hilbert space we obtain from the states which get annihilated by the Gauss and diffeomorphism constraints). The Hamiltonian constraint (its general form being like the Schrödinger equation), represents the dynamics of the theory.

Because of the technical and conceptual difficulties with the Hamiltonian constraint equation, proponents of LQG have sought alternative ways of understanding the dynamics of LQG. Here, I will focus on the idea of treating spin networks as "initial" and "final" states, and the dynamics of the theory being determined by the transition probability amplitudes, $W(s)$, between them, i.e. by taking the sum-over-histories approach. This represents a *covariant* formulation of LQG as opposed to the canonical one (although, I should point out that the covariant representation can be approached from different starting points, and the one presented here stems from the canonical formulation).

The covariant formulation of LQG (also called "spin foam theory") is a relatively new area of research and is less-developed than the canonical formulation of LQG. The aim of the formalism is to provide a means of calculating the transition amplitudes in LQG: it does this as a sum-over-histories, where the "histories" being summed-over are known as *spin foams*. A spin foam can be thought of as a world-history of a spin network, and represents a "spacetime" in the way that a spin network represents a "space". These ideas will be discussed in Sect. 7.2.3.

[10] Where $i = 1, 2, 3$ are "internal" indices that label the three axes of a local triad, and $a = 1, 2, 3$ are spatial indices. A densitised electric field has $\rho(E) = 1$.

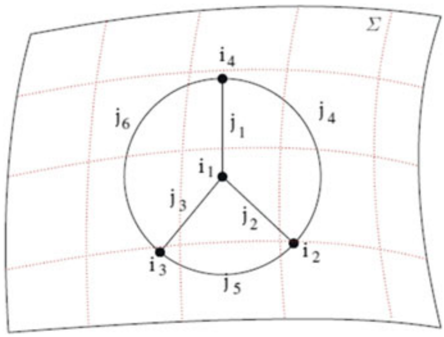

Fig. 7.4 A spin network embedded on the spatial hypersurface Σ (adapted from Nicolai and Peeters 2007, p. 156)

7.2.1 Micro-structure of Space: Spin Networks

Spin networks do not start out as abstract graphs: rather, LQG begins with a three-dimensional spatial manifold, Σ, on which the holonomies and spin networks are defined. The manifold is used to label the positions of the vertices and edges with coordinates; embedded spin network states are designated $|S\rangle$ (i.e. with a capital S rather than the lowercase s of the abstract spin network states). An embedded spin network is shown in Fig. 7.4.

The spin network wave functions only "probe" the geometry on one-dimensional sub-manifolds (i.e. along the one-dimensional edges), and are insensitive to the geometry elsewhere on Σ. For any two embedded spin network states, the scalar product is defined as,

$$\langle \Psi_{\Gamma, j_l, i_n} | \Psi'_{\Gamma', j'_l, i'_n} \rangle = \begin{cases} 0, & \text{if } \Gamma \neq \Gamma' \\ \int \prod_{j_l \in \Gamma} \mathrm{d}h_{j_l} \; \bar{\psi}_{\Gamma, j_l, i_n} \psi'_{\Gamma', j'_l, i'_n}, & \text{if } \Gamma = \Gamma' \end{cases} \tag{7.2}$$

where Ψ are the spin network wave functions, Γ the spin network graphs, j_l are the spins attached to the edges (links) and i_n the intertwiners associated to the nodes of the graph.

The *pre-kinematical Hilbert space*, $\mathcal{H}_{K*}$ is defined using the scalar product (7.2), which induces a peculiar discretisation (one entirely different from the discreteness of a lattice or the naive discretisation of space).[11] The resulting topology is similar to the discrete topology of the real line with countable unions of points as the open sets. Because the only notion of "closeness" between two points in this topology is whether or not they are coincident, *any* function is continuous in it.[12] Thus, it is already difficult to see how it would be possible to recover any conventional notion of continuity in LQG. The effect of the scalar product (7.2) means that non-coincident states are orthogonal, and the expectation values of operators that depend on some

[11] The full kinematical Hilbert space, $\mathcal{H}_K$ (which is separable) is defined once the diffeomorphism constraint has been implemented, as will be explained below.

[12] Nicolai and Peeters (2007), p. 156.

parameter do not vary continuously as the parameters upon which they depend are continuously varied.

This leads to the traditional problem mentioned above, of the spin network basis being "too large": any operation which moves the graphs around continuously corresponds to an uncountable sequence of mutually orthogonal states in $\mathcal{H}_{K^*}$. The Hilbert space does not admit a countable basis and so is non-separable. No matter how "small" the deformation of the graph in Σ, the associated elements of $\mathcal{H}_{K^*}$ always remain a finite distance apart—this means that continuous motion in "real" space gets mapped to highly discontinuous motion in $\mathcal{H}_{K^*}$.

The separable kinematical state space, $\mathcal{H}_K$, of the theory comes about once the diffeomorphism constraint has been implemented. This involves factoring out the gauge equivalent loop representations according to diffeomorphism invariance (i.e. any two graphs can be deformed into one another): the uncountable, "too large" basis is reduced once gauge redundancy is taken into account, and the non-separable Hilbert space becomes separable (Rovelli 2008). As mentioned in the introduction to this chapter, the loop representation of GR (as a Yang–Mills field theory) is only of value because of the theory's diffeomorphism invariance. Implementing diffeomorphism invariance removes the significance of the manifold Σ, and the natural basis states are abstract spin network states (also called *s*-knots), which are equivalence classes (under diffeomorphism invariance) of the embedded spin networks. Construction of the constraints makes use of the operators corresponding to the relevant physical observables.

The important operators are $\hat{\mathbf{A}}$, which measures the area of a two-dimensional surface, $S \subset \Sigma$, and $\hat{\mathbf{V}}$, which measures the volume of a three-dimensional subset of Σ. These operators, however, *cannot* be classed as physical observables, since they do not commute with the constraints. In particular, being defined for surfaces and regions on Σ, the area and volume operators are not invariant under the transformations generated by the diffeomorphism constraint. Because $\hat{\mathbf{A}}$ and $\hat{\mathbf{V}}$ are not diffeomorphism invariant, Rickles (2005, p. 423) argues that these quantities are not measurable unless they are *gauge fixed* and taken to correspond to the area and volume of some *physically defined* surface or region. Nicolai and Peeters (2007, p. 157) make a similar point, though emphasising the failure of the operators to commute with the Hamiltonian, and the necessity of the inclusion of matter in defining a physical surface or region.

Dittrich and Thiemann (2009) argue that if the discrete spectra of the area and volume operators are to be taken as representing physical discreteness of geometry, then we need to investigate whether the kinematical discreteness of the spectra survives at the gauge-invariant level once the operators have been "turned into" gauge-invariant quantities. This process depends on the mechanism used in order to turn the gauge-dependent operators into gauge-invariant ones, as well as on the interpretation of generally covariant quantum theory—a point that Rovelli (2007) emphasises. In standard quantum theory, a quantity is predicted to have discrete values if the corresponding quantum operator has a discrete spectrum. Whether this idea carries over to generally-covariant quantum field theory—a framework that is not fully developed or understood, but which LQG is modelled on—is equivocal because the distinction between the kinematics and the dynamics of such a theory is not clear cut (Rovelli 2007, p. 1).

The debate between Dittrich and Thiemann (2009) and Rovelli (2007) demonstrates that the discrete spectra of $\hat{\mathbf{A}}$ and $\hat{\mathbf{V}}$ do not necessarily represent physical discreteness of Planck-scale geometry. Dittrich and Thiemann (2009) present a simplified, non-LQG (not quantum gravitational) case-study in which the discreteness of the kinematical operators' spectra does not survive once the operators have been turned into gauge-invariant ones. Rovelli (2007) argues that not only is this example not analogous to the case in LQG, but that Dittrich and Thiemann utilise a particular interpretation of generally covariant quantum theory—one that Rovelli does not endorse.

On this interpretation, it is possible in principle for a "partial observable", $\hat{f}$, representing the area of a coordinate surface, to have a discrete spectrum, while the corresponding "complete observable", $\hat{F}$, representing the area of a physically defined surface, have continuous spectrum. Thus, on this interpretation, there is no "guarantee" that the discrete spectra of the operators of LQG, understood as partial observables, indicates physical discreteness of spacetime. On a different interpretation (the one that Rovelli recommends), the physical quantisation depends on the spectra of the kinematical operators in $\mathcal{H}_{K^*}$, and so it "follows immediately" that a physical measurement of the area or volume operators of LQG would yield a discrete value.

What is needed now, if we follow Rovelli's line of reasoning, is an examination of the motivations for preferring one interpretation over the other. Unfortunately, I cannot undertake this here. It is enough to say that it is at least plausible that the discrete spectra of the operators represent physical predictions of LQG.[13] If the spectra of $\hat{\mathbf{V}}$ and $\hat{\mathbf{A}}$ are physical predictions, then, according to LQG, space itself is discrete and combinatorial, and, because the theory has a natural cutoff at the Planck scale, there are no ultraviolet divergences. On the other hand, if the discrete spectra of the kinematical operators is not physical, then they might be understood in a similar way to the discrete elements in CDT: part of the formalism of the theory, but not themselves evidence of spacetime discreteness. In any case, the structures described by LQG are very different from our familiar conception of space.

Additionally, there is another, perhaps even stronger, way in which the spin networks differ from the emergent spacetime they are supposed to underlie: the fundamental relation of *adjacency* which is meant to correspond to the notion of two

[13] Even granting this, though, we should consider the meaning of a prediction that cannot be tested—if the discrete spectra are physical predictions, then LQG states that *if* we were able to probe extremely high energy scales *and* had a means of detecting the discrete structure of spacetime itself at these energies, then we would find spacetime to be discrete. Needless to say, this is a long shot. Also, we must remember that, because LQG is not based on any physical data, the prediction is a consequence of a particular combination of principles and assumptions (Crowther and Rickles 2014). Because of this, the distinction between *postulated* versus *predicted* discreteness is perhaps not an interesting one to push (as we might be tempted to do in comparing LQG with the "discrete" approaches, where discreteness is explicitly acknowledged as an assumption or principle of the theory).

objects being "nearby" to one another in LQG does not, typically, translate into this notion in the emergent spacetime. Recall that two nodes being linked by an edge in a graph represents two adjacent quanta of space. The idea of spacetime emerging from the more basic spin networks means that there is a "mapping" of spin network nodes onto events in the emergent spacetime.

However, two nodes that are adjacent in the basic (high-energy) description can be arbitrarily large distances away from one another as measured in the emergent metric—in other words, they will, in general, not be mapped to "nearby points" in the emergent spacetime.[14] The fact that the adjacency relations described by the fundamental spin networks do not typically feature in the emergent spacetime (i.e. are not "translated" into the corresponding spatiotemporal relations at low-energy) means that many of them (i.e. all those adjacencies which do not get translated into Planck-sized neighbourhoods in the spacetime) are suppressed at low-energy.

7.2.2 Semiclassical Limit: Weaves

Finding the low-energy limit of LQG has proven very difficult, and all attempts to recover GR from LQG have so far been unsuccessful. One obvious handicap is the fact that all such attempts have been confined to working with the kinematical Hilbert space $\mathcal{H}_K$, rather than the physical Hilbert space of the theory. Thus, there are questions regarding both the viability and the meaningfulness of relating the kinematical states to corresponding classical spacetimes (or spaces).

The most prominent of the attempts to construct semiclassical states (i.e. states in which the quantum fluctuations are minimal and the gravitational field behaves almost classically) is based on *weave states*, which were first introduced by Ashtekar et al. (1992).[15] The intuitive idea is captured by analogy: at familiar scales, the fabric of a t-shirt is a smooth, two-dimensional curved surface, but when we examine it more closely, we see that the fabric is composed of one-dimensional threads woven together.[16] The suggestion is that LQG presents a similar picture: while spacetime at large-scales has a continuous geometry, at high-energy it is revealed to be a spin network with an enormous number of nodes and links.[17]

[14] Huggett and Wüthrich (2013), Wüthrich (Forthcoming) both emphasise this point and the associated problems for our understanding of locality.

[15] Alternative methods aiming to overcome the shortcomings of weave states have, and are still, being explored; for a discussion, see Thiemann (2001, Sect. II.3) or Thiemann (2007, Sect. 11).

[16] The analogy comes from Ashtekar et al. (1992).

[17] Sometimes, the spin network is described as a very large "lattice" of Planck-sized spacing. This is misleading for a couple of reasons; firstly, it implies that the Planck-scale structure has more regularity than it actually does (it does not resemble a regularly-spaced grid, for instance). Secondly, the spin network in LQG is different to lattice QFT, or to causal set theory, where the continuum is replaced by a lattice structure. The manifold is not modified in LQG; rather, the diffeomorphism invariance of the theory is supposed to mean that the continuum structure is not physically significant. Thanks to Joshua Norton for emphasising this point.

Consider a classical three-dimensional gravitational field e, which determines a three-dimensional metric $g_{ab}(\vec{x}) = e^i_a(\vec{x})e_{ib}(\vec{x})$, and a macroscopic three-dimensional region $\mathcal{R}$ of spacetime with this metric, bounded by the two-dimensional surface $\mathcal{S}$ (the values of area and volume being large compared to the Planck scale). It is possible to construct an (embedded) spin network state $|S\rangle$ that approximates this metric at a length scale $\Delta \gg l_P$, where l_P is the Planck length. To do this involves selecting spin network states that are eigenstates of the volume and area operators for the region $\mathcal{R}$ and the surface $\mathcal{S}$ with eigenvalues that approximate the corresponding classical values for the volume of $\mathcal{R}$ and area of $\mathcal{S}$ as given by e. The classical value for the area $\mathbf{A}$ of a surface $\mathcal{S} \subset \mathcal{M}$ and the classical value for the volume of a region $\mathcal{R} \subset \mathcal{M}$ with respect to a fiducial gravitational field c^i_a are given by,[18]

$$\mathbf{A}[e, \mathcal{S}] = \int |d^2\mathcal{S}| \tag{7.3}$$

$$\mathbf{V}[e, \mathcal{R}] = \int |d^3\mathcal{R}| \tag{7.4}$$

where the relevant measures for the integrals are determined by c^i_a.

Now, we require that $|S\rangle$ is an eigenstate of $\hat{\mathbf{A}}$ and $\hat{\mathbf{V}}$, with eigenvalues given by (7.3) and (7.4), respectively, up to small corrections of the of l_P/Δ,

$$\hat{\mathbf{A}}(\mathcal{S})|S\rangle = (\mathbf{A}[e, \mathcal{S}] + \mathcal{O}(l_P^2/\Delta^2))|S\rangle \tag{7.5}$$

$$\hat{\mathbf{V}}(\mathcal{R})|S\rangle = (\mathbf{V}[e, \mathcal{S}] + \mathcal{O}(l_P^3/\Delta^3))|S\rangle \tag{7.6}$$

If an embedded spin network state $|S\rangle$ satisfies (7.5) and (7.6), then it is a *weave state* of the metric g_{ab}.[19] At length scales of order Δ or larger, the weave state is a good approximation to the corresponding classical geometry, as $|S\rangle$ determines the same volumes and areas as g_{ab}. At length scales much smaller than Δ, however, the quantum fluctuations become relevant, and the weave state can no longer be considered a valid semiclassical approximation.

Finally, it is worth pointing out that (7.5) and (7.6) do not determine the state $|S\rangle$ uniquely for a given three-metric g_{ab}. This is because (7.5) and (7.6) involve quantities that are averaged over the macroscopic surface $\mathcal{S}$ and region $\mathcal{R}$. There are many different spin network states that can represent these averaged values, whereas there is only one classical metric that corresponds to these values. In this

---

[18]This presentation is based on Wüthrich (Forthcoming, p. 26).

[19]Although embedded, rather than abstract, spin network graphs are used, this definition of weave states is able to be carried over to the diffeomorphism-invariant level of abstract spin network states (s-knots) without issue. If we introduce a map $P_{Diff} : \mathcal{H}_(K^*) \to \mathcal{H}_K$, which projects states in the pre-kinematical Hilbert space into the same elements of the kinematical Hilbert space, then the state $|s\rangle = P_{Diff}|S\rangle$ is a weave state of the classical three-geometry $[g_{ab}]$, i.e. the equivalence class of three-metrics g_{ab}, just in case $|S\rangle$ is a weave state of the classical three-metric g_{ab} (Wüthrich Forthcoming).

way, the actual microstate of a given macroscopic geometry is *underdetermined* by the Eqs. (7.5) and (7.6). However, the generic quantum state of the macroscopic spacetime is not supposed to be a weave state, but a superposition of weave states.

7.2.3 Micro-structure of Spacetime: Spin Foams

The preceding two sub-sections have described only the kinematics of the theory: since the spin networks are based in the kinematical Hilbert space rather than the physical Hilbert space, they represent microstates of *space* rather than *spacetime*. There is a covariant version of LQG which aims to discover the dynamics of the theory without engaging with the Hamiltonian constraint of canonical LQG. This formulation, which is known as *spin foam theory* describes the micro-structure of spacetime as a *spin foam*, which is a history of spin networks. Presently, this theory is also incomplete and its relation to canonical LQG not known. It should also be noted that the sum-over-histories approach described in this sub-section is not representative of the full covariant LQG program, in that its starting-point draws from the concepts and results of canonical LQG.

Recall that in quantum mechanics, a complete description of the dynamics of a particle is provided by the transition probability amplitudes, A, defined as,

$$A = \langle \psi' | e^{\frac{i}{\hbar} H_0 (t-t')} | \psi \rangle \tag{7.7}$$

where $|\psi\rangle$ is the initial quantum state prepared at t, and $|\psi'\rangle$ is the final state of the system, measured at t', and H_0 is the Hamiltonian operator. Following Feynman, this amplitude can be calculated as a sum-over-paths between the "initial" and "final" states. The same is true in LQG, where the dynamics of the theory may be described entirely by the spin network transition amplitudes $W(s', s)$, governed by the Hamiltonian operator H, which is defined on the space of the spin networks. The space of solutions of the Wheeler–DeWitt equation is the *physical Hilbert space*, denoted $\mathcal{H}$. There is an operator $P : \mathcal{H}_K \rightarrow \mathcal{H}$ that projects $\mathcal{H}_K$ on the space of solutions of the Wheeler–DeWitt equation.[20] The transition amplitudes between an "initial" spin network state $|s'\rangle$ and the "final" spin network state $|s\rangle$ (recalling that there is no external time variable in the theory) are the matrix elements of the operator P,

$$W(s, s') = \langle s | P | s' \rangle_{Hilb_K} = \langle s | s' \rangle_{\mathcal{H}} \tag{7.8}$$

The Hamiltonian operator H (in all its different versions in LQG) acts only on the nodes of the spin network graph; in the vicinity of a node, the action of H upon a generic spin network state $|s\rangle$, is to change the topology and labels of the graph. Typically, H splits a node into three nodes and multiplies the state by a number a that depends on the labels of the spin network around the node. This is illustrated in Fig. 7.5.

[20] Rovelli (2004), Sect. 1.2.3.

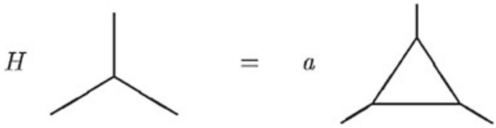

Fig. 7.5 Scheme of the action of H on a node of a spin network

Fig. 7.6 "Lime" rock candy: A slice width-wise reveals a cross-section with a lime-shaped pattern, which represents a spin network state s_1. The face at the "bottom" of the stick is s', and the one at the top is s. The whole length of the stick represents a history of spin networks, i.e. a spin foam, $\sigma = (s, s_N, \ldots, s_1, s')$ (Of course, every slice of rock candy will reveal a cross-section of essentially the same pattern, whereas "slices" of a spin foam would reveal different shaped spin networks, since the spin networks are transformed under the action of H)

The transition amplitude $W(s, s')$ can be represented as a sum-over-histories: a representation that follows from summing over different histories of sequences of actions of H that send s' to s.[21] The histories (of spin networks) being summed over are spin foams; a history of going from s' to s is a spin foam, σ, bounded by s' and s. The heuristic way to picture a spin foam is to imagine a 4-d spacetime in which the graph of a spin network s is embedded. If this graph moves along "upwards" through the "time" coordinate of the 4-d spacetime, then it "sweeps out" a "worldvolume" (a 3-dimensional version of a worldline). Actually, it is sort of like rock candy, shown in the picture, Fig. 7.6.

More schematically, Fig. 7.7 illustrates the worldsheet of a spin network that's shaped like θ. The surfaces traced out by the links of the spin network graph are called *faces*; the worldlines traced out by the nodes of the spin network graph are called *edges*. A spin foam, σ also includes a colouring, where faces are labelled by the area quantum numbers j_l and edges are labelled by the volume quantum numbers i_n.

[21] Although Rovelli (2004, p. 26) states that this is just one of several ways to arrive at the sum-over-histories representation of W.

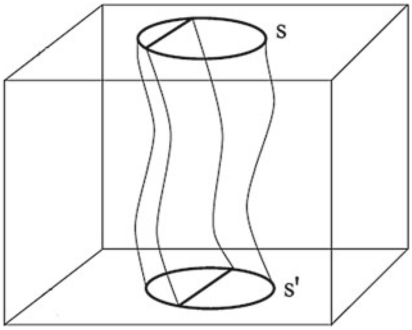

Fig. 7.7 A simple spinfoam: the worldsheet of a spin network ("colouring" of the faces and edges not indicated) (Adapted from Rovelli 2004, p. 325)

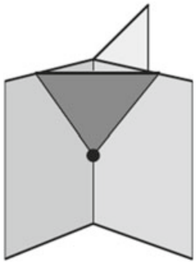

Fig. 7.8 Vertex of a spin foam (Rovelli 2004, p. 325)

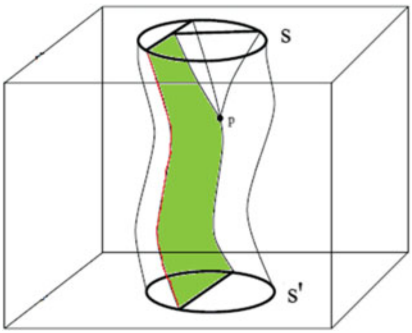

Fig. 7.9 Spinfoam with one vertex. One *face* has been *shaded green/grey*, and one *edge* is shown in *red/grey* (Adapted from Rovelli 2004, p. 326)

When H acts on a node, it results in the corresponding edge of the spin foam to branch off into three edges, in the "3-dimensional" version of the action shown in Fig. 7.5. The point where the edges branch is called a *vertex*, v, as shown in Fig. 7.8. A spin foam with a vertex is shown in Fig. 7.9.

A spin foam is a Feynman graph of spin networks. However, it differs from usual Feynman graphs in that it has one additional structure: while Feynman graphs have edges and vertices, spin foams have edges, vertices and faces. In the perturbative expansion of $W(s, s')$, each spin foam σ, bounded by s' and s, is weighted by an amplitude which is given by (a measure term $\mu(\sigma)$-times) the product over the vertices, v of a vertex amplitude, $A_v(\sigma)$. The vertex amplitude is determined by the matrix elements of H between the incoming and outgoing spin networks and depends on the labels of the faces and the edges adjacent to the vertex (Rovelli 2004, pp. 26–27). This is analogous to the amplitude of a standard Feynman vertex, which is determined by the matrix element of the Hamiltonian between the incoming and outgoing states.

The sum-over-histories is thus,

$$W(s, s') = \sum_{\sigma} \mu(\sigma) \prod_{v} A_v(\sigma) \tag{7.9}$$

Just as Feynman's sum-over-histories can be interpreted as a sum over different possible classical paths that a particle might take between two points, so too the spin foam sum (7.9) can be interpreted as a sum over spacetimes. And yet, even though Feynman's sum-over-histories can be interpreted as a sum over different possible particle trajectories, we know (in quantum theory) that there are no classical trajectories. The sum-over-histories is not itself a history. Although a single spin foam can be thought of as representing a spacetime, the theory states that spacetime is not a single spin foam, but a sum over spin foams. For this reason, Rovelli (2004, p. 31) states that the notion of spacetime disappears in quantum gravity in the same way that the notion of a particle trajectory disappears in the quantum theory of a particle. This interpretation of spin foam theory may also be applicable to the discrete quantum gravity approaches which utilise a sum-over-histories.

We can distinguish between two types of discrete quantum gravity approaches which utilise a sum-over-histories (i.e. approaches which attempt to define a discretised path integral in quantum gravity). The first group represents those approaches where spacetime is approximated by a fixed number of simplices and the integration is performed over all edge lengths: quantum Regge calculus is an example. The second group represents those approaches, including causal dynamical triangulations, where the simplices are assigned fixed edge lengths, and the sum is taken over different triangulations while keeping the number of simplices fixed (thus changing the "shape" of the triangulation but not its "volume"). Spin foam theory falls into the former category, along with quantum Regge calculus, since, in the first step of the procedure—which is calculating the partition function for a given spin foam—all spins are summed over (for the given spin foam), but there is no addition, removal or replacement of edges, vertices or faces.

In the second step (which aims to recover continuous geometry), the sum is taken over all spin foams, as in (7.9). The method by which to perform this step is not formally agreed upon; Nicolai and Peeters (2007) suggests that one way of doing it would be to weight each spin foam term in the sum according to its "shape", in order to achieve formal independence of the triangulations. This would thus resonate with CDT (Sect. 6.6), where, recall, the spacetime obtained was autonomous from the 4-simplices used to approximate the path integral, and so Nicolai and Peeters

(2007) state that we might interpret such a spin foam model as a hybrid of the two classes just distinguished.[22]

A key difference between CDT and spin foam theory is the ontological interpretation of the two approaches. In spin foam theory, following LQG, the discrete elements described by the theory are interpreted realistically, as the ultimate constituents of spacetime.[23] Recall that in CDT, however, the 4-simplices are taken simply to be mathematical tools that aid in the regularisation procedure used to define the integral—at the end of calculations, the aim is to remove the discretisation and recover a continuum theory (as in lattice QFT). Hence, spacetime according to CDT (and quantum Regge calculus) is not fundamentally discrete.[24] The fact that the discrete elements of space are interpreted realistically in LQG is the reason why the “continuum limit” cannot be recovered as it is in the other theories: the “lattice spacing” cannot be taken to zero (as described above in Sect. 7.2.2).

7.3 Emergence

Although LQG is incomplete and its physical Hilbert space undefined, there are still some potential bases for a conception of emergence in the theory. Interestingly, there does not appear to be a notion of emergence that might be based on the idea of a limiting relation in the theory (or, at least, not based on a limiting relation alone). This is because macroscopic geometry is not recovered in the limit as the density of the weave of loops goes to infinity. Originally in LQG, before the spectra of the area and volume operators had been derived, the limiting procedure was thought to run analogously to that in conventional QFT, where a continuum theory is defined by taking the limit of a lattice theory, as the lattice spacing a goes to zero. However, when the limit of the “loop constant” in LQG (the constant which was believed to be analogous to the “lattice spacing”)[25] is taken to zero, there is no increase in the accuracy of the LQG approximation to macroscopic geometry.[26]

The reason the theory fails to approximate a smooth geometry in the limit as the loop constant goes to zero is that the physical density of the loops does not increase in this limit. Instead, what occurs is that the eigenvalues of the area and volume operators increase, meaning that the areas and volumes in the region being studied

[22] Note that, in spin foam theory, this method of constructing the sum would mean a sum over spin foams with different numbers of simplices and different edge lengths, which is not how it is done in CDT.

[23] Although, as described above, there may be reason to question this interpretation of LQG.

[24] Oriti (2014) describes some further differences between LQG, Regge calculus and CDT, and explores what they might teach us about the fundamental nature of space and time, in combination with the *group field theory* (GFT) approach. GFT is essentially similiar to LQG and spin foam models, but with a key advantage in its definition of the dynamics—in particular, it prescribes a strict means by which to calculate the weights for the terms in the path integral.

[25] Please refer to Footnote 17.

[26] Rovelli 2004, p. 269.

grow larger. In other words, the loop density remains constant because we look at greater volumes. If we believe that area and volume are quantised, then this result can be readily interpreted: there is a minimum size for the loops, and thus, a minimal physical scale. The theory refuses to approximate a smooth geometry as the loop constant is taken below l_P because—on this interpretation—there is no physical length scale below l_P. Thus, such a limit is unable to serve as the means by which to recover spacetime from its fundamental spin network structure. This seems to suggest that a conception of emergence based on the idea of a limiting relation (e.g. that of Butterfield 2011a, b) will not be applicable in LQG.

Wüthrich (Forthcoming), however, suggests that perhaps the failure of the limiting procedure to recover spacetime is due to the fact that it is only capable of representing one of the two necessary transitions involved in the recovery process. The two steps of the process are: firstly, an approximating procedure which turns the quantum states into semiclassical ones, and, secondly, the limiting procedure which relates the semiclassical states to the phase space of the classical (i.e. the "emergent") theory. The associated notion of emergence comes from Butterfield and Isham (1999, 2001), where a theory T_1 emerges from another theory, T_2, *iff* T_1 can be arrived at from T_2 by either a limiting procedure or an approximation procedure, or both.

A *limiting procedure* is defined as taking the mathematical limit of some physically relevant parameter(s) in the underlying theory T_2 in order to recover the emergent theory T_1. An *approximation procedure* is defined as the process of either neglecting some physical magnitudes, and justifying such neglect, or selecting a proper subset of states in the state space of the approximating theory, and justifying such selection, or both, in order to arrive at a theory whose values of physical quantities remain sufficiently close to those of the theory to be approximated (Butterfield and Isham 1999). In LQG, Wüthrich (Forthcoming) imagines the limiting procedure to be something akin to the one mentioned above, where the semiclassical weave states are mapped to classical spacetimes. The weave states themselves are supposed to be arrived at via an approximation procedure.

The method of constructing weave states, described above, fits the definition of an approximation procedure, since they must be carefully selected to include only those states which are peaked around the geometrical values (of area and volume) determined by the fiducial metric e^i_a. This can be achieved either by neglecting all those operators constructed from connection operators (since the "geometrical" eigenstates are maximally spread in these operators), or, if this cannot be justified (as the approximation procedure requires), then only the semiclassical states that are peaked in both the connection and the triad basis, and peaked in such a way that they approximate classical states, should be considered. This approximation procedure is taken to represent whatever the physical mechanism is that drives the quantum states to semiclassical ones. We can suppose that this physical mechanism (justifying this approximation procedure) is *decoherence*.[27] This approximation procedure is necessary in addition to the limiting procedure because no limiting procedure (not even the $\hbar \to 0$ limit) can resolve a quantum superposition into a classical state.

[27] Although, as should be clear, LQG makes no reference to any such mechanism.

The $\hbar \to 0$ limit is the "textbook" means of justifying the use (or explaining the success) of "older", classical, theories for large orbits and low energies. The fact that it is typically paired with the approximation procedure corresponding to decoherence suggests that the limiting procedure may be interpreted as representing the micro/macro transition (a rescaling of the theory) while decoherence represents the quantum/classical transition—both being necessary for an account of emergence (as discussed in Sect. 1.7). This interpretation seems to accord with the claim that the $\hbar \to 0$ limit has the $N \to \infty$ limit as a special case; while the $N \to \infty$ limit implies moving to a large system of many particles, the $\hbar \to 0$ limit means moving to a description where $\hbar$ can be treated as negligible. This is certainly what is being indicated by the weave analogy in LQG.

It is worth emphasising again that both the idea of decoherence and the $\hbar \to 0$ limit are *external* to the theory itself: they have been imposed in an attempt to have LQG match up with low-energy classical physics, including GR. Decoherence and the $\hbar \to 0$ limit are among the traditional means by which quantum theories are shown to "reduce" to classical physics (although, again, the former represents a physical mechanism, and its interpretation more controversial than the mathematical limit in this case), and it would be distressing if they did not work in the context of LQG, given that the theory itself does not offer any "natural" means by which to recover spacetime: LQG does not (on its own, without the additional assumption of decoherence) explain how or why some states (those that are able to be "mapped" to classical states) are "selected".[28]

However, this might be expected given that LQG is a quantisation of GR. Perhaps LQG should be considered as a sort of "stepping stone", offering us access to the information contained in quantised GR, without (uniquely) capturing the micro-dynamics. On such an interpretation, the recovery of large-scale physics might be less important than making predictions—an interpretation along the same lines as treating GR as an EFT, as in (Sect. 5.3). Leaving these concerns aside, though, we can begin to sketch how we might understand the emergence of spacetime from LQG, based on the implications of applying the approximation and limiting procedures. These procedures lead to an underdetermination of the "more basic" quantum description by the emergent classical one.

The idea of emergence that Wüthrich utilises is inspired by Landsman (2006), who argues that while neither the limiting procedure nor decoherence is sufficient, on its own, for understanding how the classical picture emerges from the quantum world, together these procedures indicate that it comes from ignoring certain states and observables in the quantum theory. "Thus the classical world is not created by observation (as Heisenberg once claimed), but rather by the lack of it" (Landsman 2006, p. 417). On this account, the classical realm is correlated with certain "classical" states and observables, those which are robust against coupling to the environment,

[28]As explained in the Introduction (Sect. 1.7.1), there is a strong possibility that quantum gravity will provide some insight into the physical means (mechanism) by which quantum superpositions are resolved into classical states (e.g. decoherence), even in familiar quantum mechanics. LQG does not, as it stands, offer such insight, even though this physical mechanism is likely to play a role in the emergence of spacetime from LQG.

and which "survive" the approximation procedure that eliminates the non-"classical" states and observables.

Hence, this means of possibly recovering spacetime from LQG embodies an idea of underdetermination, and this provides the basis for the *autonomy* of spacetime from LQG. If this conception of emergence does apply in LQG, then the classical spacetime that emerges from LQG will be independent of many of the micro-states described by the theory. Unfortunately, this claim is a very vague one since the relation between the spin network states, weave states and spacetime being utilised is, at this stage, only a crude sketch. Nevertheless, we can see that the macro physics is also significantly *novel* compared to the micro-theory, given the substantial differences between the fundamental structures of LQG and those of GR—in particular, the failure of the relation of adjacency to map onto the corresponding notion of "closeness" in the classical geometry (as described in Sect. 7.2.1).

The failure of this relation to translate properly into the emergent spacetime suggests another possible conception of emergence associated with the idea of an approximation procedure. Because the spin networks generically give rise to geometries in which the notion of adjacency is not respected (from the point of view of the spin network), all those spin networks which correspond to geometries in which the spatial counterparts of two adjacent nodes are separated by more than a Planck length must be suppressed. In other words, we must select "classical" spin networks and ignore the rest, as accords with the definition of an approximating procedure, so long as some physical justification is provided for the neglect. Thus, it seems as though understanding this procedure and the justification for it could potentially lead to another notion of emergence in LQG.[29]

Finally, the weave states furnish yet another possible basis for emergence, where the relevant conception of emergence resembles that associated with thermodynamics, and is again related to the idea of underdetermination. The underdetermination comes about because the Eqs. (7.5) and (7.6) do not uniquely determine a weave state for a given metric. The reason for this, recall (Sect. 7.2.2), is that the equations only utilise *averaged* properties, which could be represented by a number of different microstates. This is similar to the case of thermodynamics, where an averaged, macroscopic property, such as temperature, will correspond to many different microstates in a system with a large number of micro-level degrees of freedom.

Hence, this account seems to accord with the picture of emergence associated with thermodynamics, Sect. 4.12, although there are two differences of potential relevance: firstly, emergence in thermodynamics can be connected to the idea of universality. Because there is no physical Hilbert space in LQG, it is not clear that the notion of universality makes sense within the theory (at this stage). Secondly, macroscopic geometry is supposed to correspond to a superposition of weave states (Sect. 7.2.2), whereas the macroscopic variables of thermodynamics are not typically taken to correspond to superpositions of micro-states.

[29]Recall that the idea of emergence associated with neglecting certain states has also been proposed by Bain (2013), in the context of EFT. This is discussed in Sect. 3.9, where I also tie it to the idea of underdetermination.

Most of the potential bases for a conception of emergence in LQG that have been presented here utilise the idea of underdetermination (as providing an explanation for spacetime being largely autonomous of its micro-structure). The suggestion that the micro-structure of spacetime is a superposition of microstates—which is made not only in regards to the weave states in LQG, but also in spinfoam theory (Sect. 7.2.3)—raises some interesting questions in regards to how the classical idea of underdetermination corresponds to quantum indeterminacy. This is another point where the quantum/classical transition intersects with the micro/macro transition, and perhaps the idea of decoherence will be of some help.

Given the substantial differences between LQG and GR, and the absence of any limiting procedure linking the two theories, it might seem pointless to attempt to frame a conception of emergence with GR taken to be emergent from LQG. The idea of emergence related to underdetermination that has been presented here perhaps fuels this worry—on it, *any* low-energy theory might be said to be emergent from LQG in the same way that GR is supposed to be. Of course, the only reply is that LQG is a direct quantisation of GR. It is hoped (or assumed) that GR must somehow be emergent from LQG because we are able to "go the other way" and arrive at LQG from GR. The aim of the project is not to recover some other spacetime theory from LQG, but to approximate GR in the regime where the accuracy of the latter theory has been proven. Here, it is worth pointing out, however, that a theory of quantum gravity need not be a quantisation of GR[30]: conversely, a quantisation of GR does not necessarily produce a theory of quantum gravity.

7.4 Conclusion

LQG and spin foam theory are incomplete, with no definite means by which to describe the dynamics in either theory: there are several alternative Hamiltonian operators in LQG, and a number of options for calculating the measures in the spin foam sum, but no indication that any choice is correct in either of the cases. This incompleteness means that we are unable to develop a concrete picture of how spacetime could emerge from LQG. It is likely that such a picture would involve both the micro/macro transition as well as the quantum/classical transition, where the latter perhaps will need to be understood before the former can be implemented—nevertheless, I have purposefully avoided engaging with questions related to the quantum/classical transition and the idea of decoherence here. It may be that the failure of the "continuum limit" of LQG is a consequence of not yet understanding the role of decoherence.

Because the limit in which the density of the weave states goes to infinity fails to approximate continuous geometry, it is unclear how a conception of emergence based on the idea of a limiting relation, such as described by Butterfield (2011a, b),

[30]It could be a quantisation of a theory other than GR, or it might not be a quantisation of any classical theory.

could apply in LQG. This could potentially be problematic, since (as described in Sects. 1.6 and 2.3) a limiting relation is a common means by which a newer theory is shown to relate to the older theory it is supposed to supplant, and the demonstration of the recovery of GR (which may be done through the use of a limiting relation) is generally taken as necessary for a theory of quantum gravity. An RG scaling procedure is also unable to be implemented as a means of recovering spacetime at large-distances, since the physical Hilbert space of LQG is undefined.

Although the area and volume operators of LQG have discrete spectra, the fact that they are not gauge invariant—only existing at the kinematical level—means we cannot say definitely that LQG predicts spacetime discreteness. Additionally, LQG is faced with the problem of time and the problem of space, which are also related to the difficulties of interpreting gauge invariance (amplified by tensions between quantum theory and GR, Sect. 1.5). It may be that the problems with LQG have to do with the fact that it is a quantisation of GR. This fact also makes the lack of a low-energy limit of the theory particularly worrisome—given that LQG is a quantisation of GR, we would expect it to be relatively easy[31] to recover GR through decoherence plus a semiclassical limit. This having not been done is perhaps further motivation for considering an approach to quantum gravity that does not have a quantisation of GR as its starting point.

In spite of these difficulties, a number of potential bases for emergence can be identified in LQG. The requisite criterion of *novelty* is fulfilled, since the macro-structures of GR differ in several major ways from the micro-structures described by LQG. Not only is the discreteness of the spin networks (and spin foams) of LQG a departure from the structures of GR, but the generic micro-state of spacetime is not supposed to even be a single spin network (spin foam)—rather a quantum superposition of such states. Furthermore, the fundamental relation of adjacency in a spin network, which indicates that two "quanta of space" are next to one another, is not typically preserved in the macroscopic geometry that the spin network is supposed to underlie: two adjacent quanta of space in a spin network may be arbitrarily far away from one another in the emergent space.

The idea of *autonomy* that forms the second part of the account of emergence presented here (Sect. 2.4) is furnished primarily by the underdetermination of the micro-states given the macro-states. For instance, while a weave state is constructed so as to represent the micro-state of a given (macroscopic, three-dimensional) metric, the construction of weave states means that, for a given metric there is a multitude of potential micro-states.[32] In other words, the emergent spacetime is able to be represented by a number of different spin networks, and so is independent of the details of the micro-theory. The reason for this is the fact that the weave states depend only on average values, and so this idea of emergence resembles that associated with

[31] Well, compared to other approaches!

[32] More correctly: a weave state is constructed so as to demonstrate the existence of a micro-state for a given macro-state.

hydrodynamics/thermodynamics, where a number of micro-states correspond to the same macro-state of a system. The emergent structures depend only on *collective properties* of the micro-degrees of freedom.

References

Ashtekar, A. (1986). New variables for classical and quantum gravity. *Physical Review Letters*, *57*, 2244–2247.

Ashtekar, A. (1987). New Hamiltonian formulation of general relativity. *Physical Review D*, *36*(6), 1587–1602.

Ashtekar, A., & Lewandowski, J. (1997a). Quantum theory of geometry. I: area operators. *Classical and Quantum Gravity*, *14*, A55–A82.

Ashtekar, A., & Lewandowski, J. (1997b). Quantum theory of gravity: II. volume operators. *Advances in Theoretical and Mathematical Physics*, *1*, 388–429.

Ashtekar, A., & Lewandowski, J. (2004). Background independent quantum gravity: A status report. *Classical and Quantum Gravity*, *21*, R53–R152.

Ashtekar, A., Rovelli, C., & Smolin, L. (1992). Weaving a classical geometry with quantum threads. *Physical Review Letters*, *69*, 237–240.

Bain, J. (2013). Emergence in effective field theories. *European Journal for Philosophy of Science*, *3*(3), 257–273.

Butterfield, J. (2011a). Emergence, reduction and supervenience: A varied landscape. *Foundations of Physics*, *41*(6), 920–959.

Butterfield, J. (2011b). Less is different: Emergence and reduction reconciled. *Foundations of Physics*, *41*, 1065–1135.

Butterfield, J., & Isham, C. (1999). On the emergence of time in quantum gravity. In J. Butterfield (Ed.), *The arguments of time* (pp. 116–168). Oxford: Oxford University Press.

Butterfield, J., & Isham, C. (2001). Spacetime and the philosophical challenge of quantum gravity. In C. Callender & N. Huggett (Eds.), *Physics meets philosophy at the Planck scale* (pp. 33–89). Cambridge: Cambridge University Press.

Carlip, S. (2001). Quantum gravity: A progress report. *Reports on Progress in Physics*, *64*(8), 885.

Crowther, K., & Rickles, D. (2014). Introduction: Principles of quantum gravity. *Studies in History and Philosophy of Modern Physics*, *46*, 135–141.

Dittrich, B., & Thiemann, T. (2009). Are the spectra of geometrical operators in loop quantum gravity really discrete? *Journal of Mathematical Physics*, *50*, 0125033.

Fairbairn, W., & Rovelli, C. (2004). Separable Hilbert space in loop quantum gravity. *Journal of Mathematical Physics*, *45*, 2802–2814.

Frittelli, S., Lehner, L., & Rovelli, C. (1996). The complete spectrum of the area from recoupling theory in loop quantum gravity. *Classical and Quantum Gravity*, *13*, 2921–2932.

Huggett, N., & Wüthrich, C. (2013). Emergent spacetime and empirical (in)coherence. *Studies in History and Philosophy of Modern Physics*, *44*(3), 276–285.

Jacobson, T., & Smolin, L. (1988). Nonperturbative quantum geometries. *Nuclear Physics B*, *299*, 295–345.

Landsman, N. P. (2006). Between classical and quantum. In J. Butterfield & J. Earman (Eds.), *Handbook of the philosophy of science. Philosophy of physics* (Vol. 2, pp. 417–553). Amsterdam: Elsevier B.V. arXiv:quant-ph/0506082v2

Loll, R. (1995). The volume operator in discretized quantum gravity. *Physical Review Letters*, *75*, 3048–3051.

Loll, R. (1996). Spectrum of the volume operator in quantum gravity. *Nuclear Physics B*, *460*(1), 143–154.

Nicolai, H., & Peeters, K. (2007). Loop and spin foam quantum gravity: A brief guide for beginners. In I.-O. Stamatescu & E. Seiler (Eds.), *Approaches to fundamental physics* (Vol. 721, pp. 151–184). Lecture notes in physics. Berlin: Springer. arXiv:hep-th/0601129.

Oriti, D. (2014). Disappearance and emergence of space and time in quantum gravity. *Studies in History and Philosophy of Modern Physics*, *46*, 186–199.

Penrose, R. (1971). Angular momentum: an approach to combinatorial space-time. In T. Bastin (Ed.) *Quantum Theory and Beyond* (pp. 151–180). Cambridge: Cambridge University Press.

Rickles, D. (2005). A new spin on the hole argument. *Studies in History and Philosophy of Modern Physics*, *36*, 415–434.

Rovelli, C. (2003). Loop quantum gravity. *Physics World*, *16*, 37–41.

Rovelli, C. (2004). *Quantum gravity*. Cambridge: Cambridge University Press.

Rovelli, C. (2007). Comment on "Are the spectra of geometrical operators in loop quantum gravity really discrete?" by B. Dittrich & T. Thiemann.

Rovelli, C. (2008). Loop quantum gravity. *Living Reviews in Relativity*, *11*(5). http://www.livingreviews.org/lrr-2008-5

Rovelli, C. (2011). A new look at loop quantum gravity. *Classical and Quantum Gravity*, *28*(11), 114005.

Rovelli, C., & Smolin, L. (1988). Knot theory and quantum gravity. *Physical Review Letters*, *61*, 1155–1158.

Rovelli, C., & Smolin, L. (1990). Loop space representation of quantum general relativity. *Nuclear Physics B*, *331*, 80–152.

Rovelli, C., & Smolin, L. (1995a). Discreteness of area and volume in quantum gravity. *Nuclear Physics B*, *442*, 593–622. Erratum: *Nuclear Physics B*, *456*, 753.

Rovelli, C., & Smolin, L. (1995b). Spin networks and quantum gravity. *Physical Review D*, *53*, 5743–5759.

Thiemann, T. (2001). Introduction to modern canonical quantum general relativity. arXiv:gr-qc/0110034v1

Thiemann, T. (2007). *Modern canonical quantum general relativity*. Cambridge: Cambridge University Press.

Wilson, K. G. (1974). Confinement of quarks. *Physical Review D*, *10*(8), 2445–2459.

Wüthrich, C. (Forthcoming). Raiders of the lost spacetime. In D. Lehmkuhl (Ed.) *Towards a theory of spacetime theories*. Birkhäuser. arXiv:1405.5552