Light on a Ticking Lattice: Local Discrete-Time Maxwell Evolution and Its Unavoidable Corrections

Jonathan Washburn
Recognition Physics Institute

12 October 2026

Abstract

Space is a cubic grid and time a sequence of ticks, each updating every cell from nearby cells. Can such a world carry light at one cell per tick, obey Maxwell's equations and have the right quantum fluctuations? It can, up to lattice corrections, if a tick reaches all 26 neighbours. A tick reaching only the six face neighbours cannot, whatever its weights: along a body diagonal its signals travel at most $1/\sqrt{3}$ cell per tick. We give a rule that reaches all 26. Light moves exactly one cell per tick along the axes, and no wave packet is faster. Three of Maxwell's equations hold exactly, the fourth up to a lattice term. Unavoidable corrections remain: a frequency ceiling, direction-dependent speeds, a weak signal ahead of the light front, and no vacuum state for the modes at zero frequency and at the ceiling. Every other mode has a closed-form vacuum.

1 The question

Continuum physics describes light with fields that vary smoothly in space and time. Models that build space from discrete cells and time from discrete steps have to say what light does instead. Lattice field theories, cellular automata and quantum walks are examples. The question has a precise form once the model is fixed, and this paper answers it for the simplest geometry.

Space is a cubic grid of cells, called *voxels*, each of side ℓ. Time advances in *ticks* of length τ. A *field* gives a number to every voxel at every tick, or three numbers for a vector field such as the electric field. One tick replaces the field by a weighted sum of its values at nearby voxels, with the same weights at every voxel and every tick. We call such an update a *local rule*, and the set of voxels it draws on, measured from the voxel being updated, its *stencil*. Waves also need the field one tick earlier, and Section 3 shows how such two-step rules fit this description. The rule is linear. That covers a classical field, and it covers a free quantum field too, whose field operators obey the same linear rule. Interacting quantum fields are outside this paper. We measure length in voxels and time in ticks ($\ell = \tau = 1$), so light that crosses one voxel per tick has speed one.

Which voxels can a stencil contain? Besides the voxel itself, Figure 1 names them: the six that share a face with the central voxel, the twelve that share an edge with it, and the eight that touch it only at a corner.

What the model must do. First, long waves of light should cross one voxel per tick whichever way they travel, and no wave packet, a pulse built from waves of nearly equal wave vector (wavelength and direction), should travel faster. Second, the electric and magnetic fields should obey

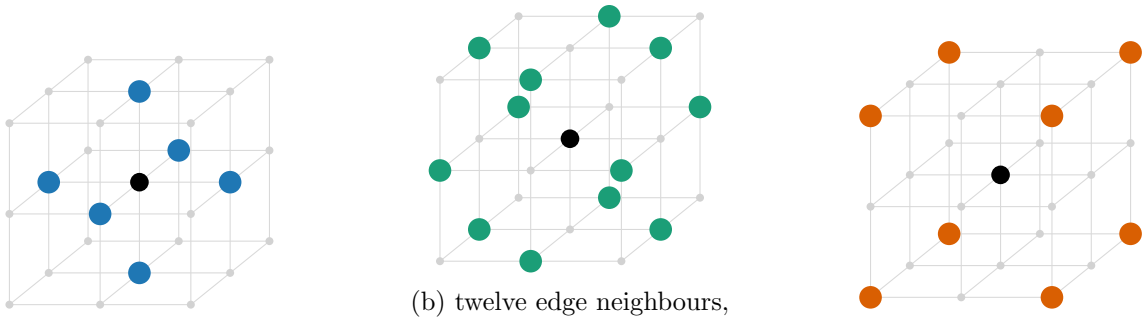

(a) six face neighbours, hops $(\pm1, 0, 0)$ and permutations

(b) twelve edge neighbours, hops $(\pm1, \pm1, 0)$ and permutations

(c) eight corner neighbours, hops $(\pm1, \pm1, \pm1)$

Figure 1: The 26 voxels around a voxel, in a $3 \times 3 \times 3$ block whose central voxel is the black dot. A hop is the displacement from the central voxel, in whole voxels along the three axes. The coloured voxels in each panel are one kind of neighbour; the small grey dots are the rest of the block.

Maxwell's equations. Third, the quantum vacuum should come out right. Quantum theory gives empty space fluctuating fields: each mode of the field (one wave with one polarization) is a quantum oscillator, and even in its lowest-energy state, the *vacuum* of the mode, its field value is not sharp but spreads over a range set by $\hbar$ (Planck's constant divided by 2π) and the mode's frequency. The model should give each mode that spread. Finally, no grid is a continuum, so some departures from continuum physics will remain. We ask which of them cannot be removed, and for each one how widely that holds (Section 8).

The answer in brief. It turns on which neighbours one tick reaches. A tick that reaches only the six face neighbours cannot, whatever weights it uses, move signals along a body diagonal of the grid (the direction $(1, 1, 1)$) faster than $1/\sqrt{3}$ voxel per tick, so it cannot carry light at one voxel per tick in every direction (Section 3). A tick that reaches all 26 neighbours can. The corner rule of Section 4 moves light exactly one voxel per tick along the three axes and, for long waves, in every direction, and no wave packet built from its waves moves faster than one voxel per tick. Among rules that treat the three axes alike, a stable rule needs some weight on the corner neighbours if its shortest wave, the one whose sign flips from voxel to voxel along all three axes, is to have a nonzero frequency (Section 4). Section 6 builds electromagnetism on the corner rule: three of Maxwell's four equations hold exactly at every tick, and the fourth holds up to a lattice term that vanishes along the axes and is of fourth order in the wave number. Section 7 computes in closed form the vacuum of each mode that has one. Section 8 lists the corrections that cannot be removed, each with its scope: a region reached that is never exactly the ball that light fills (a face rule falls short of it, and a rule that reaches beyond the faces sends a weak precursor outside it), a ceiling on frequency, wave speeds that depend on direction, a lattice term in Ampère's law, and no vacuum state for the modes at the two ends of the frequency range, the uniform modes at $\omega = 0$ and the modes at the ceiling $\omega = \pi$.

2 One dimension first

Take a single row of voxels, labelled by the integers x, and the rule

$$u(n+2, x) = u(n+1, x+1) + u(n+1, x-1) - u(n, x). \tag{1}$$

term	meaning
voxel	one cell of the cubic grid, labelled by integers $x = (x_0, x_1, x_2)$; side $\ell = 1$
tick	one time step, $\tau = 1$, labelled by an integer n
hop	a displacement $s \in \mathbb{Z}^3$ from one voxel to another
local rule, stencil	a linear update in which the new value at x is $\sum_s w_s f(x+s)$, a sum over a finite set of hops (the stencil) with the same weights w_s at every voxel and tick
face, edge, corner neighbours	the 6, 12 and 8 voxels at hops $(\pm1, 0, 0)$, $(\pm1, \pm1, 0)$, $(\pm1, \pm1, \pm1)$ and their permutations (Figure 1)
plane wave	a field $e^{i(k\cdot x - \omega n)}$ with wave vector k and frequency ω; $\lvert k\rvert$ is its wave number and its wavelength is $2\pi/\lvert k\rvert$ voxels
dispersion relation	the relation between ω and k that a rule's plane waves obey; in the continuum with light speed one it is $\omega = \lvert k\rvert$
phase speed, group speed	$\omega/\lvert k\rvert$, the speed of the crests, and $\lvert\nabla_k \omega\rvert$, the speed of a wave packet, that is, a pulse built from waves near k
zone	the cube $\lvert k_j\rvert \le \pi$ of wave vectors; on the lattice the waves with k_j and $k_j + 2\pi$ are the same function (Figure 3)
symbol	the factor $\hat{U}(k) = \sum_s w_s e^{ik\cdot s}$ by which a rule multiplies the plane wave $e^{ik\cdot x}$
lag	the relative shortfall of a wave's phase speed below the rule's light speed
stable rule	a rule whose plane waves all have a real frequency ω; in an unstable rule some waves grow exponentially, tick after tick (Section 2). At the two ends of the frequency range, $\omega = 0$ and $\omega = \pi$, the amplitude of a wave of a stable rule grows in proportion to the tick number instead of oscillating (Section 7)
mode	one plane wave with one polarization; the field is a sum of modes, each an independent oscillator
vacuum	for a mode, its lowest-energy quantum state; on the lattice, the pure Gaussian quantum state that one tick leaves unchanged (Section 7)

Table 1: Terms used throughout.

The field at x at tick $n + 2$ is built from the field at the two neighbours of x one tick earlier and the field at x itself two ticks earlier (Figure 2b). Looking back two ticks is natural for waves, whose equation is second order in time: two consecutive ticks of the field make up the state, and each tick moves that pair forward by one step.

Pulses. Let F and G be any two functions on the integers. Then $u(n, x) = F(x - n) + G(x + n)$ satisfies (1). To check, put it into the right side:

$$F(x-n)+G(x+n+2)+F(x-n-2)+G(x+n)-F(x-n)-G(x+n) = F(x-n-2)+G(x+n+2),$$

which is the left side. Every solution has this form. The rule gives each later tick from the first two, so the starting values $u(0, x)$ and $u(1, x)$ fix everything, and F and G can be chosen to match them. The conditions $u(0, x) = F(x) + G(x)$ and $u(1, x) = F(x - 1) + G(x + 1)$ give $F(x + 1) - F(x - 1) = u(0, x + 1) - u(1, x)$. This fixes F at every voxel once its values at two neighbouring voxels are chosen, and then $G = u(0, \cdot) - F$. A lump of field therefore splits into two copies, one moving right and one moving left, each advancing exactly one voxel per tick and keeping its shape (Figure 2a). This is d'Alembert's solution of the wave equation, and on the lattice it is exact.

Waves. A *plane wave* is a field $e^{i(kx - \omega n)}$. It has wavelength $2\pi/\lvert k\rvert$ voxels and oscillates with angular frequency ω per tick; k is its *wave number*. Substituting into (1) and dividing by $e^{i(kx-\omega n)}$

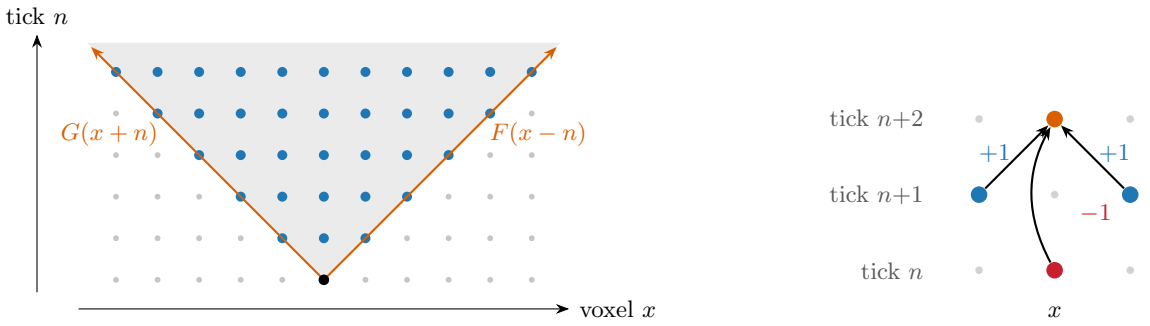

(a) light on a line (b) the rule (1)

Figure 2: The one-dimensional rule. (a) Each dot is a voxel at one tick. A lump started at the origin splits into a right-moving and a left-moving copy that advance along the two orange lines, one voxel per tick. The blue dots, which fill the shaded cone, are the voxels that the lump can have reached by each tick. (b) To get the field at x at tick $n+2$, add the field at the two neighbours of x at tick $n+1$ and subtract the field at x at tick n.

leaves $e^{-2i\omega} = 2\cos k\, e^{-i\omega} - 1$, since the two neighbours contribute $e^{ik} + e^{-ik} = 2\cos k$. Multiplying by $e^{i\omega}$ gives $e^{i\omega} + e^{-i\omega} = 2\cos k$, which is

$$\cos\omega = \cos k. \tag{2}$$

The relation between ω and k is the rule's *dispersion relation*. Two conventions fix what it says. On integer sites e^{ikx} and $e^{i(k+2\pi)x}$ are the same function, so every wave has a wave number in $[-\pi, \pi]$, and the shortest wave alternates in sign from voxel to voxel: wavelength two voxels, $|k| = \pi$. In the same way ω and $\omega + 2\pi$ give the same wave, because n is an integer, and for a real field the wave with (k, ω) is the same wave as the one with $(-k, -\omega)$. So we take $0 \le \omega \le \pi$. Then (2) gives $\omega = |k|$: every wave, whatever its wavelength, moves at one voxel per tick.

So in one dimension the tick that reaches the two neighbours of a voxel carries light perfectly.

Three dimensions. The obvious three-dimensional version adds a second difference along each axis. Write e_j for the unit step along axis j, and

$$\Delta f(x) = \sum_{j=0}^{2} \big[f(x + e_j) - 2f(x) + f(x - e_j) \big] \tag{3}$$

for the sum of the three second differences. The rule

$$u(n+2) = 2u(n+1) - u(n) + r^2\, \Delta u(n+1) \tag{4}$$

says that the second difference in time equals r^2 times the second difference in space. That is a lattice version of the wave equation with wave speed r voxels per tick, and in one dimension $r = 1$ is the rule (1). We call (4) the *face leapfrog*, because its tick reaches the six face neighbours. Every two-step rule below has the form $u(n+2) = 2S\, u(n+1) - u(n)$, where S is a local rule that says how the field one tick ahead feeds the next one; for the face leapfrog $S = 1 + \frac{r^2}{2}\Delta$.

A plane wave in three dimensions is $e^{i(k\cdot x - \omega n)}$ with a wave vector $k = (k_0, k_1, k_2)$. The second difference along axis j multiplies it by $e^{ik_j} - 2 + e^{-ik_j} = -4\sin^2(k_j/2)$, and in the same way the

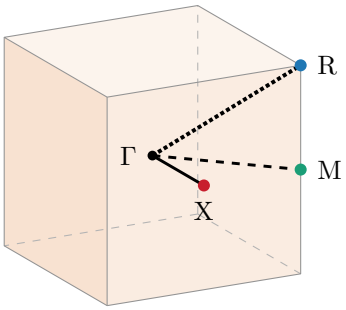

Γ $(0,0,0)$, the longest waves.
X $(\pi,0,0)$, the centre of a face of the cube: wavelength two voxels along one axis.
M $(\pi,\pi,0)$, the centre of an edge: wavelength two voxels along two axes.
R (π,π,π), a corner: wavelength two voxels along all three axes, the shortest wave there is.

Figure 3: The zone of wave vectors, $|k_j| \le \pi$, and four points in it that matter later. The path ΓX (solid) runs along an axis, ΓM (dashed) along a face diagonal and ΓR (dotted) along a body diagonal. Their lengths are π, $\sqrt{2}\,\pi$ and $\sqrt{3}\,\pi$.

second difference in time, $u(n+2) - 2u(n+1) + u(n)$, equals $-4\sin^2(\omega/2)\,u(n+1)$. So (4) requires

$$\sin^2\frac{\omega}{2} = r^2 \sum_{j=0}^{2} \sin^2\frac{k_j}{2}. \tag{5}$$

The wave vectors that matter form the cube $|k_j| \le \pi$, called the *zone* (Figure 3). A real frequency needs the right side of (5) to be at most one, since $\sin^2(\omega/2) \le 1$; otherwise the wave grows exponentially. At the corner $R = (\pi,\pi,\pi)$ of the zone, where the field alternates in sign along all three axes, the sum of sines equals 3 and the right side is $3r^2$. So the face leapfrog has a real frequency for every wave exactly when $3r^2 \le 1$. At $r = 1$, the value that was exact in one dimension, the wave at R has $\cos\omega = -5$: its amplitude is multiplied by about ten at every tick, and the rule is useless. The largest stable choice is $r = 1/\sqrt{3}$, and light then crosses only 0.577 voxel per tick. Numerical analysts know this limit from finite-difference schemes for Maxwell's equations.

Is the limit a defect of this particular rule, or of every rule whose tick reaches only face neighbours? Section 3 shows that it belongs to all of them.

3 How far a tick reaches

To speak about every local rule at once we need one description of them. Let T_s be the shift by the hop $s \in \mathbb{Z}^3$: $(T_s f)(x) = f(x+s)$. A local rule is a finite weighted sum of shifts,

$$Uf = \sum_{s\in F} w_s\, T_s f, \qquad (Uf)(x) = \sum_{s\in F} w_s\, f(x+s),$$

and the finite set $F \subset \mathbb{Z}^3$ of hops with nonzero weight is its stencil. Applying the rule twice shifts by t and then by s, which is a shift by $s+t$, so $U^2 = \sum_{s,t} w_s w_t\, T_{s+t}$. After n ticks the hops are sums of n hops from the stencil of one tick.

Definition 1. A rule is a *face rule* if every hop in its stencil is the rest hop 0 or a unit step $\pm e_j$ along an axis, which are the voxel itself and its six face neighbours. It is a *cube rule* if every hop has all three coordinates in $\{-1,0,1\}$, which are the voxel and all 26 of its neighbours. A rule *reaches* the hop h after n ticks if the weight of T_h in U^n is not zero, so that the field at x after n ticks depends on the starting field at $x+h$. The regions compared below do not change when every hop is replaced by its negative, so it does not matter whether a hop is counted from the voxel that receives the signal or from the voxel that sends it.

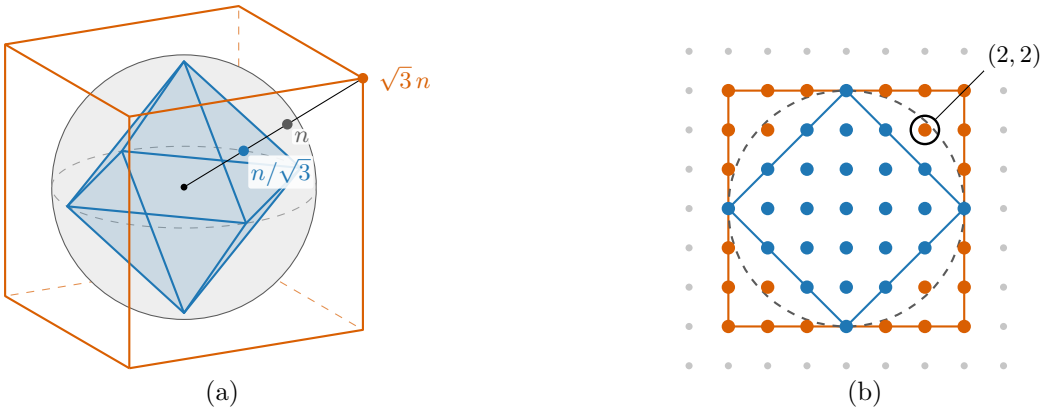

Figure 4: How far a tick reaches. (a) After n ticks a face rule reaches only the octahedron $|x_0|+|x_1|+|x_2| \le n$ (blue), a cube rule fills the cube $\max_j |x_j| \le n$ (orange), and light moving one voxel per tick fills the sphere of radius n (grey). The octahedron touches the sphere only at the six axis points. Along the body diagonal it ends at distance $n/\sqrt{3}$, the sphere at n and the cube at $\sqrt{3}\,n$. (b) The plane $x_2 = 0$ after three ticks. Blue sites are reachable by a face rule; orange sites only by a rule whose tick goes beyond the face neighbours. The circled site $(2, 2)$ is 2.83 voxels away, inside the disc that light covers in three ticks, but a face rule needs four ticks to touch it.

The weights w_s may be numbers, matrices or operators. Multicomponent fields such as the pair (E, B) or a spinor (a field with several components at each site), quantum walks (a quantum particle that hops between sites) and the linear evolution of a free quantum field are all local rules in this sense. The two-step rules used above fit too. The rule $u(n+2) = 2S\,u(n+1) - u(n)$ is a one-step rule on the pair of consecutive ticks,

$$\begin{pmatrix} u(n+2) \\ u(n+1) \end{pmatrix} = \begin{pmatrix} 2S & -1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} u(n+1) \\ u(n) \end{pmatrix},$$

with matrix weights. The entries are $2S$, -1, 1 and 0, so one tick of the pair uses the hops of S and the rest hop, and nothing else. A two-step rule is therefore a face rule or a cube rule whenever S is.

Theorem 2 (Reach)**.** *After n ticks a face rule reaches only hops h with $|h_0|+|h_1|+|h_2| \le n$. After n ticks a cube rule reaches only hops with $|h_j| \le n$ for each j.*

Proof. A hop reached after n ticks is a sum of n hops from the stencil of one tick. For a face rule each of those hops has taxicab length $|s_0|+|s_1|+|s_2|$ at most one, and the taxicab length of a sum is at most the sum of the lengths, so the sum has length at most n. For a cube rule each coordinate of each hop has size at most one, and coordinates add, so each coordinate of the sum has size at most n. □

In words: after n ticks a face rule can have spread information only to the voxels reachable by n steps along grid lines, an octahedron, while a cube rule can have spread it through the whole cube of half-width n. Figure 4 draws both beside the sphere of radius n that light covers at one voxel per tick. The diagonal site (m, m, m) lies at Euclidean distance $\sqrt{3}\,|m|$ but at taxicab distance $3|m|$, so a face rule needs at least $3|m|$ ticks to touch it. Take the site $(2, 2, 0)$: it is $\sqrt{8} \approx 2.83$ voxels from the origin, so light covers it within three ticks, yet a face rule must walk two steps along x_0 and two along x_1 and needs four. A cube rule reaches it in two ticks, by two hops of $(1, 1, 0)$.

Corollary 3 (Face rules are slow)**.** *Suppose a face rule reaches, after every number n of ticks, every body-diagonal site (m, m, m) within Euclidean distance vn of the origin. Then $v \le 1/\sqrt{3}$. In particular no face rule carries light at one voxel per tick in every direction.*

Proof. Fix n and take the largest integer m with $\sqrt{3}\,m \le vn$, so that $m > vn/\sqrt{3} - 1$. By hypothesis the rule reaches (m, m, m), and Theorem 2 then requires $3m \le n$. Together these give $\sqrt{3}\,vn - 3 < n$, that is $v < (n+3)/(\sqrt{3}\,n)$. Letting n grow gives $v \le 1/\sqrt{3}$. □

The hypothesis says what it means for a rule to carry light at speed v in a causal sense: whatever light can reach within n ticks, the rule can reach too. A site that the rule does not reach after n ticks is one where the starting field at the origin has no influence on the field at tick n, so no signal has travelled from the origin to it. Light at speed v does travel from the origin to every diagonal site inside the ball of radius vn, so a rule that matches it must reach them.

The corollary is the Courant–Friedrichs–Lewy condition [1], stated for every face rule at once. A numerical scheme cannot converge if its numerical domain of dependence, the set of starting voxels that can influence a given voxel, is smaller than the domain of dependence of the equation it approximates, which for light at speed v is a ball of radius vn after n ticks. Corollary 3 is that statement for the domain of a face rule. Yee's scheme for Maxwell's equations on a cubic grid has the stability limit $c\,\Delta t \le \Delta x/\sqrt{3}$ [2]; the corollary says that no other choice of face-neighbour weights, components or internal structure can move signals along a body diagonal faster. In physical units, a world whose tick reaches only face neighbours has light speed at most $\ell/(\sqrt{3}\,\tau)$ along a body diagonal. If light is to cross one voxel per tick, the tick must reach further, and Section 4 shows that the corners are enough.

4 A rule that reaches the corners

The idea. In one dimension the exact relation (2) can be written $\cos(\omega/2) = \cos(k/2)$: in half a tick the wave advances by the phase it picks up over half a voxel. In three dimensions a half-voxel step along every axis at once takes a voxel's centre to one of its eight corners. Averaging a plane wave over the eight corners of a voxel multiplies it by

$$P(k) = \frac{1}{8} \sum_{\eta \in \{\pm 1\}^3} e^{i\,\eta \cdot k/2} = \cos\frac{k_0}{2}\,\cos\frac{k_1}{2}\,\cos\frac{k_2}{2}. \tag{6}$$

We take the three-dimensional counterpart of the one-dimensional relation to be $\cos(\omega/2) = P(k)$. This is a guess guided by the one-dimensional case. We show next that a local rule realizes it, and then check what it delivers. Along an axis two of the three cosines equal one, so there it is the exact one-dimensional relation. In the zone P lies between 0 and 1, so the relation has a real solution ω at every wave vector, which is what a tick needs.

To write the rule as a weighted sum of neighbours, let av_j be the average along axis j with weights $\frac{1}{4}, \frac{1}{2}, \frac{1}{4}$ on $x - e_j$, x and $x + e_j$, and set

$$S = 2\,\mathrm{av}_0\,\mathrm{av}_1\,\mathrm{av}_2 - 1, \qquad u(n+2) = 2S\,u(n+1) - u(n). \tag{7}$$

Multiplying out the three averages, the weight of a hop is $2\prod_j a(s_j)$ with $a(0) = \frac{1}{2}$ and $a(\pm 1) = \frac{1}{4}$, minus one at the centre. So S puts weight $-\frac{3}{4}$ on the voxel itself, $\frac{1}{8}$ on each of the six face neighbours, $\frac{1}{16}$ on each of the twelve edge neighbours and $\frac{1}{32}$ on each of the eight corner neighbours (Figure 5). The weights add to one, so a constant field is left unchanged. We call (7) the *corner rule*. A plane wave $e^{i(k\cdot x - \omega n)}$ obeys $u(n+2) = 2Su(n+1) - u(n)$ when $e^{-2i\omega} = 2\hat{S}\,e^{-i\omega} - 1$, where $\hat{S}(k)$

(a) face leapfrog, $S = 1 + \frac{1}{6}\Delta$ (b) corner rule, $S = 2\mathrm{av}_0\mathrm{av}_1\mathrm{av}_2 - 1$

Figure 5: One tick of each rule as a stencil on the $3 \times 3 \times 3$ block around a voxel; disc area is proportional to the size of the weight. (a) The face leapfrog at its largest stable step ($r^2 = \frac{1}{3}$) puts $\frac{1}{6}$ on each face neighbour and nothing on the centre, edges or corners (grey dots). (b) The corner rule puts $\frac{1}{8}$ on faces (blue), $\frac{1}{16}$ on edges (green), $\frac{1}{32}$ on corners (orange) and $-\frac{3}{4}$ on the centre (red ring). In both, the weights sum to one, so a constant field is left unchanged.

is the factor by which S multiplies $e^{ik\cdot x}$. Multiplying by $e^{i\omega}$, as in one dimension, turns this into $\cos\omega = \hat{S}(k)$. Each av_j multiplies a plane wave by $\frac{1}{2}(1 + \cos k_j) = \cos^2(k_j/2)$, so $\hat{S} = 2P(k)^2 - 1$. Since $\cos\omega = 2\cos^2(\omega/2) - 1$, the condition is $\cos^2(\omega/2) = P^2$. In the zone every $\cos(k_j/2)$ is at least zero, and so is $\cos(\omega/2)$ for $0 \le \omega \le \pi$, which leaves

$$\cos\frac{\omega}{2} = P(k) = \prod_j \cos\frac{k_j}{2}. \tag{8}$$

A pulse is a superposition of waves with wave vectors near some k. Its crests move at the phase speed $\omega/|k|$, but the pulse as a whole, and any signal it carries, moves at the *group velocity* $v = \nabla_k\omega$, the vector with components $v_j = \partial\omega/\partial k_j$.

Theorem 4 (The corner rule)**.** *In the zone $|k_j| \le \pi$ the corner rule has the following properties.*

1. *Every wave has a real frequency $\omega = 2\arccos P(k)$, and $0 \le \omega \le \pi$.*
2. *$\omega = 0$ only at $k = 0$, and $\omega = \pi$ exactly where some $|k_j| = \pi$, that is, on the six faces of the zone.*
3. *Along an axis, $\omega(\kappa, 0, 0) = |\kappa|$ exactly.*
4. *At every $k \neq 0$ the group velocity satisfies*

$$|v|^2 = 1 - \frac{\Lambda(k)}{1 - P(k)^2}, \qquad \Lambda = (1 - c_0^2c_1^2)(1 - c_2^2) + c_2^2(1 - c_0^2)(1 - c_1^2) \ \ge 0,$$

with $c_j = \cos(k_j/2)$. Hence $|v| \le 1$ everywhere, with equality along the axes, where $\Lambda = 0$.

Proof. (1) In the zone each $c_j = \cos(k_j/2)$ is at least zero, because $|k_j|/2 \le \pi/2$. So $0 \le P \le 1$, and $\omega = 2\arccos P$ is a real number in $[0, \pi]$ that satisfies (8). (2) If $\omega = 0$ then $P = 1$. Each factor of P is at most one, so each equals one, and every $k_j = 0$. If $\omega = \pi$ then $P = \cos(\pi/2) = 0$, so some factor $\cos(k_j/2)$ vanishes, which in the zone happens only at $|k_j| = \pi$. Conversely $|k_j| = \pi$ makes $P = 0$, so $\cos(\omega/2) = 0$ and $\omega = \pi$. (3) On an axis $P = \cos(\kappa/2)$, and $2\arccos\cos(|\kappa|/2) = |\kappa|$ because $|\kappa|/2 \in [0, \pi/2]$. (4) Differentiate $\cos(\omega/2) = \prod_i c_i$ with respect to k_j:

$$-\tfrac{1}{2}\sin\tfrac{\omega}{2}\,\partial_j\omega = -\tfrac{1}{2}\sin\tfrac{k_j}{2}\prod_{i\neq j} c_i.$$

Since $\sin(\omega/2) = \sqrt{1-P^2}$, this gives $v_j = \sin(k_j/2)\prod_{i\neq j} c_i/\sqrt{1-P^2}$. Square and sum, and write $\sin^2(k_j/2) = 1 - c_j^2$:

$$(1-P^2)\,|v|^2 = \sum_j (1-c_j^2)\prod_{i\neq j} c_i^2 = c_0^2c_1^2 + c_0^2c_2^2 + c_1^2c_2^2 - 3P^2.$$

Multiplying out Λ gives $\Lambda = 1 - c_0^2c_1^2 - c_0^2c_2^2 - c_1^2c_2^2 + 2P^2$, so the right side is $1 - P^2 - \Lambda$ and $|v|^2 = 1 - \Lambda/(1-P^2)$. Each of the two terms of Λ is a product of factors between zero and one, so $\Lambda \geq 0$. If $k_2 = 0$ then $c_2 = 1$ and $\Lambda = (1-c_0^2)(1-c_1^2)$, which vanishes only when $k_0 = 0$ or $k_1 = 0$; if $c_2 \neq 1$, $\Lambda = 0$ forces $c_0 = c_1 = 1$. So $\Lambda = 0$ exactly on the three axes. □

In words: the rule is stable (every wave has a real frequency, though at $\omega = 0$ and $\omega = \pi$ the amplitude of a wave drifts instead of oscillating, as Section 7 shows); the lowest frequency belongs only to the longest waves, so no short wave masquerades as a long one, a fault that lattice theories call doubling; the highest frequency a tick can show, $\omega = \pi$, belongs to the waves on the faces of the zone; light travels at exactly one voxel per tick along the three axes; and a wave packet built from waves of any wavelength has group speed at most one voxel per tick, in every direction.

How special is the corner rule? Consider the two-step rules $u(n+2) = 2S\,u(n+1) - u(n)$ in which S has real weights on the voxel and its 26 neighbours and treats the three axes alike, and the two directions along an axis alike, so that the weight of a hop depends only on how many of its coordinates are nonzero. Such an S puts one weight w_0 on the voxel, w_1 on each of the six face neighbours, w_2 on each of the twelve edge neighbours (which share only an edge with the voxel) and w_3 on each of the eight corner neighbours (which share only a corner). Its symbol, the factor by which it multiplies $e^{ik\cdot x}$, is

$$\hat{S}(k) = a_0 + a_1\sum_j \cos k_j + a_2\sum_{i<j}\cos k_i\cos k_j + a_3\prod_j \cos k_j, \tag{9}$$

with $a_0 = w_0$, $a_1 = 2w_1$, $a_2 = 4w_2$ and $a_3 = 8w_3$: the pair of hops $\pm e_j$ contributes $2\cos k_j$, the four edge hops in a plane contribute $4\cos k_i\cos k_j$, and the eight corner hops contribute $8\prod_j \cos k_j$. As above, the plane waves of $u(n+2) = 2S\,u(n+1) - u(n)$ obey $\cos\omega = \hat{S}(k)$. We ask two things of the rule. First, a constant field is unchanged, so the weights sum to one: $a_0 + 3a_1 + 3a_2 + a_3 = 1$. Second, long waves move one voxel per tick. For small k the symbol is $\hat{S} \approx 1 - \frac{1}{2}|k|^2(a_1 + 2a_2 + a_3)$ while $\cos\omega \approx 1 - \frac{1}{2}\omega^2$, so this condition is $a_1 + 2a_2 + a_3 = 1$. The corner rule has $(a_0, a_1, a_2, a_3) = (-\frac{3}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$ and meets both.

The relation $\cos\omega = \hat{S}(k)$ has a real solution ω only if $-1 \leq \hat{S}(k) \leq 1$. Write $\sigma(k) := 4\sin^2(\omega/2) = 2 - 2\hat{S}(k)$, which is close to ω^2 when ω is small; the condition is that σ lies between 0 and 4. Let $y_j = \sin^2(k_j/2)$. Putting $\cos k_j = 1 - 2y_j$ into (9) and using the two conditions gives

$$\sigma(k) = 4\sum_j y_j - 8\,(a_2 + a_3)\sum_{i<j} y_iy_j + 16\,a_3\,y_0y_1y_2. \tag{10}$$

The first term is $\sigma_{\text{face}}(k) = 4\sum_j \sin^2(k_j/2)$, the symbol of $-\Delta$. It is the whole of σ for a face rule of light speed one, and it overshoots the allowed range of $\sigma \leq 4$ near the edges and corners of the zone: it equals 8 at M and 12 at R. The terms with a_2 and a_3, the weights on edges and corners, are what pull it back.

Theorem 5 (Edge and corner weight)**.** *Take a rule of the form* (9) *with* $a_0 + 3a_1 + 3a_2 + a_3 = 1$ *and* $a_1 + 2a_2 + a_3 = 1$.

1. *Along an axis it is exact:* $\omega(\kappa,0,0)=|\kappa|$ *for* $|\kappa|\le\pi$.

2. *If every wave has a real frequency, then* $a_2+a_3\ge\frac{1}{2}$ *and* $a_3\ge 0$.

3. *The corner rule has* $a_2+a_3=\frac{1}{2}$, *and* $\sigma=\sigma_{\text{corner}}(k):=4\big(1-P(k)^2\big)$.

4. *The only such rule with* $a_3=0$ *and a real frequency at every wave is the* edge rule, *which has* $a_0=-\frac{1}{2}$, $a_1=0$ *and* $a_2=\frac{1}{2}$. *It does have a real frequency everywhere, and at the zone corner R its frequency is* $\omega=0$*: one tick leaves the pattern* $(-1)^{x_0+x_1+x_2}$ *unchanged.*

5. *The rules with* $a_2+a_3=\frac{1}{2}$ *are* $(a_0,a_1,a_2,a_3)=(-\frac{1}{2}-t,\ t,\ \frac{1}{2}-t,\ t)$, *one for each real number* $t=a_3$. *Its symbol is* $(1-4t)$ *times the edge rule's symbol plus* $4t$ *times the corner rule's, so* $t=0$ *is the edge rule and* $t=\frac{1}{4}$ *is the corner rule. It has a real frequency at every wave vector exactly when* $0\le t\le\frac{1}{4}$, *and then* $\omega(R)=2\arcsin(2\sqrt{t})$, *which grows from* 0 *to* π *along the segment.*

Proof. (1) On an axis $y_1=y_2=0$, so (10) gives $\sigma=4\sin^2(\kappa/2)$, that is $\sin^2(\omega/2)=\sin^2(\kappa/2)$, and $\omega=|\kappa|$ in the zone. (2) At the edge midpoint M$=(\pi,\pi,0)$, $(y_0,y_1,y_2)=(1,1,0)$, and (10) gives $\sigma=8-8(a_2+a_3)$. A real frequency needs $\sigma\le 4$, which is $a_2+a_3\ge\frac{1}{2}$. At the zone corner R$=(\pi,\pi,\pi)$, $(y_0,y_1,y_2)=(1,1,1)$ and (10) gives $\sigma=12-24a_2-8a_3$, which must be at least 0, so $3a_2+a_3\le\frac{3}{2}$. With $a_2\ge\frac{1}{2}-a_3$ from the first inequality this gives $\frac{3}{2}-2a_3\le\frac{3}{2}$, that is $a_3\ge 0$. (3) With $a_2=a_3=\frac{1}{4}$, (10) reads $\sigma=4\sum_j y_j-4\sum_{i<j}y_iy_j+4y_0y_1y_2=4\big(1-\prod_j(1-y_j)\big)$, and $\prod_j(1-y_j)=\prod_j\cos^2(k_j/2)=P^2$. (4) Let $a_3=0$. Part (2) gives $a_2\ge\frac{1}{2}$, and the bound at R gives $3a_2\le\frac{3}{2}$, so $a_2=\frac{1}{2}$. The two conditions then fix $a_1=1-2a_2=0$ and $a_0=-\frac{1}{2}$. This rule has $\sigma=4\big(1-\prod_j\cos^2\frac{k_j}{2}-\prod_j\sin^2\frac{k_j}{2}\big)$. Both products are nonnegative and they add to at most one, so $0\le\sigma\le 4$ and every frequency is real. At R the symbol is $\hat{S}=-\frac{1}{2}+\frac{1}{2}\cdot 3=1$, so $\omega=0$, and the plane wave $e^{iR\cdot x}$ is the pattern $(-1)^{x_0+x_1+x_2}$. (5) With $a_2+a_3=\frac{1}{2}$ and $a_3=t$ we get $a_2=\frac{1}{2}-t$, and the two conditions give $a_1=1-2a_2-a_3=t$ and $a_0=1-3a_1-3a_2-a_3=-\frac{1}{2}-t$. Each coefficient of $(1-4t)\cdot(-\frac{1}{2},0,\frac{1}{2},0)+4t\cdot(-\frac{3}{4},\frac{1}{4},\frac{1}{4},\frac{1}{4})$ equals the one just found, so the symbol is the stated combination of the edge and corner symbols. At R the edge symbol is 1 and the corner symbol is $2P^2-1=-1$ (because $P=0$ there), so $\hat{S}(R)=1-8t$. A real frequency at R needs $|1-8t|\le 1$, that is $0\le t\le\frac{1}{4}$. For such t the symbol is a weighted average, with weights $1-4t$ and $4t$ that are nonnegative and add to one, of two symbols that lie in $[-1,1]$ (parts 3 and 4), so it lies in $[-1,1]$ at every wave vector. Finally $\sigma(R)=2-2\hat{S}(R)=16t=4\sin^2(\omega/2)$ gives $\omega(R)=2\arcsin(2\sqrt{t})$. □

In words: to move light at one voxel per tick, a rule must give up some of its face-only symbol at the midpoints of the zone's edges, and the only way to do so is weight on edges and corners. A face rule has $a_2=a_3=0$, breaks the bound $a_2+a_3\ge\frac{1}{2}$ and so has no real frequency at M; Corollary 3 reaches the same verdict on face rules from reach alone. The bound $a_2+a_3\ge\frac{1}{2}$ is met by a whole segment of stable rules (part 5), with the edge rule at one end and the corner rule at the other. The corner rule's $\sigma=4(1-P^2)$ is at most 4 for the simplest possible reason, that $P^2\ge 0$. Edges alone almost suffice, but not quite. The edge rule moves light correctly along the axes and is stable, yet it has two faults. Its hops $(\pm1,\pm1,0)$ and their permutations leave the parity of $x_0+x_1+x_2$ unchanged, so the grid splits into two checkerboard sublattices that never exchange anything, and its symbol is unchanged when k moves by $R=(\pi,\pi,\pi)$, so every wave has a doubled copy at the other end of the zone (the doubling again). Also, at R the edge rule's frequency is zero, the frequency of the longest wave: the shortest wave does not oscillate at all. Every stable rule other than the edge rule has $a_3>0$, so a stable rule whose shortest wave has a nonzero frequency

rule	wavelength	axis	face diagonal	body diagonal
corner rule	20 voxels	0	0.10%	0.14%
corner rule	10 voxels	0	0.41%	0.55%
face leapfrog, $r^2 = \frac{1}{3}$	20 voxels	0.27%	0.07%	0
face leapfrog, $r^2 = \frac{1}{3}$	10 voxels	1.11%	0.28%	0

Table 2: Lag of a plane wave: the relative shortfall of its phase speed below the rule's own light speed, computed from the exact dispersion relations (8) and (5), for waves 20 and 10 voxels long in three directions.

needs corner weight. Along the segment of part 5 the corner weight $a_3 = t$ grows from the edge rule to the corner rule, which has the most corner weight on it.

The lag. A wave of long wavelength moves at the rule's light speed c (one voxel per tick for the corner rule, r for the face leapfrog), and a shorter wave moves a little slower. Call the relative shortfall $1 - \omega/(c|k|)$ of its phase speed the *lag*. It depends on direction, and where it vanishes the rule is exact. To find it for small $|k|$, start from $\sigma = 4\sin^2(\omega/2) = \omega^2 - \omega^4/12 + \cdots$, which inverts to $\omega^2 = \sigma + \sigma^2/12 + O(\sigma^3)$. Next, $4\sin^2(k_j/2) = k_j^2 - k_j^4/12 + \cdots$, so $\sigma_{\text{face}} = |k|^2 - \frac{1}{12}\sum_j k_j^4 + O(|k|^6)$. For the corner rule (10) adds $-4\sum_{i<j} y_i y_j = -\frac{1}{4}\sum_{i<j} k_i^2 k_j^2 + O(|k|^6)$. Collect the terms of fourth order, using $|k|^4 = \sum_j k_j^4 + 2\sum_{i<j} k_i^2 k_j^2$ for $\sigma^2/12$:

$$\omega^2 = |k|^2 - \tfrac{1}{12}\sum_j k_j^4 - \tfrac{1}{4}\sum_{i<j} k_i^2 k_j^2 + \tfrac{1}{12}\Big(\sum_j k_j^4 + 2\sum_{i<j} k_i^2 k_j^2\Big) + O(|k|^6) = |k|^2 - \tfrac{1}{12}\sum_{i<j} k_i^2 k_j^2 + O(|k|^6).$$

Taking the square root,

$$\omega = |k|\Big[1 - \frac{1-Q}{48}\,|k|^2 + O(|k|^4)\Big], \qquad Q = \sum_j \frac{k_j^4}{|k|^4}. \tag{11}$$

Along an axis $Q = 1$ and the lag vanishes, as Theorem 4 says it must. Along a face diagonal $Q = \frac{1}{2}$ and the lag is $|k|^2/96$; along a body diagonal $Q = \frac{1}{3}$ and it is $|k|^2/72$.

The face leapfrog shows the opposite pattern. For it $\sigma = r^2\sigma_{\text{face}}$, and the same expansion gives $\omega = r|k|\big[1 - \frac{1}{24}(Q - r^2)|k|^2 + O(|k|^4)\big]$: a lag of $(1 - r^2)|k|^2/24$ along an axis and $(1 - 3r^2)|k|^2/72$ along a body diagonal. At its largest stable step $r^2 = \frac{1}{3}$ the diagonal lag vanishes, and in fact it vanishes exactly: on $k = (\kappa, \kappa, \kappa)$ the relation (5) reads $\sin^2(\omega/2) = \sin^2(\kappa/2)$, so $\omega = \kappa = |k|/\sqrt{3}$ with $r = 1/\sqrt{3}$. The two rules are mirror images, exact along the axes and along the body diagonal respectively. Table 2 gives the numbers from the exact relations at wavelengths of 20 and 10 voxels.

Figure 6 plots the phase and group speeds out to the edge of the zone along the three directions. The corner rule (orange) stays at one along an axis and slows along the diagonals. Its phase speed falls to $1/\sqrt{3}$ at R, where $\omega = \pi$ and $|k| = \sqrt{3}\,\pi$, and its group speed falls to zero at M and at R. The face leapfrog (blue) stays at its light speed along the body diagonal and slows along the axis. Figure 7 shows the same information as curves of constant frequency in a plane of wave vectors. In the continuum those curves are circles. A curve that lies outside its circle belongs to a rule that needs more wave number for the same frequency, which means slower waves, and the direction in which a rule's curves touch the circles is the direction in which it is exact.

5 A pulse on the lattice

Figure 8 puts both rules to work on the same task: a kick at the origin. The field is zero at tick 0 and equals a smooth lump g at tick 1, that is, $u(0) = 0$ and $u(1) = g$, and the lump then spreads

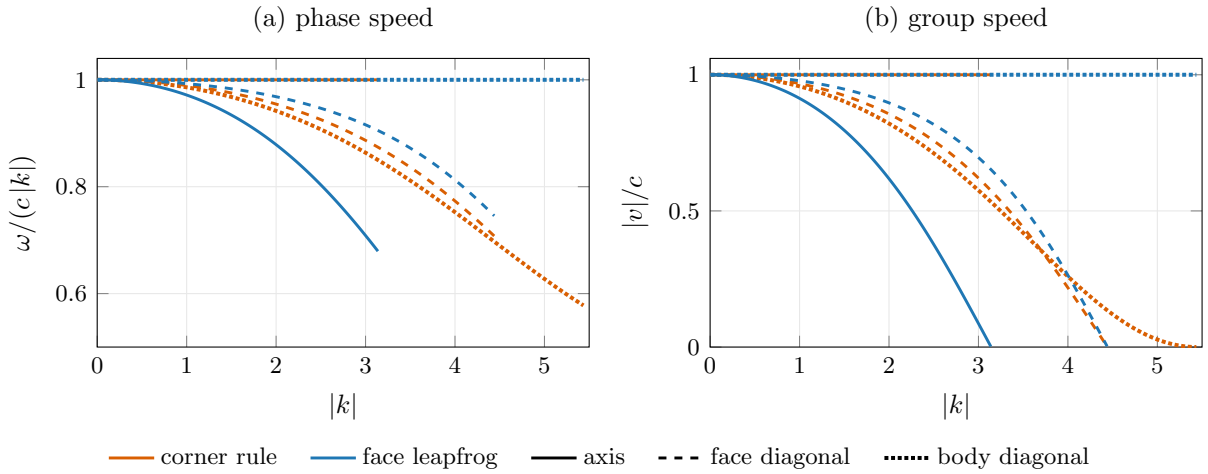

Figure 6: Wave speed against wave number along three directions, for the corner rule (orange) and the face leapfrog at its largest stable step $r = 1/\sqrt{3}$ (blue). Each rule is measured in its own light speed c: one voxel per tick for the corner rule, $1/\sqrt{3}$ for the face leapfrog. Each curve ends at the edge of the zone in its direction. (a) Phase speed. (b) Group speed. The flat lines at one are the corner rule along the axes and the face leapfrog along the body diagonal, which lie on top of each other. Neither group speed exceeds its light speed in these directions; for the corner rule Theorem 4 gives this in every direction.

outward. A lump that occupies a single voxel would contain waves of every wavelength, including the shortest, whose slow and ringing motion would hide the front. So g is a single voxel smoothed by each rule's own averaging, repeatedly. For the corner rule that is the average $\mathrm{av}_0\mathrm{av}_1\mathrm{av}_2$ of Section 4, which replaces each value by a weighted mean with weights $\frac{1}{4}, \frac{1}{2}, \frac{1}{4}$ along each axis, applied three times; its symbol $P(k)^6$ is tiny near the whole surface of the zone. For the face leapfrog it is the average $1 + \Delta/12$ applied six times, whose symbol is tiny near the zone corner R. Each rule then runs for thirty ticks on a periodic grid of 77^3 voxels, large enough that nothing wraps around. The result is a direct computation with the rules of Sections 2 and 4, with no further approximation.

Because a spherical pulse in three dimensions thins out as one over its distance from the origin, panel (c) of the figure plots distance times field. Sites along an axis are one voxel apart and sites along a body diagonal $\sqrt{3}$ apart, so a peak can be placed only to within that spacing. Under the face leapfrog the quantity peaks at $30/\sqrt{3} \approx 17.3$ voxels from the origin along the body diagonal and at 17 along an axis, and under the corner rule at 30 voxels along the axes and at 29.4 along the body diagonal. Each peak agrees to within the site spacing with the rule's front, the distance that light at the rule's speed covers in thirty ticks, so the lag of the previous section is too small to show in the position of a peak after thirty ticks. The pulse starts at tick 1 as g, which extends 6 voxels along grid lines for the face leapfrog and 3 voxels along each axis for the corner rule, and in the other 29 ticks Theorem 2 lets it spread by at most 29 more. So the face-leapfrog field lies within taxicab distance 35 of the origin and the corner-rule field within 32 along each axis, and the computed fields extend exactly that far. The corner-rule field outside the light sphere, the sphere of radius 30, is small but not zero. At voxels more than 34 from the origin, four beyond the sphere, which exist only off the axes because the field ends at 32 along them, the field is at most 1.5% of the largest field value anywhere, and those voxels hold 2×10^{-5} of the total of the squared field over all voxels. At voxels more than 38 from the origin the field is below 10^{-5} of the largest value, and they hold about 10^{-12} of the total. Each rule delivers a compact pulse along its exact direction and leaves a ringing tail along its slow one, a trail of oscillations behind the front made of the shorter waves in the pulse, which travel slower. This is the lag of the previous section at work: small for

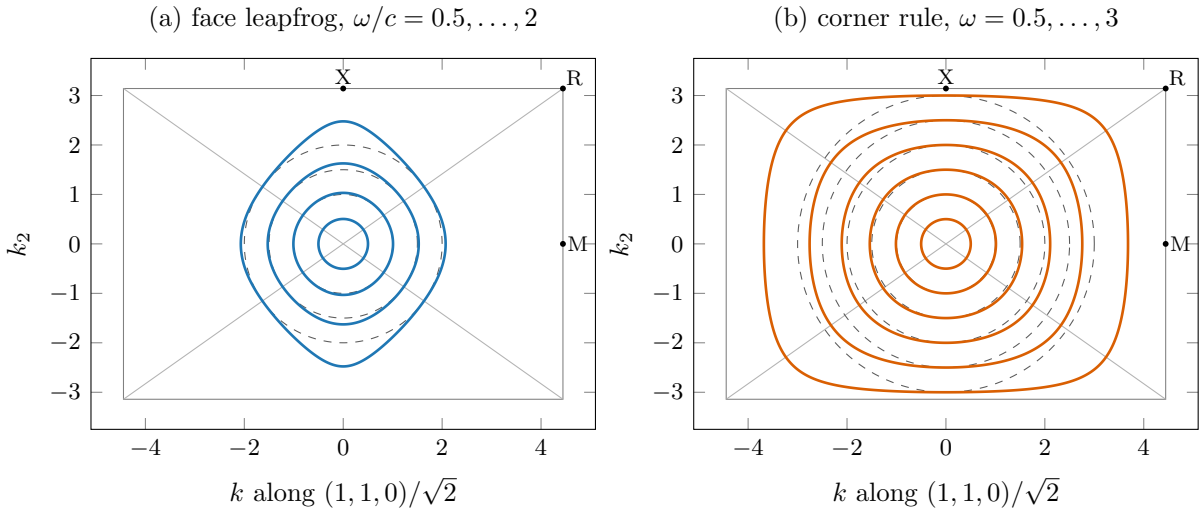

Figure 7: Curves of constant frequency in the plane of wave vectors that contains an axis (k_2, vertical), a face diagonal (horizontal) and the body diagonals (thin lines to the zone corners R). X, M and R are the points of Figure 3. Dashed circles are the continuum at the same frequencies, in each rule's own units. The rectangle is the edge of the zone. Where a curve lies outside its circle, a wave at that frequency needs more wave number than in the continuum and travels more slowly. (a) The face leapfrog matches its circles along the body diagonals and falls behind along the axis, where its frequencies stop at point X: there $\sin(\omega/2) = r = 1/\sqrt{3}$, so $\omega \approx 1.23$ and $\omega/c \approx 2.13$. (b) The corner rule matches its circles along the axis and falls behind along the diagonals; its highest curves approach the zone edge, where $\omega = \pi$.

the long waves that make the peak, larger for the short waves that make the tail. The picture is a plane section through the pulse, with colour giving $\log_{10}$ of the field's size relative to its peak, so that values as small as 10^{-12} of the peak are visible next to the front.

6 Maxwell's equations with an exact Gauss law

In empty space, with the speed of light set to one, the electric field E and the magnetic field B obey four equations: Gauss's law $\nabla \cdot E = 0$, the absence of magnetic charge $\nabla \cdot B = 0$, Faraday's law $\partial_t B = -\nabla \times E$ and Ampère's law $\partial_t E = \nabla \times B$. The first two are constraints on the fields at one instant. The last two say how the fields change, and they keep the constraints true: the divergence of a curl vanishes, so Faraday's law preserves $\nabla \cdot B$ and Ampère's preserves $\nabla \cdot E$. A vector potential A makes this economical. Put $E = -\partial_t A$ and $B = \nabla \times A$, with no scalar potential (the temporal gauge). Then $\nabla \cdot B = 0$ and Faraday's law hold identically, Ampère's law becomes $\partial_t^2 A = -\nabla \times \nabla \times A$, which is the wave equation $\partial_t^2 A = \nabla^2 A$ when $\nabla \cdot A = 0$, and Gauss's law $\nabla \cdot E = 0$ says $\partial_t(\nabla \cdot A) = 0$, so $\nabla \cdot A$ keeps its starting value, which we take to be zero.

We copy this on the lattice. For a field f, the forward and backward differences along axis j are $\partial_j^+ f(x) = f(x + e_j) - f(x)$ and $\partial_j^- f(x) = f(x) - f(x - e_j)$. From them come the divergence $\nabla^\pm \cdot \mathbf{f} = \sum_j \partial_j^\pm f_j$ and the curl $(\nabla^\pm \times \mathbf{f})_j = \partial_{j+1}^\pm f_{j+2} - \partial_{j+2}^\pm f_{j+1}$ of a vector field $\mathbf{f} = (f_0, f_1, f_2)$, with indices counted modulo three. The face Laplacian $\Delta = \sum_j \partial_j^+ \partial_j^-$ is the sum of second differences of (3). Differences are weighted sums of shifts, and shifts commute, so every one of these operators commutes with every other. Therefore $\nabla^+ \cdot (\nabla^+ \times \mathbf{f}) = 0$ exactly, for the same reason as in the continuum: each mixed difference $\partial_i^+ \partial_j^+ f_k$ appears twice, with opposite signs.

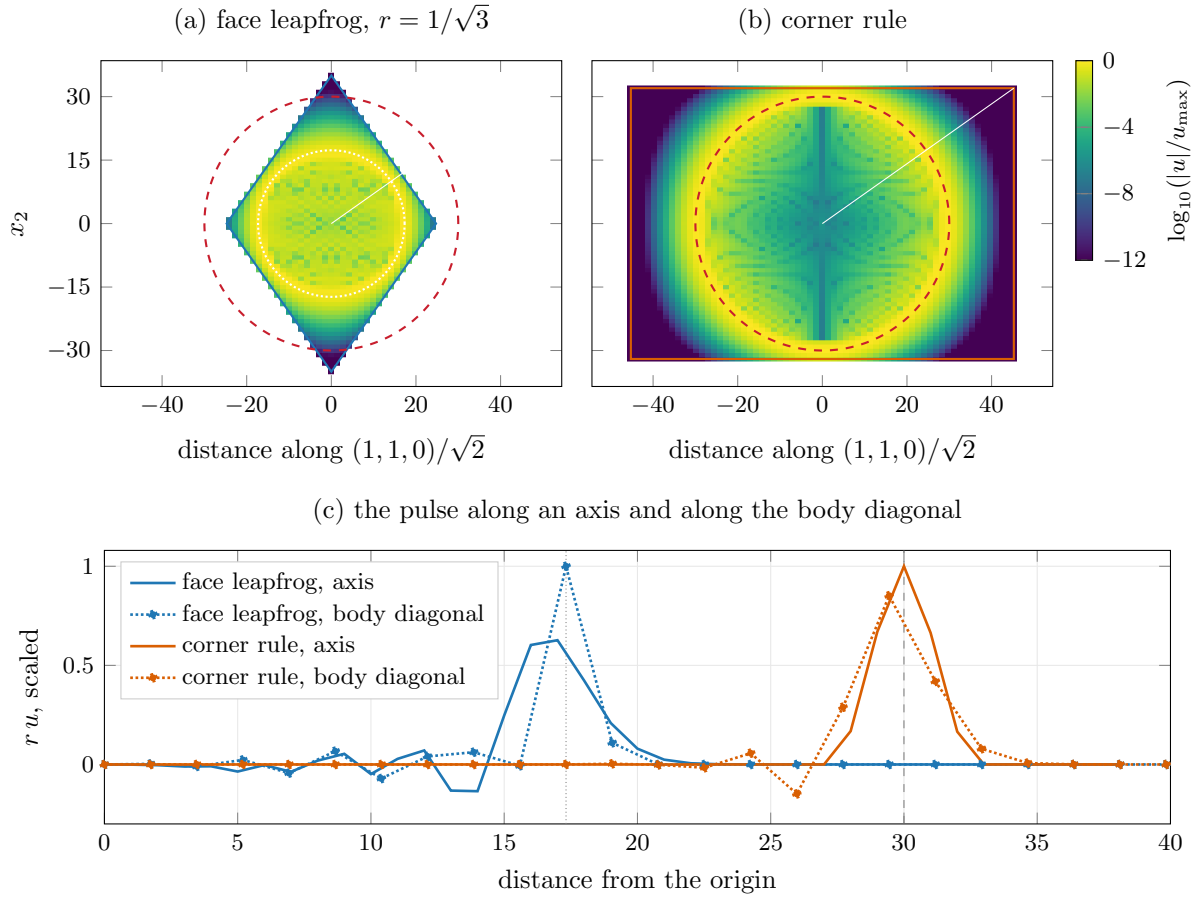

Figure 8: The same kick under the two rules, thirty ticks after launch. The field is zero at tick 0 and equals g at tick 1, with g made by applying each rule's own averaging to a single voxel, so that it carries almost nothing near the edge of the zone. (a, b) The field in the plane $x_0 = x_1$, which contains the x_2 axis (vertical), a face diagonal (horizontal) and the body diagonal (thin white line). Colour is $\log_{10}$ of the field's size relative to its peak; white means exactly zero. The red dashed circle has radius 30, the distance light covers in thirty ticks at one voxel per tick. (a) The face leapfrog: its front (dotted circle) is at $30/\sqrt{3}$, and the field vanishes outside the octahedron $\sum_j |x_j| \le 35$ (blue), the 29 ticks that Theorem 2 allows plus the 6 voxels that g occupies. (b) The corner rule: its front lies on the red circle along the axes and, to within the site spacing, along the body diagonal, and the field vanishes outside the cube $\max_j |x_j| \le 32$ (orange), the 29 ticks plus the 3 voxels that g occupies. (c) Distance times field along the x_2 axis (solid) and the body diagonal (dots, one per diagonal site), scaled for each rule so that the larger of its two peaks is one. Each rule delivers a compact pulse along its exact direction and leaves a ringing tail along its slow one.

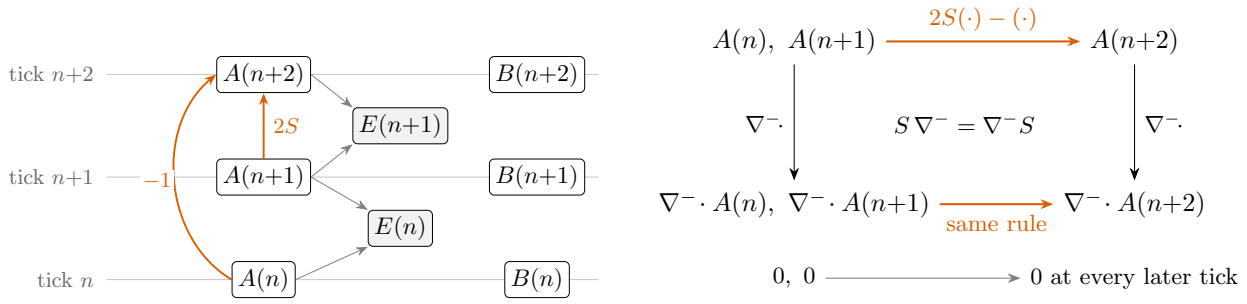

(a) where the fields live (b) why Gauss's law is exact

Figure 9: (a) The potential A and the magnetic field $B = \nabla^+ \times A$ sit on ticks; the electric field $E(n) = A(n) - A(n+1)$ sits between them. Each tick computes $A(n+2)$ from the two before it. (b) Taking the backward divergence commutes with the update, so $\nabla^- \cdot A$ obeys the same two-tick rule as A. Started from zero on two ticks, it stays zero, and so does $\nabla^- \cdot E$.

Let $A(n, x) \in \mathbb{R}^3$ be a vector potential on the lattice, and define the lattice fields

$$E(n) = A(n) - A(n+1), \qquad B(n) = \nabla^+ \times A(n), \tag{12}$$

the lattice versions of $E = -\partial_t A$ and $B = \nabla \times A$. Evolve each component of A by

$$A(n+2) = 2S\,A(n+1) - A(n), \tag{13}$$

where S is any local rule, a finite weighted sum of shifts; the update needs the potential on two ticks, $A(0)$ and $A(1)$, to start. Because shifts commute, S commutes with every lattice difference. The corner rule is one such S, and so is the face leapfrog. Figure 9 shows where the fields sit in time and why the constraint survives every tick.

Theorem 6 (Lattice Maxwell equations)**.** *Let S be any local rule and suppose $\nabla^- \cdot A(0) = \nabla^- \cdot A(1) = 0$. Then for every tick n:*

1. *$\nabla^- \cdot A(n) = 0$, and therefore Gauss's law $\nabla^- \cdot E(n) = 0$ holds exactly;*
2. *$\nabla^+ \cdot B(n) = 0$ exactly;*
3. *Faraday's law $B(n+1) - B(n) = -\nabla^+ \times E(n)$ holds exactly;*
4. *Ampère's law holds with one lattice term,*

$$E(n+1) - E(n) = \nabla^- \times B(n+1) + \mathcal{D}A(n+1), \qquad \mathcal{D} = 2(1-S) + \Delta.$$

Proof. (1) S commutes with ∇^-, so applying $\nabla^- \cdot$ to (13) gives $\nabla^- \cdot A(n+2) = 2S\,\nabla^- \cdot A(n+1) - \nabla^- \cdot A(n)$: the divergence obeys the same recursion as A. It starts from zero on two ticks, and a recursion that starts from zero stays zero. Then $\nabla^- \cdot E(n) = \nabla^- \cdot A(n) - \nabla^- \cdot A(n+1) = 0$. (2) $B = \nabla^+ \times A$ and $\nabla^+ \cdot \nabla^+ \times = 0$. (3) $B(n+1) - B(n) = \nabla^+ \times \big(A(n+1) - A(n)\big) = -\nabla^+ \times E(n)$, by linearity. (4) From (13),

$$E(n+1) - E(n) = 2A(n+1) - A(n+2) - A(n) = 2(1-S)A(n+1).$$

On the other side, the lattice identity $\nabla^- \times \nabla^+ \times \mathbf{f} = \nabla^+(\nabla^- \cdot \mathbf{f}) - \Delta \mathbf{f}$ (checked below), applied to $\mathbf{f} = A(n+1)$ with $\nabla^- \cdot A(n+1) = 0$, gives $\nabla^- \times B(n+1) = -\Delta A(n+1)$. The difference of the two sides is $\big[2(1-S) + \Delta\big]A(n+1)$.

To check the identity, note that a finite weighted sum of shifts is zero exactly when its symbol vanishes for every wave vector, so it is enough to compare both sides on plane waves $\varepsilon\, e^{ik\cdot x}$ with a constant vector $\varepsilon \in \mathbb{C}^3$. There ∇^+ acts as multiplication by the vector z with components $z_j = e^{ik_j} - 1$, and ∇^- as multiplication by $-\bar{z}$, because $1 - e^{-ik_j} = -\overline{z_j}$. The rule $p \times (q \times s) = q\,(p \cdot s) - s\,(p \cdot q)$ for vectors p, q, s, applied with $p = -\bar{z}$, $q = z$, $s = \varepsilon$, turns the left side of the identity into $-\bar{z} \times (z \times \varepsilon) = -z\,(\bar{z} \cdot \varepsilon) + |z|^2\,\varepsilon$. On the right, $\nabla^+(\nabla^- \cdot \mathbf{f})$ becomes $-z\,(\bar{z} \cdot \varepsilon)$ and $-\Delta\mathbf{f}$ becomes $\sum_j |e^{ik_j} - 1|^2\,\varepsilon = |z|^2\,\varepsilon$. The two sides agree. □

In words: if the potential starts free of electric charge, meaning $\nabla^- \cdot A = 0$ on the first two ticks, it stays so, and then Gauss's law, the absence of magnetic charge and Faraday's law are exact for every local rule, whatever its weights, because the constraints come from the structure of differences and not from the dynamics. Only Ampère's law depends on the choice of S, and the operator $\mathcal{D}$ measures how far the choice is from the lattice wave equation: it vanishes exactly when $S = 1 + \Delta/2$, the face leapfrog at $r = 1$ (Section 8). A rule whose light is slower has an exact law of its own. For the face leapfrog $2(1 - S) = -r^2\Delta$ and $\nabla^- \times B(n+1) = -\Delta A(n+1)$, so $E(n+1) - E(n) = r^2\,\nabla^- \times B(n+1)$, which is Ampère's law for light of speed r. For the corner rule $\mathcal{D}$ can be computed.

Theorem 7 (Corner-rule Maxwell waves)**.** *With S the corner rule, let k be a wave vector, let ω satisfy* (8)*, and let $\varepsilon \in \mathbb{C}^3$ satisfy $\sum_j (1 - e^{-ik_j})\varepsilon_j = 0$. Then the real wave $A(n,x) = \mathrm{Re}\left[\varepsilon\, e^{i(k\cdot x - n\omega)}\right]$ has $\nabla^- \cdot A = 0$ at every tick (it is transverse), solves* (13) *exactly, and the lattice term acts on it as multiplication by*

$$d(k) = \sigma_{\mathrm{corner}}(k) - \sigma_{\mathrm{face}}(k) = -4\,(y_0y_1 + y_0y_2 + y_1y_2) + 4\,y_0y_1y_2, \qquad y_j = \sin^2\tfrac{k_j}{2},$$

where $\sigma_{\mathrm{face}} = 4\sum_j y_j$ is the symbol of $-\Delta$ and $\sigma_{\mathrm{corner}} = 4\sin^2(\omega/2) = 4(1 - P^2)$ is the symbol of $2(1 - S)$. Along an axis $d = 0$, so Ampère's law holds exactly for waves travelling along the axes.

Proof. Write $\varepsilon_j = |\varepsilon_j| e^{i\phi_j}$. Each component of A is then a real plane wave $|\varepsilon_j| \cos(k \cdot x - n\omega + \phi_j)$. Each av_j multiplies $\cos(k \cdot x + \theta)$ by $\cos^2(k_j/2)$, so S multiplies it by $2P^2 - 1 = \cos\omega$, and $\cos(\theta - \omega) + \cos(\theta + \omega) = 2\cos\omega\cos\theta$ is the recursion. The backward divergence is $\mathrm{Re}\left[\sum_j (1 - e^{-ik_j})\varepsilon_j\, e^{i(k\cdot x - n\omega)}\right]$, which vanishes because the sum does. The face Laplacian multiplies a plane wave by $-\sigma_{\mathrm{face}}$, and $2(1 - S)$ multiplies it by $2(1 - \cos\omega) = 4\sin^2(\omega/2) = \sigma_{\mathrm{corner}}$. Writing $P^2 = \prod_j (1 - y_j)$ and expanding gives d. On an axis two of the y_j vanish. □

At small wave number $d(k) \approx -\frac{1}{4}\sum_{i<j} k_i^2 k_j^2$, a fourth-order correction to the second-order Maxwell operator, and $d \le 0$ throughout the zone because $y_0y_1y_2 \le y_0y_1$. A concrete example is $k = (\kappa, 0, 0)$ with $\varepsilon = (0, 1, 0)$, the wave travelling along x_0 and polarized along x_1: it is transverse, nonzero, obeys all four laws exactly and moves one voxel per tick. Figure 11(a) shows the size of the lattice term relative to the Maxwell operator along three directions.

7 The vacuum of each mode

Modes. Take the lattice to be periodic with N voxels along each axis, so that the allowed wave numbers $k_j = 2\pi m_j/N$ form a finite set and the potential is a finite sum. The potential is then a sum of plane waves $\varepsilon\, e^{ik\cdot x}$, one for each wave vector k, with complex amplitude vectors $\varepsilon \in \mathbb{C}^3$; a real field pairs each wave with its complex conjugate at $-k$, so counting k and $-k$ separately there is one real oscillator for each wave vector. For a wave vector $k \neq 0$ the transversality condition $\sum_j (1 - e^{-ik_j})\varepsilon_j = 0$ of Theorem 7, which is Gauss's law $\nabla^- \cdot A = 0$ for the wave, is

one linear condition on the three components of ε, so it leaves a plane of allowed vectors: two independent polarizations. At $k = 0$ every coefficient $1 - e^{-ik_j}$ vanishes and the condition says nothing, so the three components of a uniform potential are all allowed; those modes are treated below. The transverse potential therefore has one wave for each wave vector $k \neq 0$ and each of the two polarizations, together with the uniform modes. The update (13) multiplies the plane wave $e^{ik\cdot x}$ by the symbol $\hat{S}(k) = \cos\omega$. So each such *mode* has an amplitude $x(n)$, the coefficient that multiplies the mode's fixed spatial pattern at tick n (a number, not a position in space). The amplitude evolves on its own and obeys

$$x(n+2) + x(n) = 2\cos\omega\; x(n+1). \tag{14}$$

A mode is one oscillator. Quantizing the field means quantizing each mode, and the vacuum of the field is the product of the vacua of its modes.

The action of one tick. The *action* of a motion $x_0, x_1, x_2, \ldots$ of one mode is the sum over ticks of $L(x_n, x_{n+1})$, where

$$L(x, x') = \tfrac{1}{2}(x' - x)^2 - \tfrac{1}{2}\, q\, x^2, \qquad q = 2 - 2\cos\omega = 4\sin^2\tfrac{\omega}{2}, \tag{15}$$

and x and x' are the amplitudes at consecutive ticks. The motions the mode follows are those for which the action is stationary when any one x_n is varied, the lattice form of the principle of least action. Varying x_n in $L(x_{n-1}, x_n) + L(x_n, x_{n+1})$ gives $(x_n - x_{n-1}) - (x_{n+1} - x_n) - q\, x_n = 0$, which is (14); so (14) is the Euler–Lagrange equation of this action. The first term of L is the electric energy of the mode, since $x' - x$ is minus the electric amplitude between the two ticks. The second term is the counterpart of potential energy, with the coefficient q fixed by the requirement that the recursion be (14). It plays the role that $\frac{1}{2}\omega^2 x^2$ plays in the continuum Lagrangian $\frac{1}{2}\dot{x}^2 - \frac{1}{2}\omega^2 x^2$. It equals the magnetic energy $\frac{1}{2}\sigma_{\text{face}}(k)\, x^2$ of the mode, computed below, exactly when the lattice term of Ampère's law vanishes on the mode, that is, when $q = \sigma_{\text{face}}(k)$; for the corner rule $q = \sigma_{\text{corner}}(k)$, and the two differ by $d(k)$ (Theorem 7). In physical units the action carries an overall constant factor m, set by the units and the volume of the lattice; we take $m = 1$. Restoring m multiplies every variance of the field amplitude by $1/m$ and every variance of its momentum by m, so the ratios of lattice to continuum values below do not depend on it.

Position, momentum and the one-tick map. Quantum mechanics treats the amplitude x and its conjugate momentum p as operators with $[x, p] = i\hbar$. In classical mechanics the momentum conjugate to x is the derivative of the Lagrangian with respect to the velocity. A step from one tick to the next has two ends, and $L(x, x')$ gives a momentum at each: at the earlier end $p = -\partial L/\partial x = (x' - x) + q\, x$, and at the later end $p' = \partial L/\partial x' = x' - x$. That the momentum at a tick is the same whether it is read as the later end of one step or the earlier end of the next is exactly the recursion (14). Given (x, p) at one tick, the first equation fixes $x' = (1 - q)\, x + p$ and the second then gives $p' = x' - x = -q\, x + p$:

$$\begin{pmatrix} x' \\ p' \end{pmatrix} = \Phi \begin{pmatrix} x \\ p \end{pmatrix}, \qquad \Phi = \begin{pmatrix} 1 - q & 1 \\ -q & 1 \end{pmatrix}. \tag{16}$$

Φ has determinant one, as a map that preserves $[x, p] = i\hbar$ must; in one degree of freedom that is the same as conserving phase-space area. Its trace is $2 - q = 2\cos\omega$, so for $0 < q < 4$ its eigenvalues are $e^{\pm i\omega}$ and it turns phase space through the angle ω in suitable coordinates. The momentum is minus the electric amplitude between ticks: $p(n+1) = -E(n)$.

What the vacuum is. The spread of a state is recorded by the covariances $V = \langle x^2 \rangle$, $C = \frac{1}{2}\langle xp + px \rangle$ and $W = \langle p^2 \rangle$, where $\langle \cdot \rangle$ is the expected value in the state and the means of x and p are taken to be zero. The uncertainty principle requires $VW - C^2 \geq \hbar^2/4$ of every state. A *Gaussian* state is one whose Wigner function, the quantum counterpart of a joint distribution of x and p, is a bell curve, fixed by its means and covariances; it is *pure* (a single wave function, not a mixture) exactly when $VW - C^2 = \hbar^2/4$, the least spread the principle allows. In the continuum the vacuum is the lowest-energy state, and it is the pure Gaussian state that time evolution leaves unchanged. A lattice with ticks has no continuous time and no ordinary energy: its evolution is the one-tick map, and a frequency means something only up to multiples of 2π (Section 2). So we define the vacuum in the way that carries over: it is the pure Gaussian state whose covariances the one-tick map leaves unchanged. For the ordinary oscillator this returns the ground state, because the flow rotates phase space and only the circular Gaussian survives a rotation. A tick sends the covariance matrix $\Sigma = \begin{pmatrix} V & C \\ C & W \end{pmatrix}$ to $\Phi \Sigma \Phi^{\mathsf{T}}$.

Theorem 8 (Mode vacuum)**.** *Let $q = 4\sin^2(\omega/2)$ with $0 \leq q \leq 4$. The covariances invariant under the one-tick map are exactly $(V, C, W) = V\,(1,\, q/2,\, q)$ with $V > 0$, and their determinant is $V^2\, q(4-q)/4 = V^2 \sin^2 \omega$.*

1. *For $0 < q < 4$, that is $0 < \omega < \pi$, exactly one invariant covariance is pure, with determinant $(\hbar/2)^2$:*
$$\langle x^2 \rangle = \frac{\hbar}{2\sin\omega}, \qquad \langle p^2 \rangle = \hbar \tan\frac{\omega}{2}.$$
2. *For $q = 0$ or $q = 4$, that is $\omega = 0$ or $\omega = \pi$, every invariant covariance has determinant zero, below the $\hbar^2/4$ that the uncertainty principle requires of every state. So no Gaussian state, pure or mixed, is left unchanged by the tick.*
3. *For $q = 4$ the tick map is a shear: $\Phi^n(x,p) = (-1)^n\big((1+2n)\,x - n\,p,\ (1-2n)\,p + 4n\,x\big)$, so the amplitude obeys $x(n) = (-1)^n\big[x(0) + n\,(2x(0) - p(0))\big]$. For $q = 0$ it is the shear of a free particle, $\Phi^n(x,p) = (x + n\,p,\ p)$, and $x(n) = x(0) + n\,p(0)$.*

Proof. Write out $\Phi\Sigma\Phi^{\mathsf{T}}$ with the rows $(1-q,\ 1)$ and $(-q,\ 1)$ of Φ. The pp entry is $q^2V - 2qC + W$, and setting it equal to W gives $q\,(qV - 2C) = 0$, so $C = \frac{q}{2}V$ when $q \neq 0$. The xp entry is $-q(1-q)V + (1-2q)C + W$; setting it equal to C and using $C = \frac{q}{2}V$ gives $W = qV$. The xx entry, $(1-q)^2V + 2(1-q)C + W$, then equals V identically. For $q = 0$ the pp equation says nothing, the xp entry gives $W = 0$ and the xx entry then gives $C = 0$, which is $V\,(1, q/2, q)$ at $q = 0$ again. The determinant is $VW - C^2 = qV^2 - \frac{q^2}{4}V^2 = V^2\,\frac{q(4-q)}{4}$, and with $q = 4\sin^2(\omega/2)$ we get $q(4-q)/4 = 4\sin^2\frac{\omega}{2}\cos^2\frac{\omega}{2} = \sin^2\omega$. For $0 < q < 4$ the determinant is positive, and setting it to $(\hbar/2)^2$ fixes $V = \hbar/(2\sin\omega)$, and then $W = qV = \hbar\,\dfrac{4\sin^2(\omega/2)}{2\sin\omega} = \hbar\tan\dfrac{\omega}{2}$. For $q = 0$ and $q = 4$ the determinant is $V^2\,q(4-q)/4 = 0$, which is less than $\hbar^2/4 > 0$, so no invariant covariance describes a state. (A state left unchanged must also have unchanged means m, which means $\Phi m = m$; at $q = 4$ the matrix $\Phi - I$ has determinant 4, so $m = 0$. At $q = 0$ the means need not vanish, but the covariance already excludes every state.) For (3), at $q = 4$ we have $\Phi = \begin{pmatrix} -3 & 1 \\ -4 & 1 \end{pmatrix}$, and $N = \Phi + I = \begin{pmatrix} -2 & 1 \\ -4 & 2 \end{pmatrix}$ satisfies $N^2 = 0$. Since $\Phi = -(I - N)$, $\Phi^n = (-1)^n(I - nN)$, which is the stated formula. At $q = 0$, $\Phi = I + N_0$ with $N_0 = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$ and $N_0^2 = 0$, so $\Phi^n = I + nN_0$. □

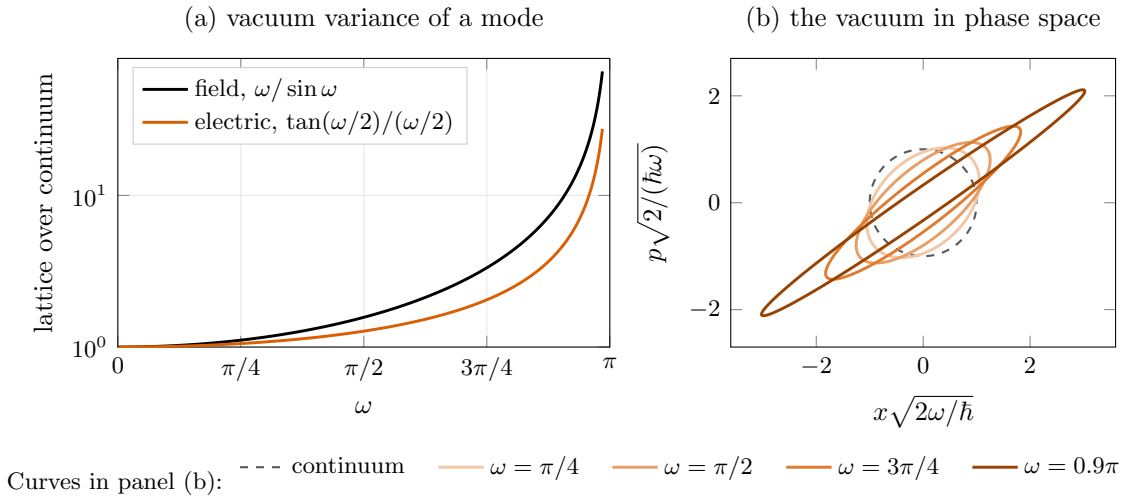

Figure 10: (a) Vacuum variances of a mode with tick frequency ω, divided by those of the continuum oscillator at the same frequency: the field variance $\hbar/(2\sin\omega)$ over $\hbar/(2\omega)$, and the electric variance $\hbar\tan(\omega/2)$ over $\hbar\omega/2$. Both start at one and grow without bound as ω approaches π, the frequency of the zone-face modes. (b) The vacuum of Theorem 8 as an uncertainty ellipse, in coordinates scaled so that the continuum ground state at the same frequency is the unit circle (dashed). Every ellipse has the area of the circle, because each vacuum is pure. As ω approaches π the ellipse stretches and thins toward a line; at $\omega = \pi$ no pure invariant state remains.

In words: the tick map alone fixes the vacuum's spread. As $\omega \to 0$ the two variances become $\hbar/(2\omega)$ and $\hbar\omega/2$, those of the continuum oscillator. At finite ω both are larger, and both grow without bound as ω approaches π. At $\omega = \pi$ the tick map is no longer a rotation of phase space but a shear, like the motion of a free particle with a sign flip at every tick, and the spread of any state grows without limit, which is why no state is left unchanged. At $\omega = 0$ it is a shear without the sign flip, the motion of a free particle, and again no state is left unchanged; this is the case of the uniform modes, whose potential is constant in space and whose electric field is constant, with no oscillation. Figure 10 shows how the variances depart from the continuum as the frequency rises, and the shape of the vacuum in the phase plane.

Electric and magnetic fluctuations. The vacuum of a mode fixes $\langle x^2\rangle$ and $\langle p^2\rangle$, the field and electric spreads. The magnetic field of the mode follows from its amplitude. For a transverse mode with unit polarization ε the magnetic amplitude is $\nabla^+ \times A_k = z \times \varepsilon\, x$, where z is the vector with components $z_j = e^{ik_j} - 1$ (the same z as in the proof of Theorem 6). Lagrange's identity $|z \times \varepsilon|^2 = |z|^2|\varepsilon|^2 - |\bar{z}\cdot\varepsilon|^2$ gives $|\nabla^+ \times A_k|^2 = \sigma_{\text{face}}(k)\,|A_k|^2$, because the transversality condition is $\bar{z}\cdot\varepsilon = 0$ and $|z|^2 = \sum_j |e^{ik_j} - 1|^2 = \sigma_{\text{face}}(k)$. So the magnetic spread is $\sigma_{\text{face}}(k)\langle x^2\rangle$ and the electric spread is $\langle p^2\rangle = q\langle x^2\rangle$, and for the corner rule $q = \sigma_{\text{corner}}(k)$. Three consequences follow.

Corollary 9 (Corner-rule vacuum)**.** *The corner rule has three vacuum properties.*

1. *The ratio of magnetic to electric vacuum variance in a mode is* $\sigma_{\text{face}}(k)/\sigma_{\text{corner}}(k) = 1 - d(k)/\sigma_{\text{corner}}(k)$. *Along the axes it is exactly one, so electric and magnetic fluctuations are equal there.*

2. *Along an axis the electric variance* $\hbar\tan(\kappa/2)$ *exceeds the continuum value* $\hbar\kappa/2$ *for every* $0 < \kappa < \pi$, *by the relative amount* $\kappa^2/12 + O(\kappa^4)$.

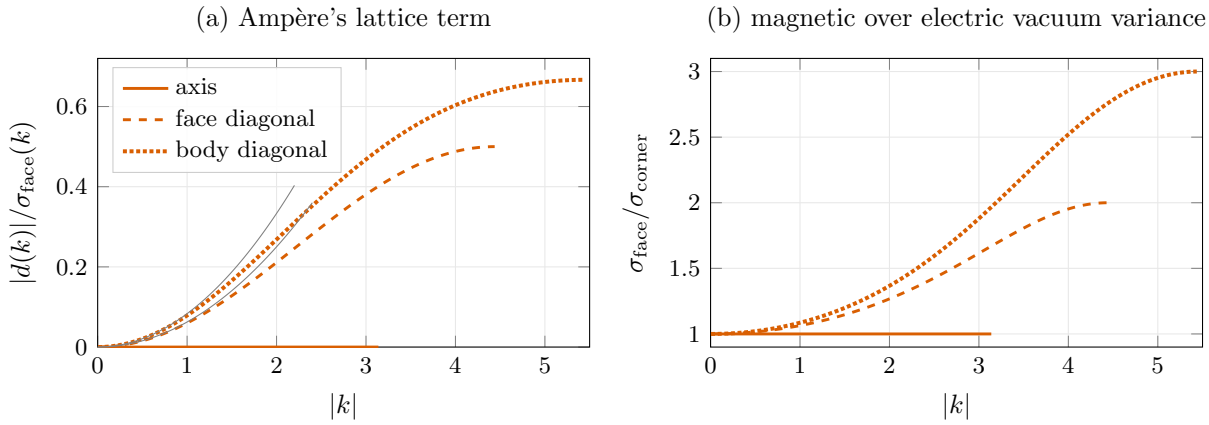

Figure 11: Where the corner rule departs from the continuum Maxwell field, along an axis (solid), a face diagonal (dashed) and the body diagonal (dotted). (a) Ampère's lattice term relative to the Maxwell operator. It is zero along the axes and grows like $|k|^2/16$ and $|k|^2/12$ along the two diagonals at small wave number (thin grey curves). (b) The ratio of magnetic to electric vacuum variance in a mode. It is exactly one along the axes and above one elsewhere, because $d \le 0$.

3. On the zone faces, where $P(k) = 0$, the modes have $q = 4$ and no pure vacuum exists.

Proof. (1) The magnetic variance is $\sigma_{\text{face}}\langle x^2\rangle$ and the electric variance is $\langle p^2\rangle = q\langle x^2\rangle$ with $q = \sigma_{\text{corner}}$, so their ratio is $\sigma_{\text{face}}/\sigma_{\text{corner}}$; and $\sigma_{\text{corner}} = \sigma_{\text{face}} + d$ by Theorem 7. (2) On an axis $\omega = \kappa$ and $\tan y > y$ for $0 < y < \pi/2$, and $\tan y = y(1 + y^2/3 + \cdots)$ with $y = \kappa/2$. (3) $P = 0$ gives $q = 4\big(1 - P^2\big) = 4$, and Theorem 8 applies. □

Along the face and body diagonals everything has a closed form. Write $y = \sin^2(t/2)$ for the wave vectors $(t, t, 0)$ and (t, t, t):

	face diagonal $(t,t,0)$	body diagonal (t,t,t)
σ_{face}	$8y$	$12y$
σ_{corner}	$8y - 4y^2$	$12y - 12y^2 + 4y^3$
$\lvert d\rvert/\sigma_{\text{face}}$	$y/2$	$y - y^2/3$
$\sigma_{\text{face}}/\sigma_{\text{corner}}$	$2/(2-y)$	$3/(3 - 3y + y^2)$

Table 3: Exact lattice quantities for the corner rule along the diagonals. As the wave vector approaches the zone points M and R ($y \to 1$), the Ampère term tends to $\frac{1}{2}$ and $\frac{2}{3}$ of σ_{face}, and the magnetic variance to 2 and 3 times the electric one.

The shortest waves. A mode on a zone face has $\omega = \pi$: its amplitude flips sign at every tick, and by Corollary 9 it has no pure vacuum; no state at all is left unchanged by its tick (Theorem 8). Modes just inside the zone do have a vacuum, but a large one. To see how large, take a periodic lattice with an odd number N of voxels along each axis, so the wave numbers are $k_j = 2\pi m_j/N$ with integers $|m_j| \le (N-1)/2$. None of them equals π, so every mode has $P(k) = \prod_j \cos(k_j/2) > 0$, and every mode with $k \neq 0$ has a vacuum. The uniform modes, $k = 0$, have $\omega = 0$ and no vacuum (Theorem 8). Near the zone faces P is small. By Theorem 8 with $\cos(\omega/2) = P$, the electric variance of a mode is $\hbar\tan(\omega/2) = \hbar\sqrt{1-P^2}/P$, which lies between $\hbar(1/P - 1)$ and $\hbar/P$, and so grows without bound as $P \to 0$.

The electric field at a single voxel receives a contribution from every mode. Write the field as a Fourier sum over the N^3 wave vectors, with a factor $N^{-3/2}$ so that the transform is orthonormal. In a mode whose polarization vector ε has unit length, component j of the electric field carries the share $|\varepsilon_j|^2/N^3$ of the mode's variance, and the three components together carry $\sum_j |\varepsilon_j|^2/N^3 = 1/N^3$ of it. Each wave vector has two orthogonal unit polarizations, each an independent mode, so the electric variances of the three field components at one voxel add up to exactly $(2/N^3)\sum_{k\neq 0} \hbar \tan(\omega_k/2)$, a sum over the modes that have a vacuum. Adding the term $0 = \hbar \tan 0$ for $k = 0$ does not change it, so this is two times the average $\bar{V}$ of the per-mode variance over all N^3 wave vectors, and the bounds above give $\hbar(\bar{U} - 1) \leq \bar{V} \leq \hbar\,\bar{U}$, where $\bar{U}$ is the average of $1/P$. Because $1/P$ is a product of three one-dimensional factors $\sec(\pi m_j/N)$, its average is the cube of a one-dimensional average: $\bar{U} = S_N^3$ with $S_N = \frac{1}{N}\sum_{|m|\leq (N-1)/2} \sec(\pi m/N)$. Two elementary inequalities on $[0, \pi/2)$ bound the secant. One is $\sec x \geq \frac{2}{\pi - 2x}$, which is $\sin y \leq y$ at $y = \pi/2 - x$. The other is $\sec x \leq \frac{\pi}{\pi - 2x}$, which is Jordan's inequality $\sin y \geq 2y/\pi$ at the same y. At $x = \pi m/N$ the denominator is $\pi - 2x = \pi(N - 2m)/N$, and $N - 2m$ runs through the odd numbers $1, 3, 5, \ldots, N - 2$ as m runs from $(N - 1)/2$ down to 1. Write O for the sum $1 + \frac{1}{3} + \frac{1}{5} + \cdots + \frac{1}{N-2}$ of their reciprocals. Summing the two inequalities over m (the term $m = 0$, where $\sec 0 = 1$, is kept only in the upper bound) gives

$$\frac{4}{\pi}\,O \;\leq\; S_N \;\leq\; \frac{1}{N} + 2\,O.$$

Because $\frac{1}{2j} \leq \frac{1}{2j-1} \leq \frac{1}{j}$, O lies between half of the harmonic sum $H = 1 + \frac{1}{2} + \cdots + \frac{1}{(N-1)/2}$ and all of it, and H lies between $\ln \frac{N+1}{2}$ and $1 + \ln \frac{N-1}{2}$. Together,

$$\frac{2}{\pi} \ln \frac{N+1}{2} \;\leq\; S_N \;\leq\; 3 + 2 \ln N.$$

The average per-mode variance therefore lies between $\hbar\big[(\frac{2}{\pi} \ln \frac{N+1}{2})^3 - 1\big]$ and $\hbar\,(3 + 2 \ln N)^3$. It grows like $(\ln N)^3$, one logarithm for each axis, and without bound as the number of voxels increases at fixed voxel size. A field averaged with a smooth profile over a region of many voxels is different. The average weights each mode by the profile's Fourier transform, and a profile that varies slowly over many voxels contains almost no waves of one-voxel wavelength, so its weight is small near the zone boundary. The variance of the averaged field is then set by the long-wavelength modes, where the per-mode values above approach the continuum ones.

8 Which corrections cannot be removed

Some of the differences from the continuum belong to the corner rule: another stencil would move them around. Others come from putting a local rule on a grid at all, and do not depend on the stencil. This section collects five of the second kind. None is tied to the corner rule's weights, and each paragraph names the class of rules for which it is proved.

No rule reaches exactly the ball that light fills, tick after tick. In the continuum, a disturbance at the origin has reached, after n units of time, exactly the points within distance n: a ball of radius n. A local rule has its own region, the voxels that it reaches after n ticks in the sense of Section 3. Theorem 2 confines the region of a face rule to an octahedron and the region of a cube rule to a cube. Neither is a ball, and no rule does better: no local rule reaches exactly the voxels of the ball of radius 1 after one tick and exactly those of the ball of radius 2 after two. The voxels within distance one of the origin are the voxel itself and its six face neighbours, so a rule that reaches exactly those after one tick is a face rule. After two ticks Theorem 2 confines it

to the octahedron $|x_0| + |x_1| + |x_2| \le 2$, which leaves out the voxel $(1, 1, 1)$, although that voxel is only $\sqrt{3} < 2$ from the origin. So at one tick or at two, the rule's region and the ball differ. The argument does not use the kind of weights, which may be numbers, matrices or operators.

A face rule falls short of the ball, which is Corollary 3. A rule that reaches beyond the face neighbours sends some influence outside it; the argument takes real weights. Let s be a hop of its stencil of greatest Euclidean length that is not a face hop, so $|s| \ge \sqrt{2}$. Any other hop t of the stencil is no longer than s, so $t \cdot s \le |t|\,|s| \le s \cdot s$, with equality only when $t = s$. A chain of n hops that ends at the voxel $n\,s$ has total dot product $n\,s \cdot s$ with s, and each hop of the chain contributes at most $s \cdot s$, so each contributes exactly that and is the hop s. The voxel $n\,s$ is therefore reached with weight exactly w_s^n, which is not zero, at distance $n|s| \ge \sqrt{2}\,n$ from the origin, outside the ball of radius n that light fills. We call the field that a rule delivers outside the light ball a *precursor*. For a two-step rule the weights are the 2×2 matrices of Section 3, in which a hop of S with weight w appears as $\begin{pmatrix} 2w & 0 \\ 0 & 0 \end{pmatrix}$, and the same argument gives the weight $(2w_s)^n$ in the upper left entry. In the corner rule the longest hops are the eight corner hops, each of weight $w_s = \frac{1}{32}$, so the far corner (n, n, n) of the cube is reached with weight $(2 \cdot \frac{1}{32})^n = 16^{-n}$. The corner rule can therefore influence a voxel on a body diagonal after fewer ticks than light needs, up to $\sqrt{3}$ voxels per tick, but with an amplitude that falls as 16^{-n} at the far corner. A wave packet still moves at group speed at most one (Theorem 4). For the smooth pulse of Figure 8(b), the field at voxels more than four voxels outside the light sphere is at most 1.5% of the largest field value anywhere, and at voxels more than eight outside it the field is below 10^{-5} of that value (Section 5). A bound on the precursor of an arbitrary starting field is not derived here.

Frequencies stop at π. Section 2 took $0 \le \omega \le \pi$ because a tick records the phase of a wave only modulo 2π: the sequences $e^{-i\omega n}$ and $e^{-i(\omega+2\pi)n}$ are the same at integer n, and a real field cannot tell the wave (k, ω) from $(-k, -\omega)$. So no tick rule can show an oscillation faster than a sign flip at every tick, which is $\omega = \pi$. The continuum relation $\omega = |k|$ asks for more at every wave vector with $|k| > \pi$, and the zone contains such vectors, up to $|k| = \sqrt{3}\,\pi \approx 5.44$ at the corner R. So no tick rule, whatever its stencil, matches the continuum dispersion relation on the whole zone. Within $|k| \le \pi$ the cap does no harm, and the next correction acts there.

Exact speed along the axes forces speeds that depend on direction. Take a two-step rule $u(n + 2) = 2S\,u(n + 1) - u(n)$ in which S has a finite stencil with real weights, equal at s and at $-s$, and suppose that it moves light exactly one voxel per tick along the axes, so that $\cos\omega = \hat{S}(\kappa, 0, 0) = \cos\kappa$ for $|\kappa| \le \pi$. The corner rule is one such rule, and by Theorem 5 so is every rule of the family (9) that meets the two conditions of Section 4. Can the frequency also equal $|k|$ in every other direction? Take the face-diagonal waves $k = (t, t, 0)$, whose wave number is $|k| = \sqrt{2}\,t$, with $0 \le t \le \pi/\sqrt{2}$ so that $|k| \le \pi$. If the speed did not depend on direction, the symbol at $(t, t, 0)$ would equal the symbol at the axis point $(\sqrt{2}\,t, 0, 0)$ of the same wave number, which is $\cos(\sqrt{2}\,t)$. With weights equal at s and $-s$ the symbol is $\hat{S}(k) = \sum_s w_s \cos(k \cdot s)$, so on the face-diagonal slice it is $\varphi(t) = \sum_s w_s \cos\big((s_0 + s_1)\,t\big)$, a finite sum of cosines of integer multiples of t. Both $\varphi(t)$ and $\cos(\sqrt{2}\,t)$ equal their power series for every t. Their difference is then a function given by a convergent power series, and if it vanishes on an interval, every derivative of it vanishes at a point of the interval, all its power-series coefficients about that point are zero, and it vanishes for every t. So $\varphi(t) = \cos(\sqrt{2}\,t)$ for all t. But $\varphi(t + 2\pi) = \varphi(t)$, because each cosine has integer frequency, whereas $\cos(\sqrt{2}\,(t + 2\pi)) = \cos(\sqrt{2}\,t)$ for all t would need $\cos(2\sqrt{2}\,\pi) = 1$, that is, $\sqrt{2}$ equal to an integer. The face-diagonal frequency therefore differs from the axis frequency at some wave

number up to π. How the difference is distributed depends on the stencil. At small wave number the corner rule is exact along the axes and lags by $|k|^2/72$ along the body diagonal, while the face leapfrog at $r = 1/\sqrt{3}$ is exact along the body diagonal and lags most along an axis (Section 4). At a wavelength of 20 voxels the corner rule's body-diagonal wave is 0.14% slower than light, and the axis wave is not slower at all (Table 2).

Modes with $\omega = \pi$ have no vacuum. A rule that is exact along the axes has $\cos\omega = \cos\pi = -1$ at the wave vector $(\pi, 0, 0)$, the point X of Figure 3, so $\omega = \pi$ there. Section 7 defines a mode's vacuum as the pure Gaussian state that one tick leaves unchanged. For $\omega = \pi$, that is $q = 4\sin^2(\omega/2) = 4$, Theorem 8 shows that no Gaussian state at all, pure or mixed, is left unchanged, and that one tick shears phase space, so that the amplitude grows in proportion to the number of ticks. The argument uses only $q = 4$, so it covers every rule that is exact along an axis. A periodic lattice with an even number of voxels along an axis contains such wave vectors, with $k_j = \pi$, and has modes without a vacuum. With an odd number along every axis it contains none, and every mode with $k \neq 0$ has a vacuum (Section 7). The uniform modes are the exception: they have $\omega = 0$, that is $q = 0$, and Theorem 8 shows that one tick leaves no Gaussian state of them unchanged either.

Ampère's law carries a lattice term. Theorem 6 shows that Gauss's law, the absence of magnetic charge and Faraday's law are exact for every local rule. Ampère's law, $E(n+1) - E(n) = \nabla^- \times B(n+1)$ with the coefficient one that light of one voxel per tick requires, is exact only if the operator $\mathcal{D} = 2(1 - S) + \Delta$ of that theorem is zero on every field, and this happens only for $S = 1 + \Delta/2$, the face leapfrog (4) with $r = 1$. Its symbol is $1 - \frac{1}{2}\sigma_{\text{face}}$, which equals -5 at the zone corner R, where $\sigma_{\text{face}} = 12$, and a real frequency needs a symbol between -1 and 1. So that rule is unstable, and every stable rule of the form (13), which applies one rule S to each component of A, leaves a nonzero lattice term in that law. A rule whose light is slower has an exact law of its own: for the face leapfrog at any r, $\mathcal{D} = (1 - r^2)\Delta$ and $E(n+1) - E(n) = r^2\,\nabla^- \times B(n+1)$ (Section 6). For the corner rule the term is $d(k)$ of Theorem 7, of fourth order in the wave number and zero on the axes.

9 Relation to earlier work

The face-rule bound is the Courant–Friedrichs–Lewy condition [1], and Yee's staggered scheme [2], the standard method for simulating Maxwell's equations on a cubic grid, is its best-known instance in electromagnetism. Yee's scheme also keeps the charge-free Gauss law and $\nabla \cdot B = 0$ exactly at every step. Theorem 6 extends that property to every local tick rule, including those that move light at one voxel per tick.

Stencils that reach beyond the faces have been studied in computational electromagnetics and seismology as a way to relax the time-step limit. Cole's nonstandard finite differences allow a time step about 36% longer in three dimensions and reduce the dispersion error [3, 4]. Forgy and Chew obtain isotropic dispersion and a larger stable step on an overlapped lattice [5]. Saenger, Gold and Shapiro's rotated staggered grid differentiates along cell diagonals and is stable for time steps about $\sqrt{3}$ times longer than the standard grid in three dimensions [8]. O'Reilly and Beck give a 27-point Laplacian L whose face, edge and corner weights are $\frac{3}{13}$, $\frac{3}{26}$ and $\frac{1}{13}$, in the ratio $1 : \frac{1}{2} : \frac{1}{3}$, and which is stable at Courant number one, that is, one voxel per tick, for the wave equation [6]. Used in the leapfrog $u(n+2) = 2u(n+1) - u(n) + L\,u(n+1)$, it is the two-step rule with $S = 1 + \frac{1}{2}L$, which puts the weight $-\frac{9}{13}$ on the voxel and $\frac{3}{26}$, $\frac{3}{52}$ and $\frac{1}{26}$ on each face, edge and corner neighbour.

With $a_1 = 2w_1$, $a_2 = 4w_2$ and $a_3 = 8w_3$, that is $(a_0, a_1, a_2, a_3) = (-\frac{9}{13}, \frac{3}{13}, \frac{3}{13}, \frac{4}{13})$ in the family (9). Both conditions of Section 4 hold, and $a_2 + a_3 = \frac{7}{13}$ lies above the bound $\frac{1}{2}$ of Theorem 5. These schemes are engineered for accuracy at a given cost. The corner rule is another member of the same family, chosen because its frequency, its group velocity and its mode vacuum all have closed forms (Theorems 4 and 8); it is used here to ask what a rule of this kind has to give up (Section 8).

On the quantum side, a quantum cellular automaton is a quantum evolution built from identical local unitary steps on a lattice, the quantum counterpart of the local rules here. Bialynicki-Birula showed that the Weyl, Dirac and Maxwell equations can be realized as such automata [9]. D'Ariano and Perinotti found that among the cubic lattices only the body-centred cubic one admits a non-trivial isotropic Weyl automaton, an automaton whose long-wavelength limit is the Weyl equation for a massless spin-$\frac{1}{2}$ particle [7], and Bisio, D'Ariano and Perinotti built the free electromagnetic field from pairs of such automata [10]. Their photon has a dispersive speed and a small longitudinal component. The body-centred cubic lattice joins each site to its eight nearest neighbours, the corners of the cube centred on it. The corner rule uses those eight hops together with the faces and edges, and Section 10 shows that its corner hops alone connect the sites of one such lattice. The potential formulation used here keeps the field exactly transverse, and the vacuum of Section 7 is computed in closed form for each mode.

10 Consequences

For a theory in which space is built from voxels and time from ticks, with light crossing one voxel per tick, the result is a constraint on the structure of a tick. It must reach beyond the six face neighbours (Corollary 3). Within the family of Section 4, a stable rule whose shortest wave, at the zone corner R, has a nonzero frequency must give weight to the corner neighbours (Theorem 5). Corner hops have a name in crystallography. Start at a voxel and move by corner hops only. Each hop changes every coordinate by one, so after any number of hops the three coordinates, counted from the starting voxel, are all even or all odd. The sites reached form a body-centred cubic lattice: the cubic lattice of side two with an extra point at the centre of every cube. The nearest neighbours of any of its sites are exactly the eight corner neighbours.

A tick that reaches the corners can carry light at one voxel per tick. With the corner rule on the cubic grid, Gauss's law, the absence of magnetic charge and Faraday's law are exact, which holds for every local rule, and Ampère's law holds up to a fourth-order lattice term that vanishes along the axes. Every mode with a frequency strictly between 0 and π has a vacuum with a closed form; the modes at the two ends, the uniform modes at $\omega = 0$ and the modes at $\omega = \pi$, have none. Table 4 sets the two ways to tick side by side.

	face leapfrog, $r = 1/\sqrt{3}$	corner rule
neighbours one tick reaches	6	26
region reached after n ticks	octahedron $\sum_j \lvert x_j\rvert \le n$	cube $\max_j \lvert x_j\rvert \le n$
light speed, voxels per tick	$1/\sqrt{3}$	1
direction without dispersion	body diagonal	axes
long-wave lag along an axis	$\lvert k\rvert^2/36$	0
long-wave lag along the body diagonal	0	$\lvert k\rvert^2/72$
edge and corner weight, $a_2 + a_3$	0	$\frac{1}{2}$, the least for a stable rule of the family (9) at unit speed
Gauss's law	exact	exact

Table 4: The two ways to tick. A lag is the relative shortfall of the phase speed below the rule's own light speed, as in (11).

Which neighbours a single tick connects is the decisive structural choice, and it leaves a signature inside the model. The corner rule is exact along its axes and lags most along the body diagonal. The face leapfrog at its largest stable step is exact along the body diagonal and lags most along an axis. At a wavelength of 20 voxels the lag is at most 0.27% for either rule (Table 2), and at long wavelength it grows as the square of the wave number.

References

[1] R. Courant, K. Friedrichs and H. Lewy, Über die partiellen Differenzengleichungen der mathematischen Physik, *Math. Ann.* **100**, 32–74 (1928).

[2] K. S. Yee, Numerical solution of initial boundary value problems involving Maxwell's equations in isotropic media, *IEEE Trans. Antennas Propag.* **14**, 302–307 (1966).

[3] J. B. Cole, A high-accuracy realization of the Yee algorithm using non-standard finite differences, *IEEE Trans. Microwave Theory Tech.* **45**, 991–996 (1997).

[4] J. B. Cole, High-accuracy Yee algorithm based on nonstandard finite differences: new developments and verifications, *IEEE Trans. Antennas Propag.* **50**, 1185–1191 (2002).

[5] E. A. Forgy and W. C. Chew, A time-domain method with isotropic dispersion and increased stability on an overlapped lattice, *IEEE Trans. Antennas Propag.* **50**, 983–996 (2002).

[6] R. C. O'Reilly and J. M. Beck, A family of large-stencil discrete Laplacian approximations in three dimensions, preprint, University of Colorado Boulder (2006).

[7] G. M. D'Ariano and P. Perinotti, Derivation of the Dirac equation from principles of information processing, *Phys. Rev. A* **90**, 062106 (2014).

[8] E. H. Saenger, N. Gold and S. A. Shapiro, Modeling the propagation of elastic waves using a modified finite-difference grid, *Wave Motion* **31**, 77–92 (2000).

[9] I. Bialynicki-Birula, Weyl, Dirac, and Maxwell equations on a lattice as unitary cellular automata, *Phys. Rev. D* **49**, 6920–6927 (1994).

[10] A. Bisio, G. M. D'Ariano and P. Perinotti, Quantum cellular automaton theory of light, *Ann. Phys.* **368**, 177–190 (2016).