Energy-momentum diffusion from spacetime discreteness

Lydia Philpott[*] and Fay Dowker

Blackett Laboratory, Imperial College, London SW7 2AZ, UK

Rafael D. Sorkin

Perimeter Institute for Theoretical Physics, Waterloo ON, Canada

Abstract

We study potentially observable consequences of spatio-temporal discreteness for the motion of massive and massless particles. First we describe some simple instrinsic models for the motion of a massive point-particle in a fixed causal set background. At large scales, the microscopic swerves induced by the underlying atomicity manifest themselves as a Lorentz invariant diffusion in energy-momentum governed by a single phenomenological parameter, and we derive in full the corresponding diffusion equation. Inspired by the simplicity of the result, we then derive the most general Lorentz invariant diffusion equation for a massless particle, which turns out to contain two phenomenological parameters describing respectvely diffusion and drift in the particle's energy. The particles do not leave the light cone however: their worldlines continue to be null geodesics. Finally, we deduce bounds on the drift and diffusion constants for photons from the blackbody nature of the spectrum of the cosmic microwave background radiation.

[*]Electronic address: `l.philpott06@imperial.ac.uk`

I. INTRODUCTION

The search for a theory of quantum gravity is not, as yet, motivated by experimental results. We currently have no unambiguously relevant quantum gravitational phenomena to guide us in developing candidate theories, though it has long been suggested that a nonzero value of the cosmological constant of order 10^{-120} could have a quantum gravitational origin [1–4]. Outside of cosmology, black hole thermodynamics is often mentioned as one example of a realm where concepts of general relativity and quantum mechanics must both come into play - but experimental black hole physics is out of our reach for now and even if analogue models of black holes in condensed matter systems could be tested this would only probe the semiclassical regime and not full quantum gravity. The existing approaches to quantum gravity have therefore been developed with the hope that the confrontation with experiment can be postponed. At the present time, however, the growing number of different approaches means that the importance of testing ideas against observation, if at all possible, is greater than ever.

Experimental verification of quantitative and unexpected predictions is of the utmost importance in the development of a successful new theory. An example pertinent to the current paper is if we were to find observational evidence that spacetime is fundamentally discrete, then that would have a major impact on the direction of quantum gravity research. What form might such evidence take; what could be the Brownian motion of our age? To answer that question requires the development of phenomenology that draws on essential aspects of a discrete theory of quantum gravity which turns out to be achievable in the causal set approach.

Causal set theory is a discrete, Lorentz invariant approach to quantum gravity [5–7]. For reviews and further references see, for example, [8–10]. It is a work in progress: a quantum causal set dynamics still eludes us. Without a quantum dynamics it seems at first sight premature to develop causal set phenomenology but the kinematics of causal set theory is so concrete that we are able to make some progress in this direction.

A causal set is a locally finite partial order and is the kinematical basis for the theory. One could state the central hypothesis as that spacetime *is* a causal set, or, if one wanted to hedge one's bets whilst the foundations of quantum theory are laid, that causal sets are the histories in a sum-over-histories quantum theory of spacetime.

In detail, a causal set is a set C endowed with a binary relation $\prec$ satisfying:

1. transitivity: if $x \prec y$ and $y \prec z$ then $x \prec z$, $\forall x, y, z \in C$;
2. reflexivity: $x \prec x$, $\forall x \in C$;
3. acyclicity: if $x \prec y$ and $y \prec x$ then $x = y$, $\forall x, y \in C$;
4. local finiteness: $\forall x, z \in C$ the set $\{y \mid x \prec y \prec z\}$ of elements is finite.

Our observed continuum Lorentzian manifold, it is assumed, arises as an approximation to an underlying causal set. The partial order gives rise to the causal ordering of events in the approximating continuum spacetime, and the number of elements comprising a spacetime region gives the volume of that region in fundamental units which we take to be of order the Planck volume. The above rules of correspondence give an essentially unique way to associate a class of causal sets to a given continuum spacetime via a process known as 'sprinkling' defined as follows. Given a Lorentzian manifold, (M, g), points are selected from M randomly via a Poisson process in which the probability measure is equal to the spacetime volume measure in some fundamental units. The selected points are the elements of a causal set once they have been endowed with the partial order induced by the spacetime causal order. The number of points chosen from any region of the manifold will be approximately equal to the volume of the region (in fundamental units) up to Poisson fluctuations. For more details on sprinklings see the reviews mentioned above, for a proof of the Lorentz invariance of the process see [11]. We thus have a straightforward way to construct a causal set that could be the discrete underpinning of a particular continuum spacetime and, consequently, a starting point to develop the phenomenology of discrete spacetime.

An obvious place to look for consequences of causal set theory is in the behaviour of particles. If the underlying spacetime is a discrete structure rather than a continuous manifold, free particles might no longer be able to follow precise timelike geodesics. Intuitively, the underlying discreteness could cause the particles to 'swerve' and indeed a model of particle behaviour illustrating this was proposed in [12]. There, a classical particle is modelled in the simplest possible way, as a point with no internal structure. The causal set, C, considered is a sprinkling into Minkowski spacetime and a particle trajectory consists of a *chain* of elements where a chain is a totally ordered subset of C. The trajectory is constructed iteratively, where the trajectory's past determines its future, but only a certain proper time

τ_f (the 'forgetting time') into the past is relevant. If the particle has reached an element e_n with four-momentum p_n, the next element e_{n+1} is chosen such that

- e_{n+1} is in the causal future of e_n and within a proper time τ_f,
- the momentum change $|p_{n+1} - p_n|$ is minimised.

Here the momentum p_{n+1} is defined to be proportional to the vector between e_n and e_{n+1}. Heuristically, the trajectory tries to stay as straight as possible at each step. In this simple model the discreteness of a causal set results in random fluctuations in the momentum of a particle.

One could object that this model is not intrinsic to the causal set as it makes use of information in the continuum manifold to define the momentum change. However, similar models can be defined with no reference to the continuum. Two such models are proposed in Section II. One of our main claims is that, whatever the microscopic model of particle motion, if it is Lorentz invariant and gives rise to small, random fluctuations in the momentum of the particle then it can be approximated by a continuum description as a diffusion in momentum space. In Section III we support this claim by giving the derivation of the diffusion equation for massive particles introduced in [12]. We also derive the particle diffusion equation in a more useful cosmic time form and without the original assumption of spatial homogeneity.

In Section IV we explore the case of massless particles on a causal set and obtain diffusion equations for the momentum of massless particles in the continuum approximation. Bounds are placed on the constants in the massless particle diffusion equation in Section V by considering the effect of momentum diffusion on the spectrum of the Cosmic Microwave Background. We will use units in which $c = h = G = 1$ — which we will refer to as "Planck units". Fundamental units are related to Planck units by a, yet to be determined factor of order 1. Boltzmann's constant is also set to one, $k_B = 1$.

II. INTRINSIC MODELS FOR MASSIVE PARTICLES

As mentioned above, the original microscopic model in [12] depended on information from the continuum Minkowski spacetime whereas a better model ought to be intrinsic to the causal set itself and rely only on the order relation. Two slightly different intrinsic models will be described in this section, to give an idea of the wealth of possibilities available.

We first recall some causal set definitions. Let C be a causal set.

- A *link* is an irreducible relation i.e. a pair of distinct elements a, b such that $a \prec b$ and there exists no distinct c such that $a \prec c \prec b$.

- A *chain* is a totally ordered subset of C. An *n-chain* is a chain with n elements and its *length* is $n - 1$, the number of links.

- A *longest chain* between two elements is a chain whose length is maximal amongst chains between those endpoints. There may be more than one longest chain between two elements.

- On a causal set the closest approximation we have to a timelike geodesic between two elements is a longest chain. For two causal set elements a and b the length of a longest chain between a and b will be denoted $d(a, b)$. For sprinklings into Minkowski spacetime, in the asymptotic limit of large distances, $d(a, b) \sim \alpha T$ where T is the proper time between a and b and α is a (dimension dependent) constant [13].

- There is a link between elements a and b iff $d(a, b) = 1$.

- A *path* is a chain consisting entirely of links i.e. a set of elements $a \prec b \prec c \prec d \prec \dots$ such that $d(a, b) = 1$, $d(b, c) = 1$, $d(c, d) = 1 \dots$.

A. Model 1

If a dynamical rule for particle motion is to be intrinsic to the causal set background it can no longer refer to a forgetting time τ_f. This instead becomes a 'forgetting number', an integer $n_f >> 1$. In Intrinsic Model 1 a particle trajectory is a chain, $\dots e_{n-2} \prec e_{n-1} \prec e_n \dots$ which is determined by the following (Markov of order 2) process:

Given a partial particle trajectory $\dots e_{n-1}, e_n$ the next element e_{n+1} is chosen such that

- $d(e_n, e_{n+1}) = n_f$,

- $d(e_{n-1}, e_{n+1}) = 2n_f$,

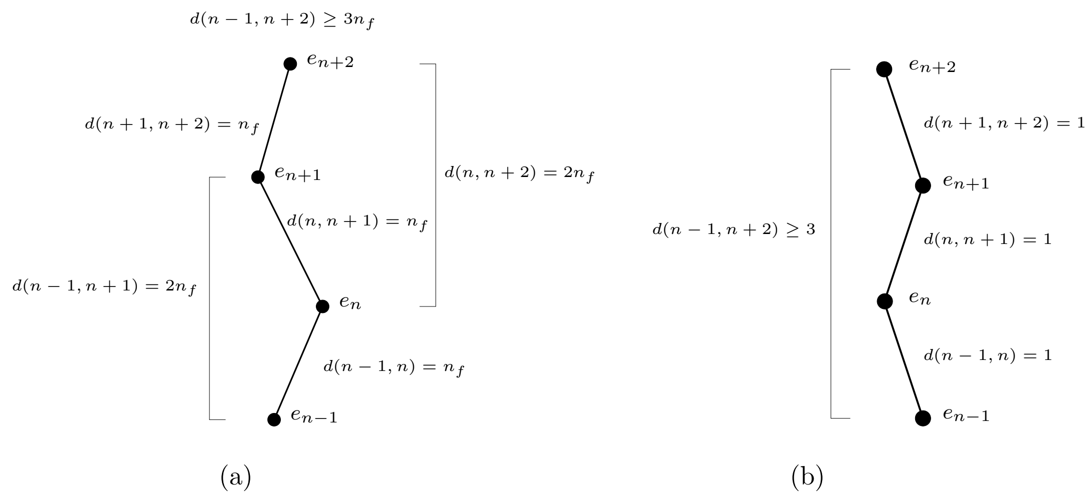

FIG. 1: A trajectory constructed using (a) model 1 and (b) model 2.

(see Figure 1(a)). These requirements do not guarantee the existence of a unique such e_{n+1}. However there will almost surely be finitely many eligible elements and we therefore construct the trajectory by choosing an element uniformly at random from these.

The particle trajectory should swerve a little, but remain approximately straight so long as n_f is large, since, in that case, the results of Brightwell and Gregory [13] show that the expected position of e_{n+1} is close to the hyperboloid of points proper distance n_f/α from e_n and to the hyperboloid of points proper distance $2n_f/\alpha$ from e_{n-1}.

In this model we can consider the trajectory as consisting of just the elements $\ldots e_{n-1}, e_n, e_{n+1} \ldots$ or of the "filled in chain" consisting of a (randomly chosen) longest chain (of length n_f) between e_{n-1} and e_n, another longest chain between e_n and e_{n+1} (also length n_f) and so on. By imposing $d(e_{n-1}, e_{n+1}) = 2n_f$ we have forced the chain of length $2n_f$ that we have between e_{n-1} and e_{n+1} also to be a longest chain. The trajectory is thus approximately geodesic over all $\{e_{n-1} : e_{n+1}\}$ segments. The trajectory consisting of longest chains between e_{n-1} and e_n, e_n and e_{n+1}, and e_{n+1} and e_{n+2} is not, however, necessarily a longest chain between e_{n-1} and e_{n+2}.

Possible variations on this model include choosing, at random, the forgetting number at each step so that the mean is n_f with some fixed variance.

B. Model 2

The trajectory is explicitly constructed as a path in this model i.e. $d(e_n, e_{n+1}) = 1$ for all n. Given a partial particle trajectory $\ldots e_{n-n_f}, \ldots, e_{n-1}$ the next element e_n is chosen such that

- $d(e_{n-1}, e_n) = 1$,
- $d(e_{n-n_f}, e_n) + \ldots + d(e_{n-2}, e_n) + d(e_{n-1}, e_n)$ is minimised,

(see Figure 1(b)). Note that this minimisation does not necessarily yield a unique e_n, in which case we construct the trajectory by choosing an element uniformly at random from those eligible. Also, if the past trajectory has length less than n_f the minimisation is done over all elements available.

Each element is linked to the previous i.e. $d(e_{n-1}, e_n) = 1$ so we know there exists a chain (our trajectory) of length n_f between e_{n-n_f} and e_n. The maximal chain length, $d(e_{n-n_f}, e_n)$, must therefore be greater than or equal to n_f. If we choose e_n to minimise $d(e_{n-n_f}, e_n)$ we ask that the trajectory be as close as possible to geodesic between e_{n-n_f} and e_n while fulfilling $d(e_{n-1}, e_n) = 1$. Minimising the sum of the partial lengths distributes the geodesic property along the path.

III. THE CONTINUUM APPROXIMATION FOR MASSIVE PARTICLES

In simple models of swerves, such as those described above, if there is an appropriate separation of scales so that $1 << n_f << n_{\mathrm{macro}}$, then the change in position and momentum will likely be small at each step, and therefore at the macroscopic scale of many (n_{macro}/n_f) steps the process can be described approximately by a diffusion equation. Although none of the models described above can be considered completely realistic (for example the particles are classical and zero-size) we claim that provided the underlying process for particle propagation is Lorentz- and translation-invariant, it will always give rise to the same diffusion equation, namely the equation written down in [12]. As promised there, we present below the full derivation of this equation, supporting our claim that the continuum model is universal and independent of the discrete microscopic details. We derive the equation initially with the particle's proper time playing the role of independent variable; we then obtain the equivalent equation referred to cosmic time, by expressing both in terms of a conserved current in a certain space of 8 dimensions. This Lorentz invariant process was first considered by Dudley [15], though one of his diffusion equations conflicts with ours. Without any imposition of Lorentz invariance, a general formalism for describing diffusion in Minkowski space was set up by Schay [14].

A. The diffusion equation for a massive particle

We use the general formalism of [16], which deals with stochastic evolution on a manifold of states. The state space, $\mathcal{M}$, of the swerving particle of mass m is $\mathcal{M} = \mathbb{M}^4 \times \mathbb{H}^3$, where $\mathbb{H}^3$ is the mass shell. The coordinates on $\mathbb{M}^4$ are the usual Cartesians $\{x^\mu\}$, $\mu = 0, 1, 2, 3$ and indices are raised and lowered with $\eta_{\mu\nu}$, the Minkowski metric. The spatial coordinates on $\mathbb{M}^4$ will be written as $\{x^i\}$. Cartesian coordinates in momentum space are p_μ and whenever they are used it will be understood that p_μ lies on the mass shell which is the hyperboloid in momentum space defined by $p_\mu p^\mu + m^2 = 0$. $p^0 = E$ is the energy and $p = \sqrt{p_1^2 + p_2^2 + p_3^2}$ is the norm of the three momentum. The 3 coordinates on $\mathbb{H}^3$ will be written abstractly as p^a. We denote the coordinates on $\mathcal{M}$ collectively as $X^A = \{x^\mu, p^a\}$ and in what follows capital letters A, B will be used to indicate general indices on $\mathcal{M}$; μ, ν are indices on $\mathbb{M}^4$; i, j are spatial indices on $\mathbb{M}^4$; a, b are indices on $\mathbb{H}^3$.

The metric on $\mathcal{M}$ is the product of the Minkowski metric $\eta_{\mu\nu}$ on $\mathbb{M}^4$ and the Lobachevski metric g_{ab} on $\mathbb{H}^3$. This is the unique Poincaré invariant metric (up to an overall constant). The "density of states", n, plays a role in the formalism of [16], and by symmetry, it must be proportional to the volume measure on $\mathcal{M}$, so $n \propto \sqrt{g}$ where $g = \det(g_{ab})$. The "entropy scalar", s, is given by $s = \ln(n)$ (Boltzmann's constant has been set to 1).[21]

A process that undergoes stochastic evolution on a manifold of states, $\mathcal{M}$, in time parameter T, can be described by a current, J^A and a continuity equation [16]:

$$J^A = -\partial_B \left(K^{AB}\rho\right) + v^A \rho \tag{1}$$

$$\frac{\partial \rho}{\partial T} = -\partial_A J^A. \tag{2}$$

Here the probability density for the system is given by $\rho = \rho\left(X^A, T\right)$, a scalar density on $\mathcal{M}$. The coefficients K^{AB} are given by

$$K^{AB} = \lim_{\Delta T \to 0+} \left\langle \frac{\Delta X^A \Delta X^B}{2\Delta T} \right\rangle , \tag{3}$$

where $< \cdot >$ denotes expectation value in the process in which the particle starts at a definite point of $\mathcal{M}$ (page 146 of [16]). K^{AB} is a symmetric, positive semi-definite matrix which transforms as the components of a tensor on $\mathcal{M}$. The coefficients v^A are

$$v^A = \lim_{\Delta T \to 0+} \left\langle \frac{\Delta X^A}{\Delta T} \right\rangle , \tag{4}$$

and do not transform as a vector on $\mathcal{M}$, but can be combined with K and the entropy scalar s to form a true vector u^A,

$$u^A = v^A - \partial_B K^{AB} - K^{AB} \partial_B s \,. \tag{5}$$

The current and continuity equations can be reexpressed in terms of the true vector u^A:

$$\frac{\partial \rho}{\partial T} = \partial_A \left(K^{AB} n \partial_B \left(\frac{\rho}{n} \right) - u^A \rho \right) . \tag{6}$$

To find the diffusion equation for our particle process, therefore, we need to determine K^{AB} and u^A.

Requiring the equation to be Poincaré invariant is a very stringent condition and proves to be sufficient for us to determine K^{AB} and u^A, up to the choice of one constant parameter. This means that the resulting equation is very robust and independent of the details of the underlying particle model so long as it is Poincaré invariant.

Consider the process referred to τ, proper time along the worldline of the particle. Then

$$K^{\mu\nu} = \lim_{\Delta \tau \to 0+} \left\langle \frac{\Delta x^\mu \Delta x^\nu}{2 \Delta \tau} \right\rangle . \tag{7}$$

$\Delta x^\mu = \frac{1}{m} p^\mu \Delta \tau$ at every step of the process and so $K^{\mu\nu} = \frac{1}{2} \lim_{\Delta \tau \to 0} p^\mu p^\nu \Delta \tau = 0$. Since K^{AB} is positive definite, this implies that $K^{\mu A} = 0$ and the only nonzero components are K^{ab}. The only Lorentz invariant tensor on $\mathbb{H}^3$ is proportional to the metric, g^{ab} and the coefficient is independent of x^μ by translation invariance. So we have

$$K^{AB} = \begin{pmatrix} 0 & 0 \\ 0 & k g^{ab} \end{pmatrix} \tag{8}$$

where $k > 0$ is a constant.

Now consider

$$v^\mu = \lim_{\Delta \tau \to 0+} \left\langle \frac{\Delta x^\mu}{\Delta \tau} \right\rangle , \tag{9}$$

which, by the above is $v^\mu = p^\mu/m$. The components of the true vector u^μ are equal to v^μ because $K^{\mu A} = 0$. There is no Lorentz invariant vector on $\mathbb{H}^3$ and so $u^a = 0$:

$$u^A = (p^\mu/m, 0) \ . \tag{10}$$

We can now write down the proper time diffusion equation from (1) and (2):

$$\frac{\partial \rho_\tau}{\partial \tau} = k\ \partial_a \left(g^{ab}\sqrt{g}\partial_b \left(\frac{\rho_\tau}{\sqrt{g}} \right) \right) - \frac{1}{m} p^\mu \partial_\mu \rho_\tau \ . \tag{11}$$

If we define a scalar $\overline{\rho} = \rho_\tau/\sqrt{g}$ we obtain the equation in reference [12]:

$$\frac{\partial \overline{\rho}}{\partial \tau} = k\ \nabla^2_H \overline{\rho} - \frac{1}{m} p^\mu \partial_\mu \overline{\rho} \tag{12}$$

where ∇^2_H is the Laplacian on $\mathbb{H}^3$.

B. Diffusion in cosmic time for massive particles

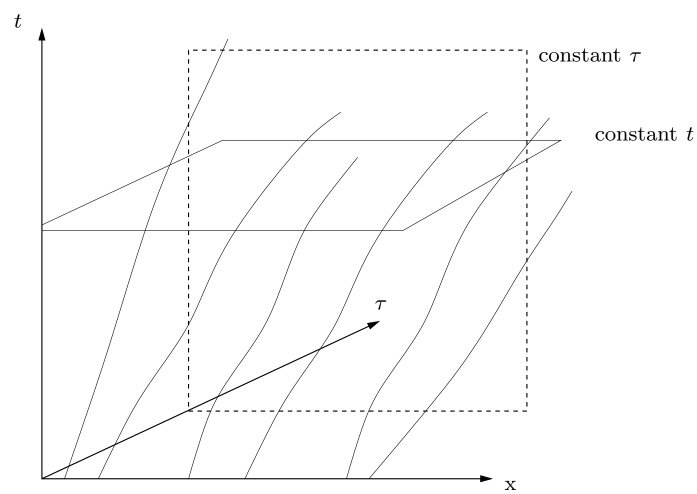

FIG. 2: Particle trajectories as flowlines in $\mathcal{M}' = \mathbb{M}^4 \times \mathbb{H}^3 \times \mathbb{R}$ (where we have suppressed two spatial dimensions).

Given an initial distribution of particles, for instance from an astronomical source, the above equation is not very useful for predicting the results of observations. Even if particles all leave the source at the same time with same momentum, the momentum variation induced by the swerves will result in particles arriving after different proper times and at different observatory times. The proper time that elapses along the particles' worldlines from source to detector is not observable. To compare the swerves model with experiment and observation it is necessary to describe the evolution of the distribution in time in the rest frame of our detector, which time we refer to as *cosmic time*.

A first step in this direction was to look at the nonrelativistic limit of the proper time diffusion equation, when proper time and cosmic time are comparable. The nonrelativistic limit in fact proves sufficient to place very strong bounds on the value of the diffusion constant and severely limit any observable effects (see [12] and [17]).

In the fully relativistic case, Dowker et al. wrote down the diffusion equation in terms of cosmic time for the special case of an initially spatially homogeneous distribution [12]. We will now give the derivation of the cosmic time evolution equation for the general case of spatially inhomogeneous distributions.

The conversion between proper time and cosmic time is possible because both are good time parameters along all possible particle worldlines, which are causal. If we visualise our diffusion process as a collection of such worldlines through spacetime and momentum space, both cosmic time, $t = x^0$, in our chosen frame and proper time τ increase monotonically along each trajectory. Adding proper time to our state space by assuming that the particle starts at parameter $\tau = 0$ and cosmic time $t = 0$, the process is represented by flowlines in $\mathcal{M}' = \mathbb{M}^4 \times \mathbb{H}^3 \times \mathbb{R}$ (see Figure 2) and along each flowline, both τ and t are good time parameters. The proper time diffusion equation we have found describes the evolution of the distribution on constant τ hypersurfaces in $\mathcal{M}'$. What we want is to obtain the diffusion equation for evolution of the distribution on constant t hypersurfaces integrated over all proper times.

First we put t and τ on an equal footing by considering the larger space $\mathcal{M}'$ and defining a new current component

$$J^\tau(t, x^i, p^a, \tau) = \rho_\tau \ . \tag{13}$$

If we denote coordinates on this extended space, $\mathcal{M}' = \mathcal{M} \times \mathbb{R}$, by $X^\alpha = \{X^A, \tau\}$ then the continuity equation (2) can be written

$$\partial_\alpha J^\alpha = 0 \ . \tag{14}$$

Using equation (1) (and still treating τ as our time-parameter) we can express the t

component of the current in terms of J^τ (a.k.a. ρ_τ).

$$\begin{aligned} J^t(t, x^i, p^a, \tau) &= -\partial_B \left(K^{tB} J^\tau\right) + v^t J^\tau \\ &= v^t J^\tau \\ &= \gamma J^\tau \end{aligned} \tag{15}$$

where $\gamma = \partial t / \partial \tau$ is the usual relativistic gamma factor. The remaining components of the current can now be written in terms of J^t. The spatial components are:

$$\begin{aligned} J^i(t, x^i, p^a, \tau) &= -\partial_B \left(K^{iB} J^\tau\right) + v^i J^\tau \\ &= v^i J^\tau \\ &= \frac{p^i}{m} \frac{J^t}{\gamma}. \end{aligned} \tag{16}$$

In the case of the p components the algebra is simpler if we first note that we can express (1) in the form (cf. (6))

$$J^A = -K^{AB} \left(n\, \partial_B \left(\frac{\rho}{n}\right)\right) + \rho u^A, \tag{17}$$

and so

$$\begin{aligned} J^a(t, x^i, p^a, \tau) &= -k g^{ab}\, n\, \partial_b \left(\frac{J^t}{\gamma\, n}\right) \\ &= -k g^{ab} \sqrt{g} \partial_b \left(\frac{J^t}{\gamma \sqrt{g}}\right) \ . \end{aligned} \tag{18}$$

The metric g^{ab} that appears here is the Lobachevski metric on $\mathbb{H}^3$.

Since τ is unobservable we need to integrate J over τ and we denote the integrated current by $\bar{J}$. Integrating the t component of the current over proper time from zero to infinity, gives us the probability density on a hypersurface of constant t:

$$\begin{aligned} \rho_t &= \bar{J}^t(x^i, p^a, t) \\ &\equiv \int J^t d\tau. \end{aligned} \tag{19}$$

The components of the new current can be written:

$$\begin{aligned}
\bar{J}^i(x^i,p^a,t) &\equiv \int J^i d\tau \\
&= \int \frac{p^i J^t}{m\gamma} d\tau \\
&= \frac{p^i}{m}\frac{\bar{J}^t}{\gamma} = \frac{p^i}{m}\frac{\rho_t}{\gamma} \qquad (20)\\
\bar{J}^a(x^i,p^a,t) &\equiv \int J^a d\tau \\
&= -kg^{ab}n\partial_b\left(\frac{\bar{J}^t}{\gamma n}\right) \\
&= -kg^{ab}n\partial_b\left(\frac{\rho_t}{\gamma n}\right) \ . \qquad (21)
\end{aligned}$$

If we integrate the continuity equation over τ we obtain

$$[J^\tau]_0^\infty + \partial_t \bar{J}^t + \partial_i \bar{J}^i + \partial_a \bar{J}^a = 0 \ . \tag{22}$$

$J^\tau|_{\tau=0}$ is zero for all $t>0$ and J^τ tends to zero as τ goes to infinity for finite t. So for all $t>0$ we have

$$\partial_t \bar{J}^t + \partial_i \bar{J}^i + \partial_a \bar{J}^a = 0 \tag{23}$$

which gives the cosmic-time diffusion equation

$$\frac{\partial\rho_t}{\partial t} = -\frac{p^i}{m\gamma}\partial_i\rho_t + k\ \partial_a\left(g^{ab}\sqrt{g}\partial_b\left(\frac{\rho_t}{\gamma\sqrt{g}}\right)\right) . \tag{24}$$

This is a powerful phenomenological model because it depends on only one parameter, the diffusion constant k. Data can therefore strongly constrain k.

We note that this solves a problem posed by Dudley [15]. We also point out that Dudley's equation for the spatially homogeneous distribution on page 267 of [15] is inconsistent with our equation (24). (Equation (3.60) of [14] also differs from the 1+1 dimensional analog of (24).)

IV. MASSLESS PARTICLES

If an underlying spacetime discreteness results in diffusion in momentum and spacetime for massive particles, it is interesting to consider whether a similar diffusion occurs for massless particles. For massive particles the concrete models for particle dynamics on a causal set described above motivated the derivation of the diffusion equation (11). The case of massless particles on a causal set background is rather different. If we consider a sprinkling into Minkowski spacetime, for any given element, p, there will almost surely be no element sprinkled on the future light cone of p. The analogue of the future light cone of p in a sprinkling of Minkowski spacetime is the set of all elements preceded by and linked to p. The elements are distributed, roughly, near the hyperboloid one Planck unit of proper time to the future of p. Although the whole light cone thus has a good causal set analogue[22], the easiest analogue of a null ray is only a single link, making it hard to see how to construct a discrete Markovian process that would result in a close-to-null trajectory.

Modeling the propagation of massless point particles on a causal set as an approximately local process is therefore problematic. It is hoped that in the future, the study of massless *fields* on a causal set will enable us to model massless particle propagation as wave packets, say. In the meantime, however, lack of knowledge of the exact nature of massless particle propagation on the discrete level, does not mean we cannot derive a diffusion equation to describe the potential effect of discreteness on photons in the continuum approximation. We can arrive at a massless diffusion equation in two ways: using the stochastic evolution on a manifold of states procedure as for the massive particle case, or simply taking a $m \to 0$ limit of the diffusion equation for massive particles. It turns out that the second method gives an incomplete result.

The state space in the massless case differs from the massive case. For massive particles we had a probability distribution on $\mathbb{M}^4 \times \mathbb{H}^3$. For massless particles $\mathbb{H}^3$ becomes the 'light cone' in momentum space defined by $p_\mu p^\mu = 0$. This cone will be denoted $\mathbb{H}^3_0$. If we assume that the photons under consideration are well described in a geometrical optics approximation so they have definite spacetime worldlines and momenta, our state space will be[23] $\mathbb{M}^4 \times \mathbb{H}^3_0$. Since proper time vanishes along a lightlike worldline, it is no longer a suitable time parameter for our diffusion process. We define, instead, an affine time, λ, along any photon worldline by

$$dx^\mu = p^\mu d\lambda \ .$$

Notice that the normalization of this affine parameter is not arbitrary. It is fixed by its relation to the particle's four-momentum, or geometrically, to its de Broglie wavelength. Under the latter interpretation, the affine parameter along a photon worldline γ measures the area swept out in spacetime by a vector connecting γ to a neighboring null geodesic that trails it by one wavelength.

In the massive particle case we equated the density of microstates, n, to the determinant of the metric on our state space: $n \propto \sqrt{g}$. In the massless case, the metric induced on $\mathbb{H}^3_0$ degenerates, but $\mathbb{H}^3_0$ still possesses a Lorentz invariant measure of volume (unique up to a constant factor). The four dimensional volume element d^4p of momentum space, together with the masslessness constraint, $p^\mu p_\mu = 0$, lets us construct on $\mathbb{H}^3_0$ the invariant volume element $d^4p\delta(p^\mu p_\mu) = d^3p/2p^0$, i.e. $n \propto 1/p^0$ in Cartesian coordinates. It will be more useful, however, to work in polar coordinates on $\mathbb{H}^3_0$: $\{p, \theta, \phi\}$ where p is the magnitude of the three momentum and θ and ϕ are the usual polar angles in momentum space. In these coordinates, the density of states is $n \propto p \sin\theta$. There is also a (unique up to a constant factor) invariant vectorfield on $\mathbb{H}^3_0$ which is the momentum itself, p^a i.e. the vector with components $(p, 0, 0)$ in polar coordinates. This is absent in the massive case, where the momentum vector does not lie in the mass shell. Finally, although there is no invariant metric on $\mathbb{H}^3_0$, there is an invariant symmetric 2-tensor, $p^a p^b$ (unique up to a constant factor).

We first consider the process in affine time, λ. As with the massless case, we begin with the current and continuity equations, (1) and (2), and determine K^{AB} and u^A. Using the formulae (3) and (4) with $T = \lambda$ we find

$$K^{\mu\nu} = \lim_{\Delta\lambda \to 0+} p^\mu p^\nu \Delta\lambda = 0. \tag{25}$$

K is positive semidefinite so $K^{\mu a} = 0$, and finally K^{ab} must be Lorentz invariant and translation invariant so

$$K^{AB} = \begin{pmatrix} 0 & 0 \\ 0 & k_1 p^a p^b \end{pmatrix} \tag{26}$$

where $k_1 \geq 0$ is a constant.

To determine u^A we again look individually at the components in spacetime and momentum space. As before, the spacetime component $u^\mu = v^\mu$ by (5), and $v^\mu = p^\mu$ by (4). In contrast to the massive case, there can be nonzero components of u^A in the momentum space directions because the momentum itself is an invariant vector. The momentum direction components are thus given by $u^a = k_2 p^a$, where k_2 is a constant. Working in polar coordinates the 'position' vector p^a on the cone $\mathbb{H}^3_0$ is simply $(p, 0, 0)$ where $p^2 = {p_0}^2$. Thus $u^A = (p^0, p^1, p^2, p^3, k_2 p, 0, 0)$ on $\mathbb{M}^4 \times \mathbb{H}^3_0$.

Substituting the forms for K^{AB} and u^A into (6) we obtain the massless particle affine time equation:

$$\begin{aligned}\frac{\partial \rho_\lambda}{\partial \lambda} &= \partial_A \left(K^{AB} n \partial_B \left(\frac{\rho_\lambda}{n} \right) - u^A \rho_\lambda \right) \\ &= -p^\mu \frac{\partial \rho_\lambda}{\partial x^\mu} + k_1 \frac{\partial}{\partial E} \left(E^3 \frac{\partial}{\partial E} \left(\frac{\rho_\lambda}{E} \right) \right) \\ &\quad -k_2 \frac{\partial}{\partial E} \left(E \rho_\lambda \right) \end{aligned} \tag{27}$$

where we have replaced p by energy $E = p$.

We see that the Lorentz invariance means that any diffusion in photon momentum cannot change the direction of the photon and so it always propagates on the light cone, at the speed of light. However the energy of the photon does undergo a diffusion. Notice also that there are two parameters, making this a less powerful phenomenological model than the massive particle model which has a single parameter. There is not only a diffusion term but an independent drift term, arising from the existence of an invariant vector on $\mathbb{H}^3_0$, and we will see that this leads to the existence of power law equilibrium solutions. Note that taking the $m \to 0$ limit of (11) would have resulted in (27) with $k_2 = 0$ because there is no invariant vector in the massive case. (As is familiar in another context, the case of zero photon mass is thus, here also, a sort of singular limit of the massive case.)

A. Cosmic time process

Again, in order to make contact with observations, we need to obtain the cosmic time diffusion equation for massless particles, for which we use the same argument as in the massive case. First we assume that $\lambda = 0$ at $t = 0$. Let

$$J^\lambda(t, x^i, p^a, \lambda) = \rho_\lambda \ . \tag{28}$$

We can then express the t component of the current J in terms of J^λ, and the remaining components of the current in terms of J^t.

$$J^t(t,x^i,p^a,\lambda) \;=\; pJ^\lambda\;; \tag{29}$$

$$J^i\left(t,x^i,p^a,\lambda\right) \;=\; \frac{p^i}{p}J^t\;; \tag{30}$$

$$\begin{aligned} J^a\left(t,x^i,p^a,\lambda\right) \;&=\; -k_1 p^a p^b p\sin\theta\,\partial_b\left(\frac{J^t}{p^2\sin\theta}\right) \\ &\quad + \frac{J^t}{p}k_2 p^a\;. \end{aligned} \tag{31}$$

J^a is proportional to p^a and in polar coordinates the vector $p^a = (p,0,0)$, so there is only one nonzero component of J^{p^a} in the radial (energy) direction:

$$J^p\left(t,x^i,p^a,\lambda\right) \;=\; -k_1 p\,\frac{\partial J^t}{\partial p} + (2k_1+k_2)\,J^t \tag{32}$$

The affine time of flight is unobservable so we integrate over it. Defining

$$\bar{J}^t(t,x^i,p^a) \;=\; \int_0^\infty J^t(t,x^i,p^a,\lambda)d\lambda, \tag{33}$$

we integrate the other current components over λ to obtain

$$\bar{J}^i(t,x^i,p^a) \;=\; \frac{p^i}{p}\bar{J}^t \tag{34}$$

$$\bar{J}^p(t,x^i,p^a) \;=\; -k_1 p\frac{\partial \bar{J}^t}{\partial p} + (2k_1+k_2)\,\bar{J}^t. \tag{35}$$

Imposing the continuity equation, gives us the massless particle cosmic time diffusion equation in terms of the scalar density $\bar{J}^t$, which we rename ρ_t:

$$\begin{aligned} \frac{\partial \rho_t}{\partial t} \;&=\; -\partial_i J^i - \partial_a J^a \\ &=\; -\frac{p^i}{E}\partial_i\rho_t - (k_1+k_2)\,\frac{\partial\rho_t}{\partial E} + k_1 E\frac{\partial^2\rho_t}{\partial E^2}\;, \end{aligned} \tag{36}$$

where $E = p$ is the energy.

So, for a massless particle in a geometric optics approximation, we expect that an underlying discreteness can induce fluctuations in the energy of the particle, but without affecting the direction of propagation. The diffusion governed by k_1 causes a distribution of energies that is initially sharply peaked to spread over time. The second constant k_2 results in an independent drift of the spectrum to higher or lower energies depending on its sign.

It is interesting that negative values of k_2 allow for power law equilibrium solutions of (36). Set $\partial_\mu \rho_t = 0$. Then the equilibrium distributions satisfy

$$-\left(k_1 + k_2\right) \frac{\partial \rho_t}{\partial E} + k_1 E \frac{\partial^2 \rho_t}{\partial E^2} = 0 \ . \tag{37}$$

This has a power law solution

$$\rho_t \propto E^{\frac{2k_1 + k_2}{k_1}} \ . \tag{38}$$

When the parameters are such that the exponent is less than -2 (and so k_2 must be negative because k_1 is positive) then this solution is normalisable if it is cut off at small energies. We conjecture that if $(2k_1 + k_2)/k_1 < -2$ any normalised distribution will tend at late times to this power law equilibrium solution at large energies. This is interesting because physical processes that result in power law distributions across a wide energy range are few and far between — the Fermi mechanism of statistical acceleration of charged particles by random magnetic fields, proposed as the source of high energy cosmic rays, is the only well-known mechanism.

Placing bounds on the parameters k_1 and k_2 is the next step.

V. BOUNDING THE CONSTANTS k_1 AND k_2

In developing a phenomenological model, one aims to provide a model for currently unexplained observations or suggest new observations that might be made to test a theory. But before proposing new observations, one should of course constrain one's model as tightly as possible, based on what is already known. Our model has two parameters: a positive diffusion constant k_1 and a 'drift' constant k_2, which may be either positive or negative. To place the strongest bounds on the values of these parameters, it seems sensible to look at photons that have been travelling for a very long time and thus have had the maximum possibility to experience any underlying discreteness. The cosmic microwave background (CMB) seems an ideal testing ground in this sense. Not only are its photons the "oldest" we can observe, but its spectrum has been determined with great precision. Most of the photons in the CMB have been "free streaming" for approximately 13.7 billion years, or on the order of 10^{60} Planck times. When the universe became transparent at recombination, they would have had a blackbody spectrum with a temperature $3000\,^oK$ (see for example [18]). Current observations yield a temperature of $2.728 \pm 0.004\,^oK$ and measure the spectrum to be Planckian (blackbody) over the $2-21cm^{-1}$ frequency range to within a weighted rms deviation of only 50 parts per million (ppm) of the peak brightness [19]. Since our diffusion would have distorted the energy distribution, the fact that the CMB photons have travelled so far but remained so perfectly thermal will allow us to constrain our parameters very tightly.

A. Simulations

Our derivation of the massless cosmic time diffusion equation assumed spacetime to Minkowskian. Therefore we will first consider a simplified model that ignores the expansion of the universe, and consequently assumes that, in the absence of diffusion, the temperature of the CMB would remain constant from the surface of last scattering to today. This will give us an order of magnitude bound on the parameters. In Section VI the cosmic expansion will be incorporated.

The initial Planckian spectrum, expressed as a number density of photons per unit spatial volume per unit energy, is

$$\rho(E, t=0) = 8\pi \frac{E^2}{\exp\left(\frac{E}{T}\right) - 1} \tag{39}$$

with a temperature $T = 2.728^o K$. According to our model, this distribution evolves via the homogeneous massless cosmic time diffusion equation

$$\frac{\partial \rho}{\partial t} = -\left(k_1 + k_2\right) \frac{\partial \rho_t}{\partial E} + k_1 E \frac{\partial^2 \rho_t}{\partial E^2} \ . \tag{40}$$

Using the `MATLAB` numerical pde solver `pdepe`, this equation was integrated over a time interval equal to that since the surface of last scattering.

Although only the $2-21cm^{-1}$ region of the spectrum is needed to compare with the reported rms deviation, these evolutions were run over a larger range of frequencies to capture more of the spectrum and allow the implementation of a boundary condition at $E = 0$. What boundary condition is appropriate? What happens to a photon as its momentum approaches zero? Do photons leak away through the tip of the null cone in momentum space? Physically, the photon concept employed by our model breaks down as the wavelength tends to infinity, because the geometrical optics approximation fails. (Moreover, our affine parameter λ fails to be well defined physically, since it reaches infinity in a finite time if $E \to 0$, and thus cannot remain approximately constant over the photon wave-packet.) This suggests that the so called "absorbing boundary condition", $\rho(E) = 0$, is appropriate at $E = 0$, and this is what was used in all our simulations. In fact, the current is

$$J = (2k_1 + k_2)\rho - k_1 E \partial\rho/\partial E \ , \tag{41}$$

so (as long as $\partial\rho/\partial E$ remains finite) any linear combination of $\rho = 0$ with the "reflecting boundary condition", $J = 0$, is equivalent at $E = 0$.

The evolved spectrum was converted from a number density per unit volume per unit frequency to a spectral radiance — energy per unit area per unit time per unit frequency per steradian — as used in the analysis of the COBE FIRAS data. This allows us to compare the deviation from Planckian with the quoted 50ppm of the peak brightness.

A Planck spectrum was fit to the evolved spectral radiance using the least squares method. By looking for the best fit Planck spectrum rather than comparing with the initial $2.728^o K$ spectrum, we allowed for the possibility that the diffusion changes the temperature of the CMB in a way that may be reconciled with observation. As it happens, we found that the temperature of the best fit Planck spectrum was very close to the initial temperature in cases where the deviation is within the allowed tolerance. For example the choice of parameters $k_1 = 5 \times 10^{-97}$ and $k_2 = 1 \times 10^{-96}$ gives a best fit temperature of $2.7281^o K$, indistinguishable from the current observed temperature of $2.728 \pm 0.004^o K$. Finally the rms deviation between the fitted Planckian spectrum and the evolved spectrum in the $2 - 21 cm^{-1}$ frequency range (energy range $4 \times 10^{-23} - 4 \times 10^{-22} J$) was calculated with all points weighted equally. This result was compared to the allowed tolerance of 50 parts per million of the peak brightness. This process was repeated for a range of values of the parameters k_1 and k_2.

B. Results

We first place bounds on the diffusion and drift constants separately, varying k_1 with $k_2 = 0$ and varying k_2 with $k_1 = 0$.

When $k_1 = 0$ we can solve the equation exactly:

$$\rho(E, t) = \rho_0(E - k_2 t) \tag{42}$$

so the spectrum just translates at a constant speed. For k_2 negative, this is inconsistent with the boundary condition $\rho = 0$ at $E = 0$. However, in this case one can implement an

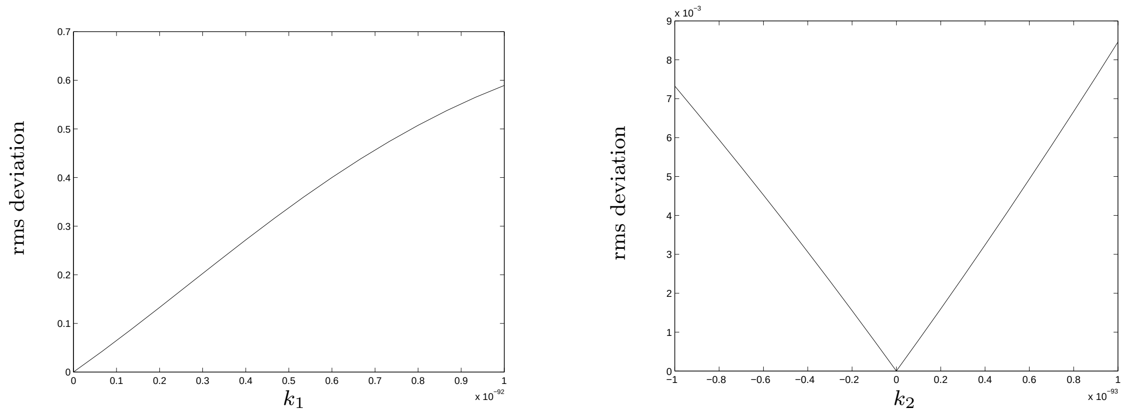

(a)

(b)

FIG. 3: The rms deviation between the evolved spectrum and a bestfit Planck spectrum (as a proportion of the spectrum peak):(a) varying k_1 with $k_2 = 0$, and (b) varying k_2 with $k_1 = 0$.

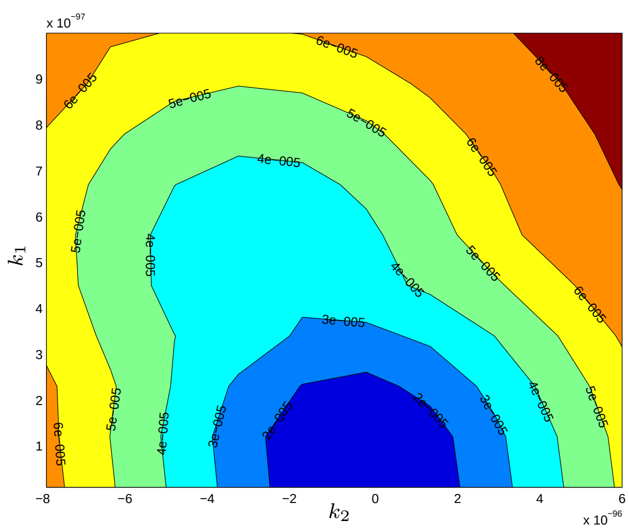

FIG. 4: The rms deviation of the simulated spectrum from Planckian as a proportion of the spectrum peak, varying both k_1 and k_2. Values of k_1 and k_2 within the $5e{-}05$ contour give a spectrum that is Planckian to within 50ppm of the peak.

absorbing boundary condition trivially: simply cut off the translated distribution at $E = 0$. This is what we did to generate the solution plotted in Figure 3(b).

One might be concerned that deviations within the allowed tolerance would be so small as to approach the level of the numerical errors in the simulations. The exact solution for $k_1 = 0$ provides us with a means of demonstrating that this is not the case[24]. When we compare the exact solution with the numerical solution for $k_1 = 0$ the errors introduced by the numerical integration can be seen to be several orders of magnitude smaller than the deviation from Planckian. For example if $k_2 = 4 \times 10^{-96}$ the rms deviation from the best fit Planck spectrum is 5×10^{-101} (5×10^{-5} peak brightness) for both the exact and the numerical solution. The rms deviation *between* the exact and numerical solution is 4×10^{-104}. If $k_2 = -4 \times 10^{-96}$ the rms deviation from the best fit Planck spectrum is also 5×10^{-101} (5×10^{-5} peak brightness) while the deviation between the exact and numerical solutions is again 4×10^{-104}. This also demonstrates that the $\rho = 0$ boundary condition we imposed on the numerical solution, although inconsistent with the exact solution when $k_2 < 0$, does not introduce noticeable errors for the values of k_1 that we are concerned with.

When $k_2 = 0$ with $k_1 > 0$ we can only solve the equation numerically. The results for both cases are displayed in Figures 3(a) and 3(b). We see that the deviation from Planckian increases approximately linearly with increasing magnitude of the parameters. (Notice that figure 3(b) was drawn from the exact solution, the graph taken from the numerical solution is indistinguishable.) The simulations suggest that for the deviation from Planckian of the CMB to be within the allowed 5×10^{-5} of the peak brightness the diffusion constant k_1 must be less than approximately 7×10^{-97} if $k_2 = 0$, and the drift parameter k_2 must fall within the range $-4 \times 10^{-96} < k_2 < 4 \times 10^{-96}$ if $k_1 = 0$. Converting to SI units we have the bounds:

$$k_1 < 3 \times 10^{-44} kgm^2 s^{-3} \tag{43}$$

$$-1 \times 10^{-43} < k_2 < 1 \times 10^{-43} kgm^2 s^{-3}. \tag{44}$$

Similar bounds apply when we let both k_1 and k_2 be nonzero. The general situation is displayed in Figure 4, from which one can read off the values of k_1 and k_2 for which the deviation from blackbody is less than 5×10^{-5} of the peak brightness when we allow both constants to vary.

In the units used here, the bounds on the parameters are very small. However we can get a handle on where these numbers come from by rescaling the energy, setting $E' = sE$ with s chosen so that $sT = 1$ when T is the CMB temperature. This means that $s \sim 10^{32}$ in Planck units. We rescale $\rho' = \rho/s$ so the initial spectrum is:

$$\rho'_0(E') = 8\pi \frac{1}{s} \frac{E'^2}{s^2(e^{E'/T'} - 1)} \tag{45}$$

where $T' = sT = 1$. If $k_2 = 0$, then we can also rescale the time, setting $t' = s\, k_1 t$ to obtain the diffusion equation

$$\frac{\partial \rho'}{\partial t'} = -\frac{\partial \rho'}{\partial E'} + E' \frac{\partial^2 \rho'}{\partial E'^2} . \tag{46}$$

If we now evolve ρ' until it differs from ρ'_0 by 50 ppm and take the value, t'_f of t' when this happens, t'_f must, for consistency with the data, be greater than or equal to $s\,k_1 t$ where t is the age of the universe, and so in Planck units $k_1 \leq 10^{-60} 10^{-32} t'_f$. We see that the order of magnitude bound found above will result if $t'_f \sim 10^{-4}$, which is indeed about the (rescaled) time at which one would have expected the deviation to reach 50 ppm. A similar order of magnitude estimate follows from the geometric interpretation of our affine parameter λ as an area, if one notes that the product of the photon wavelength ($\sim 1cm$) with the Hubble radius is around $10^{32} 10^{60} \sim 10^{92}$ in Planck units.

VI. EXPANDING UNIVERSE

In Section V we ignored the effect of the expansion of the cosmos on the CMB and assumed that it remained at a temperature of $\sim 2.7^o K$ from the surface of last scattering to today. This is of course not the case. At the surface of last scattering the CMB had a temperature of about $3000^o K$. As the universe expanded the individual photons were stretched along with the space, and correspondingly diluted, leaving us with the $2.7^o K$ spectrum observed today. We will now show that the expansion has essentially no effect on our model in the sense that the distribution in the expanding universe can be deduced easily from the nonexpanding one and that the bounds derived from the nonexpanding simulation change only slightly.

The redshifting effect of the expansion (but not the dilution) can be added to the model by adding to v a vector which has a single component in the E direction:

$$\Delta v^E = \frac{dE}{dt} = -E \frac{\dot{a}}{a} , \tag{47}$$

where $a(t)$ is the cosmic scale factor. This changes the continuity equation (2) to

$$\begin{aligned} \frac{\partial \rho_t}{\partial t} &= -\partial_i J^i - \partial_a J^a & (48) \\ &= -\frac{p^i}{E} \partial_i \rho_t - (k_1 + k_2) \frac{\partial \rho_t}{\partial E} \\ &\quad + k_1 E \frac{\partial^2 \rho_t}{\partial E^2} + \frac{\dot{a}}{a} \frac{\partial}{\partial E} (\rho E) . & (49) \end{aligned}$$

A solution of this equation, for $k_1 = k_2 = 0$ is

$$\rho_0\,(E,t) = 8\pi \frac{a^3}{a_0^3} \frac{E^2}{e^{\frac{E}{T_0}\frac{a}{a_0}} - 1}\ , \tag{50}$$

where a_0 is the scale factor at time t_0. If we multiply this distribution by $\frac{a_0^3}{a^3}$, which dilutes the photons according to the expansion, it becomes exactly the Planck distribution for temperature $T = T_0 \frac{a_0}{a}$.

If we define a new variable $\widetilde{E} = \frac{a}{a_0} E$ and a new density function $\widetilde{\rho}(\widetilde{E}) = \frac{a_0}{a}\rho(E)$ (this being just the transformation of a scalar density under a rescaling of coordinates: $\rho\, dE = \widetilde{\rho}\, d\widetilde{E}$) the distribution $\rho_0(E,t)$ (50) becomes

$$\widetilde{\rho}_0\left(\widetilde{E}, t\right) = 8\pi \frac{\widetilde{E}^2}{\exp\left(\frac{\widetilde{E}}{T_0}\right) - 1} \tag{51}$$

which is constant in time.

We now transform our diffusion equation to the rescaled quantities $\widetilde{\rho}$ and $\widetilde{E}$.

Starting with (49), we have:

$$LHS = \left(\frac{\partial}{\partial t} + \frac{\dot{a}\widetilde{E}}{a}\frac{\partial}{\partial \widetilde{E}}\right)(a\widetilde{\rho}) \tag{52}$$

$$= \dot{a}\widetilde{\rho} + a\dot{\widetilde{\rho}} + \dot{a}\widetilde{E}\widetilde{\rho}' \tag{53}$$

$$RHS = -(k_1+k_2)a\,(a\widetilde{\rho})' + k_1 a^2 \widetilde{E}\widetilde{\rho}'' + \dot{a}\left(\widetilde{\rho}\widetilde{E}\right)' \tag{54}$$

$$= -(k_1+k_2)a^2\widetilde{\rho}' + k_1 a^2 \widetilde{E}\widetilde{\rho}'' + \dot{a}\widetilde{\rho} + \dot{a}\widetilde{E}\widetilde{\rho}' \tag{55}$$

where dot denotes time derivative and prime denotes derivative with respect to $\widetilde{E}$. This gives

$$\frac{\partial\widetilde{\rho}}{\partial t} = -(k_1+k_2)a\frac{\partial\widetilde{\rho}}{\partial\widetilde{E}} + k_1 a\widetilde{E}\frac{\partial^2\widetilde{\rho}}{\partial\widetilde{E}^2}. \tag{56}$$

Choosing t' such that $\frac{dt'}{dt} = a$, we obtain

$$\frac{\partial\widetilde{\rho}}{\partial t'} = -(k_1+k_2)\frac{\partial}{\partial\widetilde{E}}\widetilde{\rho} + k_1\widetilde{E}\frac{\partial^2}{\partial\widetilde{E}^2}\widetilde{\rho} \tag{57}$$

which is the same as (40), the nonexpanding diffusion equation.

That we can find expanding solutions from static ones is due to the scale-invariance of the null cone $\mathbb{H}_0^3$: its geometrical structures are invariant under $E \to \widetilde{E} = \text{const} \times E$.

For a matter dominated FRW universe $a \sim t^{2/3}$ i.e. $a(t) = t^{2/3}/t_0^{2/3}$, where t_0 is the current value of t (and the current value of a is 1). We have $\frac{dt'}{dt} = a$ which integrates to

$$t' = \frac{3}{5}\frac{t^{\frac{5}{3}}}{t_0^{\frac{2}{3}}} + \text{const} \; . \tag{58}$$

If the range for t is 10^{60} then the range for t' is 3/5 of this. So the simulations we would need to do for the expanding case are the same as for the nonexpanding case but for only 3/5 of the time. This doesn't affect the order of magnitude of the bounds.

VII. DISCUSSION

The work presented here illustrates the familiar fact that considerations of symmetry can bring forth a fairly unique phenomenological model, even when relatively little is known about the deeper reality the model is meant to represent. Starting from the assumption of an underlying spatiotemporal discreteness that nevertheless respects Lorentz invariance in the continuum approximation, we argued that particle momenta would be subject to stochastic variations, and that if these variations were small, their effects would be describable on large scales as a diffusion in momentum space. The assumption of Lorentz symmetry lends the resulting models their power (by limiting the number of parameters), and it sets them apart from the majority of quantum gravity phenomenological models, which break Lorentz invariance.

For particles without internal degrees of freedom, we have seen that even in the absence of a definite microscopic theory, an effective diffusion model can be derived based on the assumed invariance alone. One can also imagine applying this idea more generally, including for example the polarisation of photons or neutrinos.

In the case of massive particles, if one of the explicit microscopic models is fixed upon, then the diffusion strength, k, will be a function of the forgetting time (number). This forgetting time sets the scale shorter than which the dynamics is nonlocal: at much larger scales the model is effectively local. In more realistic, more quantal models, the diffusion scale might also depend on such dimensionless numbers as the ratio of the mass of the particle to the Planck mass and properties of the particle's wave packet. The same possibilities exist for the massless case. Thus we would expect the diffusion and drift parameters, k_1 and k_2, to depend on some non-locality scale in the underlying physics, and they could also depend on features of the wave packet associated with the photon, for example the ratio of the (peak) wavelength to the length or the packet. In seeking an underlying model of photons, the Lorentz invariant, nonlocal D'Alembertian that has recently been discovered for scalar field propagation on causal set backgrounds [8, 20] could be valuable. Using it to evolve a wave packet of a massless scalar field, one could ask whether the resulting propagation exhibited any momentum diffusion or drift, and if so, what sets the scale of these phenomena.

The parameters of our model are constrained by the blackbody character of the CMB radiation. Since most observational astrophysics and cosmology relies on electromagnetic radiation, there are a host of other observations that could also be brought to bear, given that our model entails a broadening of spectral lines as well as a distance-dependent shift in energy. For example, if the diffusion constant were set to zero, it would be easy to work out how the drift would affect absorption spectra from distant objects. It seems likely however that the bounds set here will be among the most stringent.

Acknowledgments

We thank Joe Henson for invaluable help with the expanding case and Carlo Contaldi for useful discussions. LP acknowledges the support of a TEC doctoral scholarship. FD is supported in part by Marie Curie Research and Training Network "Random Geometry and Random Matrices: From Quantum Gravity to Econophysics" (MRTN-CT-2004-005616) and the Royal Society grant IJP - 2006/R2. Research at Perimeter Institute for Theoretical Physics is supported in part by the Government of Canada through NSERC and by the Province of Ontario through MRI. This research was partly supported by NSF grant PHY-0404646.

[1] R. D. Sorkin, in *Proceedings of the ninth Italian Conference on General Relativity and Gravitational Physics, Capri, Italy, September 1990*, edited by R. Cianci, R. de Ritis, M. Francaviglia, G. Marmo, C. Rubano, and P. Scudellaro (World Scientific, Singapore, 1991), pp. 68–90.

[2] R. D. Sorkin, in *Relativity and Gravitation: Classical and Quantum, Proceedings of the SILARG VII Conference, Cocoyoc, Mexico, December 1990*, edited by J. C. D'Olivo,

E. Nahmad-Achar, M. Rosenbaum, M. P. Ryan, L. F. Urrutia, and F. Zertuche (World Scientific, Singapore, 1991), pp. 150–173.

[3] R. D. Sorkin, Int. J. Theor. Phys. **36**, 2759 (1997), gr-qc/9706002.

[4] M. Ahmed, S. Dodelson, P. B. Greene, and R. Sorkin, Phys. Rev. **D69**, 103523 (2004), astro-ph/0209274.

[5] L. Bombelli, J.-H. Lee, D. Meyer, and R. Sorkin, Phys. Rev. Lett **59**, 521 (1987).

[6] G. 't Hooft, in *Recent Developments in Gravitation (Proceedings of the 1978 Cargese Summer Institute)*, edited by M. Levy and S. Deser (Plenum, 1979).

[7] J. Myrheim, *Statistical geometry* (1978), cERN preprint TH-2538.

[8] J. Henson, in *Approaches to Quantum Gravity: Towards a New Understanding of Space and Time*, edited by D. Oriti (Cambridge University Press, 2006), gr-qc/0601121.

[9] F. Dowker, Contemp. Phys. **47**, 1 (2006).

[10] R. D. Sorkin, in *Lectures on Quantum Gravity, Proceedings of the Valdivia Summer School, Valdivia, Chile, January 2002*, edited by A. Gomberoff and D. Marolf (Plenum, 2005), gr-qc/0309009.

[11] L. Bombelli, J. Henson, and R. D. Sorkin (2006), gr-qc/0605006.

[12] F. Dowker, J. Henson, and R. D. Sorkin, Mod. Phys. Lett. **A19**, 1829 (2004), gr-qc/0311055.

[13] G. Brightwell and R. Gregory, Phys. Rev. Lett **66**, 260 (1991).

[14] G. Schay, Ph.D. thesis, Princeton University (1961).

[15] R. Dudley, Arkiv för Matematik **6**, 241 (1965).

[16] R. Sorkin, Ann. Phys. **168**, 119 (1986).

[17] N. Kaloper and D. Mattingly, Phys. Rev. **D74**, 106001 (2006), astro-ph/0607485.

[18] E. W. Kolb and M. S. Turner, *The Early Universe* (Addison-Wesley, 1990).

[19] D. J. Fixsen et al., Astrophys. J. **473**, 576 (1996), astro-ph/9605054.

[20] R. D. Sorkin (2007), gr-qc/0703099.

[21] For present purposes the identification of n with a density of microscopic states is unnecessary. What is relevant is that a probability-density proportional to n be in equilibrium (i.e. time-independent).

[22] Strictly speaking, it would be more correct to identify the "light cone" with a "virtual boundary" separating the future of p from the set of elements spacelike to it.

[23] More correctly this is the state-space for a massless particle of spin zero. For a true photon,

the state-space would be enlarged so as to describe also the polarization.

[24] We can also note that solving our diffusion equation in `Mathematica` using `NDSolve` yields the same results as discussed here, suggesting the bounds we obtain are robust, do not depend on the particular method of solving the equation, and are not a consequence of numerical error or a particular choice of integration step size.