

Every Crystal Has Its Own Speed of Light

Brecht Corbeel

Absolute Digital Publishers

absolutedigitalpublishers.com

September 20, 2026

Abstract. A Lieb–Robinson bound gives every lattice system an emergent causal structure: a material constant c_{LR} outside whose light-cone all influence decays exponentially, now visible directly as the measured cone slopes of quantum-gas and trapped-ion quenches. Taking that emergent-relativity reading seriously but honestly yields three products. First, a census of c_{LR} across nine measured platforms spanning nine and a half orders of magnitude on one axis, from the ≈ 1.1 mm/s doublon–holon cone of a Bose–Hubbard gas to the $\approx 3 \times 10^6$ m/s Dirac cone of graphene — every number from a verified source, and the honest span, not the round eleven orders one is tempted to claim. Second, a transfer theorem that classifies exactly which special-relativistic kinematic facts replay inside a crystal: front-speed causality and effective causal order transfer exactly (with Klein tunneling, already measured in graphene, the flagship of massless-cone kinematics); internal-clock time dilation $\gamma_{\text{eff}} = (1 - v^2/c_{\text{eff}}^2)^{-1/2}$ transfers with a computed band-curvature correction $\propto (k/\Lambda_{\text{curv}})^2$; and boost invariance fails provably, because the lattice is an ether that selects a frame. Third, the in-crystal Michelson–Morley: the ether drift is real and, in a flowing medium, first-order and large, $\delta c/c \simeq v_d/c_{\text{eff}}$ — the inverse of the historical null — yet an interferometer built entirely of the emergent quasiparticles nulls anyway, its arms contracting by the same acoustic γ_{eff} , so the lattice frame is detectable only with an external clock and ruler. We confront one published dataset — the trapped-ion α -scan cone exponents $\zeta(\alpha)$ of Richerme et al. (2014) — against the long-range Lieb–Robinson thresholds: the measured cones are super-linear as the theorems require and violate the nearest-neighbour bound, but the data lie below every threshold (so they cannot discriminate between them) and one point, XY at $\alpha = 1.19$, is explained by no known theory. Finally, the crystal is the one place we know relativity is emergent, and its window quality ($\sim 10^{-2}$) set next to the vacuum’s Lorentz-violation bounds ($\sim 10^{-16}$) makes the gap itself the argument: our vacuum’s window is at least fourteen orders better than any crystal’s, a fact naive emergent-spacetime pictures must explain rather than assume.

Keywords: Lieb–Robinson bounds; emergent causality; light-cone spreading; emergent Lorentz invariance; Klein tunneling; Dirac materials; Lorentz-violation bounds.

Series: The Crossed Fields Papers, No. 5.

1 A 1972 speed limit is now a number you can read off a graph

Lieb and Robinson proved in 1972 that local dynamics on a lattice carry a finite group velocity: outside a cone whose slope is set by that velocity, the commutator of two observables — and with it any physical influence one can exert on the other — decays exponentially in distance [41]. Read as a bookkeeping bound this is a locality theorem for quantum spin systems. Read as physics it says something sharper: every lattice system has its own speed of light. There is an emergent material constant, which we write c_{LR} , that plays inside the crystal the causal role the vacuum speed c plays outside it — different for a cold-atom gas than for graphene, but in each case a hard front past which nothing propagates. The thesis of this paper is that the second reading is no

longer a metaphor. The cone is a slope in published data; c_{LR} is a measured quantity; and taking the emergent-relativity reading seriously, but honestly, yields three concrete products rather than one loose analogy.

For roughly three decades the theorem was almost unused outside mathematical physics [10]. The quantum-information revival around the turn of the century turned it into a workhorse: a nonvanishing gap forces exponential clustering of ground-state correlations [28, 46], the bound limits how fast entanglement and topological order can be generated [8], and a self-contained modern re-derivation now ties the short-range cone, its long-range generalizations, and the bosonic subtleties into one framework [14]. Then the bound became an experimental fact. Cheneau et al. watched a correlation cone open at a few millimetres per second in a Bose–Hubbard gas [11]; Jurcevic et al. and Richerme et al. mapped cones in trapped-ion chains, including the long-range interactions that bend the cone and, below a threshold, break

the linear picture altogether [31, 51]. A theorem about the abstract propagation of operators had become a slope one fits to a set of points.

We are careful about what in this paper is inherited and what is ours. The bound is Lieb and Robinson’s; its modern tightenings and information-theoretic meaning belong to Hastings, to Nachtergaele and Sims, and to the quantum-information community [14, 28, 46, 65]. The measured cones are Cheneau’s, Jurcevic’s, and Richerme’s [11, 31, 51]. That a linear low-energy dispersion is an emergent Lorentz window is the commonplace of the Dirac-materials community — the Fermi velocity of graphene, $v_F \approx c/300$, is the standard example [9, 48] — and Volovik’s program pushes the same reading to full emergent gravity, which we cite but do not endorse [63]. Klein tunneling is Katsnelson, Novoselov, and Geim’s prediction and Young and Kim’s measurement [32, 67]. The “emergent ether” framing has its own literature: the acoustic Lorentz–FitzGerald contraction is exact in Barceló and Jannes [3], and Liberati’s Lorentz-violation reviews supply the anisotropy formalism our in-crystal Michelson–Morley section borrows [40].

The honesty is not decoration; it changes three of the four results. The census is nine and a half orders, not eleven, because the measured extremes do not reach the round number. The confrontation is a threefold verdict — a qualitative win, an inability to discriminate between competing theories, and a point none of them explains — not a headline agreement. And the ether experiment is an inversion: we run Michelson–Morley in a medium we already know has a rest frame, so the live question is not whether it registers but what a registration means. Stated carelessly each of these becomes an overclaim; stated carefully each is a result, and the grading in each sentence is what keeps the difference visible.

What we contribute is four things, and we state their register plainly. First, a census of c_{LR} across nine measured platforms under one extraction convention — the cone-front group velocity at the relevant filling or state — ours as a compilation. It spans nine and a half orders of magnitude, from the ≈ 1.1 mm/s doublon–holon cone of a Bose–Hubbard gas to the $\approx 3 \times 10^6$ m/s Dirac cone of graphene. We flag at once that this is nine and a half orders, not the round eleven the one-sentence thesis first reached for: the verified numbers do not support eleven, and we quote what they give (Section 3). Second, a transfer theorem, ours as packaging with proofs at a stated order: which special-relativistic kinematic facts replay inside a crystal — front-speed causality and effective causal order exactly, with Klein tunneling the measured flagship of massless-cone kinematics; internal-clock time dilation $\gamma_{\text{eff}} = (1 - v^2/c_{\text{eff}}^2)^{-1/2}$ with a computed band-curvature correction $\propto (k/\Lambda_{\text{curv}})^2$ — and which fail provably, boost invariance first among them, because the lattice is an ether that selects a frame (Section 4). Third, the in-crystal Michelson–Morley: its design and its anisotropy sensitivity are ours, together

with the honest inversion that an interferometer built entirely of the emergent quasiparticles nulls even though the drift is real and first-order (Section 5). Fourth, a confrontation, ours as analysis of published data: we set Richerme’s α -scan cone exponents against the long-range Lieb–Robinson thresholds and report agreement and its limits as found, including one point no theory explains (Section 5).

Two commitments govern the argument. Every analogical step between crystal kinematics and special relativity is graded in the sentence where it is made — exact where it follows from locality alone, structured where it holds only inside the emergent window and carries a correction, speculative where it reaches toward fundamental physics. And every number is attributed at the point of use, with the paper’s own seeded computation reported in Appendix C. Section 2 states the bound and narrates the measured cones; the rest is the census, the transfer theorem, the ether experiment, and the confrontation, in that order.

2 The bound is a theorem about operators; the cone is a slope in data

The founding statement is unambiguous. For a finite-range interaction on a lattice the time-evolved observables cannot correlate faster than a fixed velocity, and the failure outside the cone is exponential.

Theorem 2.1 (Lieb–Robinson, 1972). *For a finite-range interaction Φ on a ν -dimensional lattice there exist a finite group velocity V_Φ and a strictly positive increasing function $\mu(v)$ such that*

$$\lim_{|t| \rightarrow \infty, |x| > v|t|} e^{\mu(v)|t|} \|[\tau_t^\Phi \tau_x(A), B]\| = 0$$

for every $v > V_\Phi$ and all strictly local observables A, B ; the asymptotic cone is $C_\Phi = \{(x, t) : |x| > V_\Phi|t|\}$ [41].

The theorem admits a two-body extension to interactions decaying as $e^{-\alpha|x|}$, which still gives an exponentially clean cone; a slower-than-exponential decay gives a commutator bound whose tail is linked to that of the interaction rather than a hard front [41]. Three consequences make the bound physical rather than formal. It generates an effective light cone with exponentially decaying tails for any local Hamiltonian, so information that leaks past the front is exponentially small [8]. A nonvanishing spectral gap turns the bound into exponential clustering of ground-state correlations, the standard companion result [28, 46]. And the velocity that appears is an upper limit set by the couplings, not a measured speed: the tightest short-range bounds make it S -independent for the spin- S Heisenberg model and scale it as $\sqrt{N}$ rather than N for the N -orbital Hubbard model [65], so a census of *measured* cone speeds is

reporting a different object from these worst-case constants, and we keep the two apart throughout.

The bound’s reach is wider than correlation spreading. It caps the rate at which quantum information scrambles: for one-dimensional power-law interactions the operator-growth speed between two isolated qubits stays finite for sufficiently large α , so even the butterfly effect has a front [13]. It bounds the simulatability of dynamics and the growth of entanglement, and it fixes a minimum time for state transfer across a chain — the same inequality read as a resource limit rather than a locality guarantee [14, 59]. Improved power-law bounds sharpen the constant without changing the structure [19], and the pedagogical re-derivations make explicit that all of it descends from the one 1972 estimate [27, 47]. The early route to bosonic systems recovered a meaningful analogue of the bound by restricting attention to quasi-local, finite-energy observables [54], a foreshadowing of the sharper bosonic results the census must respect.

One class of platforms carries a load-bearing caveat before any number is quoted. The bosonic local Hilbert space is unbounded, and a naive Lieb–Robinson bound can fail because site occupations are not capped. For Bose–Hubbard dynamics the rigorous guarantee is weaker than the 1972 finite-dimensional theorem: under the physically standard restriction to states of bounded per-site occupation, the wavefront grows at most as $t \log^2 t$ — almost linear, only logarithmically worse than a clean cone [35]. Independent routes confirm a finite ballistic bound for macroscopic transport in the same model [20] and a bound that scales linearly in the number density rather than staying a filling-independent constant [66]. The Bose–Hubbard census entry is therefore a measured quasiparticle velocity sitting on top of an almost-linear, state-restricted theoretical floor, and we say so where it appears.

2.1 What the cold-atom and ion-chain quenches actually measured

Cheneau et al. prepared a one-dimensional Bose–Hubbard chain of ^{87}Rb in an optical lattice of period $a_{\text{lat}} = 532$ nm, quenched from a deep Mott insulator, and watched a correlation cone open [11]. The measured propagation velocity is reported dimensionless: $v \hbar / (J a_{\text{lat}}) = 5.0(2)$, $5.6(5)$, and $5.0(2)$ at final ratios $U/J = 5.0$, 7.0 , and 9.0 , with $J/\hbar = 691(23)$, $522(14)$, and $422(11)$ s^{-1} respectively [11]. Converted through $J/\hbar$ and a_{lat} these are ≈ 1.84 , 1.56 , and 1.12 mm/s: the cone is millimetres per second, not the “ ~ 6 mm/s” a dimensionless reading invites. The authors interpret their analytic maximum $v_{\text{max}} \approx (6J a_{\text{lat}}/\hbar)[1 - 16J^2/(9U^2)]$ as an effective Lieb–Robinson bound for this model, asymptoting to a doublon group velocity $4J a_{\text{lat}}/\hbar$ faster than the holon’s $2J a_{\text{lat}}/\hbar$ through Bose-enhanced tunneling [11] — so the “cone speed” here is specifically a doublon–holon pair velocity at a chosen quench, which is exactly

the extraction convention the census fixes.

Trapped-ion chains put the same physics at metres per second and, crucially, let the interaction range tune. Jurcevic et al. encoded a long-range Ising model on 7–15 $^{40}\text{Ca}^+$ spins with couplings $J_{ij} \sim 1/|i-j|^\alpha$, α adjustable from 0 (infinite range) to ≈ 3 (short range), and extracted the maximum magnon group velocity from the fitted dispersion; for the shortest range it agrees with the nearest-neighbour model, and as α decreases it is consistent with divergence — no finite maximum group velocity, the light-cone picture breaking down [31]. That $v_g^{\text{max}}(\alpha)$ is read from a figure rather than printed as a number, a usability point we honour by treating it as figure-extracted.

The cone has since been seen well beyond these first two platforms. A quenched one-dimensional Bose gas shows thermal correlations emerging locally and then spreading in a light-cone pattern, rather than the whole system relaxing at once [38]; the trapped-ion toolchain has scaled to 53 spins with single-shot readout [68]; and a nine-transmon superconducting lattice has mapped an information cone through out-of-time-ordered correlators, on a Manhattan-distance metric because its qubits are wired rather than spatially extended [7]. Each confirms the front as a robust, platform-independent feature, and each shows that “the cone speed” is only as sharp as the extraction convention — which is exactly what the census fixes next.

2.2 Below a threshold in the interaction range, the cone bends and then breaks

Richerme et al. made the breakdown quantitative on an 11-ion $^{171}\text{Yb}^+$ chain, fitting the light-cone boundary as $r \sim t^\zeta$ at fixed correlation contour for four interaction ranges $\alpha \in \{0.63, 0.83, 1.00, 1.19\}$ in both an Ising and an XY model, and comparing against the nearest-neighbour bound $v_{\text{LR}} = 12eJ_{\text{max}}$ [51]. The propagation exceeds that bound outright for $\alpha < 1$: the cone is not merely bent but supersonic relative to any nearest-neighbour sound speed. This is the dataset Section 5 confronts, and its exponents are stated there; here it fixes the qualitative fact that a long-range crystal need not have a finite c_{LR} at all.

Physically, cone-bending is the long tail of $1/r^\alpha$ reaching across the chain faster than any local process could carry a signal: below threshold, distant sites correlate before a nearest-neighbour front could arrive, so the boundary of appreciable correlation is not a straight line $r \propto t$ but a curve $r \sim t^\zeta$ with $\zeta > 1$. The exponent, not a velocity, is then the observable — which is why the confrontation of Section 5 is framed in exponents rather than in speeds.

The theory of where the cone survives is a hierarchy, not a single threshold, and its conventions must not be conflated. In one dimension, Tran et al. establish three

linear light cones at distinct thresholds: a genuine Lieb–Robinson cone for arbitrary operators and all states at $\alpha > 3$, a Frobenius cone for typical states at $\alpha > 5/2$, and a free cone governing single-particle propagation at $\alpha > 2$ [60]. Kuwahara and Saito prove a strictly linear front for $\alpha > 2D + 1$ (that is, $\alpha > 3$ in one dimension) and construct an explicit state-transfer protocol that violates linearity below it, making the threshold tight from both sides [34]. Foss-Feig et al. bracket the intermediate regime: for $\alpha > 2D$ the cone is algebraic, becoming linear only as $\alpha \rightarrow \infty$, while for $\alpha > D$ the influence time grows at least as $(\alpha/v) \log r$, ruling out exponential-in-time growth without yet delivering a hard cone [22]. And the threshold cuts the other way at small α : Eisert et al. construct Hamiltonians with $\alpha < D$ in which quasi-locality itself breaks down and correlations spread with no causal region at all [17], a counterpoint that the algebraic-cone theorems do not contradict because they operate in the α -large regime where a supersonic but bounded region still exists [25]. Richerme’s four ranges all sit below $\alpha = 2$, in the regime where every one of these results agrees that no finite-velocity linear cone is guaranteed — which is why the confrontation can be decisive against a naive bound and, as we will see, indecisive between the competing long-range theories.

3 Nine platforms, one convention, nine and a half orders of magnitude

The bound and the measurements combine into a single compilation. We state it as a proposition because both halves are precise: the inequality is the standard Lieb–Robinson form, and the span is a computed number, not a rhetorical one.

Proposition 1 (The bound and the census). *For local lattice dynamics there are constants $C, \xi > 0$ and a velocity c_{LR} with*

$$\| [A_x(t), B_y] \| \leq C e^{-(d(x,y) - c_{\text{LR}}t)/\xi},$$

so that outside the cone $d(x,y) > c_{\text{LR}}t$ all influence decays exponentially; c_{LR} is bounded above by the microscopic couplings and read off empirically as a measured cone slope. Under one extraction convention — the cone-front group velocity at the relevant filling or state — the empirical c_{LR} of real measured systems spans nine and a half orders of magnitude, from the ≈ 1.1 mm/s Bose–Hubbard doublon–holon cone [11] to the $\approx 3 \times 10^6$ m/s graphene Dirac cone [18, 48], every value from a verified source.

Table 1 is the compilation and Figure 1 plots it on a single logarithmic axis. The ratio of the two extremes is $10^{9.44}$ — 3×10^6 m/s over 1.1×10^{-3} m/s — so the honest span is nine and a half orders, or nearly ten (Appendix C).

Remark 1. The one-sentence thesis of this series first reached for eleven orders. The verified numbers do not reach it. Eleven was the round figure a quick estimate suggests when the top of the range is pushed to $\sim 10^7$ m/s and the bottom to $\sim 10^{-4}$ m/s; the measured doublon–holon cone is 1.1 mm/s, not a tenth of that, and the graphene value is a few $\times 10^6$ m/s, not 10^7 . We quote nine and a half and let the round number go.

The rows fall into two families with a gap between them. The engineered quantum simulators — Bose–Hubbard gases, trapped-ion and Rydberg spin chains — live at 10^{-3} to 3×10^{-2} m/s, because their emergent dynamics run on slow tunneling and spin-exchange rates times micron-scale lattice constants [11, 31, 42, 51]. The natural condensed-matter carriers — magnons, phonons, spinons, and Dirac and Weyl electrons — live at 10^2 to 3×10^6 m/s, set by exchange energies and band widths times sub-nanometre spacings [16, 36, 43, 48, 55]. Between $\sim 3 \times 10^{-2}$ and $\sim 2 \times 10^2$ m/s no platform in the census sits — a genuine empty band, not an artefact of sampling: it is the seam between systems slow enough to watch atom-by-atom under a microscope and systems fast enough to require spectroscopy, and nothing yet measured is emergent-relativistic in that middle decade. Nothing forbids it in principle — a heavily loaded magnonic or phononic device, or a simulator driven faster than present tunneling rates allow, could land there — so the gap is a statement about what has been measured, not about what is possible, and it is the census’s most obvious target for a future entry.

Every entry is one number extracted one way, and the way matters as much as the number. For the Bose–Hubbard row we quote the doublon–holon pair velocity at the stated quench, 1.1–1.8 mm/s across $U/J = 9$ to 5 [11]; for the trapped-ion row the magnon group velocity from the fitted dispersion, which diverges as the interaction range grows [31, 51]. Two rows are derived conversions rather than printed velocities and are labelled as such. The spinon value follows from the measured exchange $J = 33.5$ meV through the lower two-spinon boundary $\omega_l(k) = (\pi/2)J|\sin k|$, whose small- k slope $(\pi/2)J$ is the spinon velocity; multiplied by the c -axis spacing this is $\approx 3 \times 10^4$ m/s, a verifier’s conversion, not a number in the source [36]. The graphene row is deliberately a range, not a constant: the Fermi velocity that sets the Dirac cone runs from the conventional $\approx 1 \times 10^6$ m/s at ordinary density up to $\approx 3 \times 10^6$ m/s near the neutrality point, where weakly screened electron–electron interactions renormalize it [18, 48]. The Weyl/Dirac row is carried from the sibling paper’s own pool rather than re-derived here [16]. The superconducting-qubit row is flagged in the table and the figure: it is a wired circuit whose cone is a dimensionless slope in Manhattan distance, converted to a velocity only by assuming a chip pitch [7], and it is included to show the platform’s place in the ordering, not as a crystal speed on the same footing as the others. The full conventions and caveats —

Table 1: The census of emergent cone speeds c_{LR} under one extraction convention: the cone-front group velocity at the relevant filling or state. Ranges are where a platform tunes (dispersion regime, filling, or interaction range); single values are quoted points. Per-row provenance and caveats are in Appendix B; the span is computed in Appendix C.

Platform	System (convention)	Cone speed [m/s]	How extracted	Source
Bose-Hubbard cold atoms	^{87}Rb 1D chain, quench (doublon-holon pair)	1.1×10^{-3} – 1.8×10^{-3}	dimensionless $v\hbar/(Ja_{\text{lat}}) = 5.0$ – 5.6 , with $J/\hbar$ and $a_{\text{lat}} = 532$ nm	[11]
Trapped-ion spin chains	$^{171}\text{Yb}^+ / ^{40}\text{Ca}^+$, long-range Ising/ XY	2×10^{-3} – 1×10^{-2}	magnon group velocity from the fitted dispersion; super-linear (divergent) as $\alpha \rightarrow 0$	[31, 51]
Rydberg arrays	^{87}Rb tweezers, Ising anti-ferromagnet	3×10^{-3} – 3×10^{-2}	correlation-growth delay, short-time expansion	[5, 42]
Magnons (YIG)	yttrium iron garnet films and waveguides	2×10^2 – 2×10^4	spin-wave group velocity, dipolar to exchange regime	[6, 55]
Acoustic phonons (Si)	crystalline silicon, longitudinal	8.4×10^3	standard longitudinal sound speed	[43]
Superconducting qubits [†]	3×3 transmon hard-core Bose-Hubbard lattice	5×10^3 – 1.5×10^4	$v_l = 8.433$ km/s dimensionless OTOC cone, $J/2\pi = 8.1$ MHz converted by chip pitch	[7]
Spinons (KCuF_3)	1D spin- $\frac{1}{2}$ Heisenberg chain	3×10^4	$v = (\pi/2)J$, $J = 33.5$ meV, times c -axis spacing (derived conversion)	[36]
Weyl/Dirac semimetals	TaAs-class nodal fermions	1×10^5 – 1×10^6	node Fermi velocity (sibling reference pool)	[16]
Graphene electrons	monolayer Dirac carriers	1×10^6 – 3×10^6	Dirac v_F , running with density toward the node	[18, 48]

[†]A wired circuit, not a spatially extended crystal: its light cone is a dimensionless Manhattan-distance-versus-time slope, and the metres-per-second figure is that slope converted with the chip pitch, not a directly comparable material velocity [7].

including the bosonic almost-linear floor under the cold-atom row, the regime dependence of the YIG number, and the state dependence of ion-chain cones under confinement [57] — are collected in Appendix B.

As a contribution the census is deliberately modest in kind and specific in content: it is ours as a compilation, not as a measurement, and its only methodological claim is that one convention, applied to nine independently published systems, orders them on a single axis without special pleading. The ordering itself carries information. That the engineered simulators sit nine orders below the natural carriers is not an accident of units but a consequence of energy scales — nano-electronvolt tunnelings against electronvolt band widths — and that both families obey the same Lieb–Robinson inequality is why a single axis is meaningful at all. The span is the headline; the convention is what makes the span defensible.

4 Which of relativity’s kinematic facts a crystal is allowed to keep

A census of speeds invites a category error: if a crystal has a speed of light, does it have the rest of special relativity? The answer is not uniform, and the point of the transfer theorem is that it is not. Some facts follow from locality alone and transfer exactly; some hold only inside the emergent-Lorentz window and carry a correction we can compute; some fail as a matter of proof. We state the classification and then argue each register at its own order, with the proofs and counterexamples in Appendix A and the three columns drawn as Figure 2.

Proposition 2 (The transfer theorem). *Let a crystal carry a Lieb–Robinson velocity c_{LR} and, in a low-energy window, a linear dispersion $\varepsilon(k) = \hbar c_{\text{eff}} k$. The special-relativistic kinematic facts then partition into three registers. (i) Transfers exactly: front-speed causality (the c_{LR} cone), effective causal order for operations spacelike-separated in the c_{LR} metric, and massless-cone kinematics wherever the dispersion is linear. (ii) Transfers with a stated correction: the internal clock of a quasiparticle wavepacket at group velocity $v < c_{\text{eff}}$ runs slow by $\gamma_{\text{eff}} = (1 - v^2/c_{\text{eff}}^2)^{-1/2}$, with a leading correction $\propto (k/\Lambda_{\text{curv}})^2$ from band curvature. (iii) Fails provably: boost invariance, relativity of simultaneity beyond the low-energy window, and universality of the energy–momentum relation across bands.*

4.1 What transfers exactly, because it is only locality

The first register is exact in the strong sense: it needs no emergent Lorentz symmetry at all, only the Lieb–Robinson structure of Theorem 2.1. Front-speed causality is the bound itself — c_{LR} is a hard limit and influence outside the cone is exponentially small [8, 41]. Effective causal order follows immediately: two operations whose supports are spacelike-separated in the c_{LR} metric commute up to the exponential tail, so a well-defined partial order of “what can affect what” exists inside the crystal exactly as a light-cone order exists in Minkowski space [8]. These are graded exact because they are theorems about commutators, not analogies; the crystal keeps them for the same reason the vacuum does, that signalling is capped by a finite speed. Massless-cone




Figure 1: Every lattice has its own speed of light (Proposition 1). The emergent cone speed c_{cryst} — the Lieb–Robinson velocity, read off as a measured light-cone slope — across nine measured platforms, spanning nine and a half orders of magnitude on one axis, from the mm/s doublon–holon cone of a Bose–Hubbard gas [11] to the $\sim 10^6$ m/s Dirac cone of graphene [18, 48]. Segments show the measured range where a platform tunes (dispersion regime, filling, or interaction range); points are single quoted values. The engineered quantum simulators (mm/s–cm/s) and the natural condensed-matter carriers (10^2 – 10^6 m/s) occupy two bands separated by a gap no present platform fills. [†]The superconducting-qubit entry is a wired circuit; its cone is dimensionless and is converted here with the chip pitch, not a crystal velocity. Every number is from a verified source; values compiled in the seeded computation of Appendix C.

kinematics is the third exact entry, but with a scope condition: wherever the dispersion is genuinely linear, $\varepsilon(k) = \hbar c_{\text{eff}} k$, the low-energy excitations obey the massless Dirac equation with c_{eff} in place of c , and every kinematic consequence of that equation replays. The scope condition is the whole content — it holds in the window and nowhere else — but inside it the transfer is not approximate. One reads the emergent metric off directly: the cone slope is c_{eff} , and the causal order it induces is the same partial order a Minkowski cone induces, so the geometry is measured rather than borrowed by analogy.

4.2 Klein tunneling is the transfer already measured

The flagship of the massless-cone register is Klein tunneling, and it is a flagship precisely because it is the crystal keeping a relativistic fact that is essentially unreachable in the vacuum. Katsnelson, Novoselov, and Geim predicted that graphene’s chiral massless carriers, governed by $\hat{H}_0 = -i\hbar v_F \boldsymbol{\sigma} \cdot \nabla$ with $v_F \approx 10^6$ m/s playing the role of c , tunnel through an electrostatic barrier with perfect transmission at normal incidence regardless of the barrier’s height or width — the opposite of ordinary tunneling, which decays exponentially [32]. The same paper makes the point that gives the transfer its force: the genuine Klein paradox for real electrons needs a potential drop of order mc^2 across a Compton wavelength, hence fields above 10^{16} V/cm, so the effect is a tabletop phenomenon in graphene and almost nowhere

else [32]. Young and Kim measured it. In a graphene p – n junction with top gates under 20 nm wide they resolved the phase-sensitive signature — a shift in the conductance-oscillation fringe pattern at a characteristic field $B^* \approx 0.3$ T — that distinguishes perfect from merely near-perfect transmission at normal incidence, which a bulk-resistance measurement cannot [67]. Bulk-resistance evidence corroborates it more weakly [56], and two further platforms confirm the same physics independently: a trapped-ion simulation of the Klein paradox [24] and a quasirelativistic Bose–Einstein condensate in a bichromatic lattice that transmits only in the linear-dispersion case [53]. Three platforms, one relativistic fact, transferred exactly into matter where the vacuum forbids it — this is the strongest single entry in the theorem.

The window that makes such transfers possible is itself an object one can engineer, not merely find. In a tunable honeycomb optical lattice the Dirac points — the emergent massless cones — can be created, moved through the Brillouin zone, and, beyond a critical anisotropy, merged and annihilated [58]: the designer’s inverse of the census, in which the speed of light is dialled rather than measured, and the programme the sibling paper develops [16].

4.3 Time dilation transfers, but only inside a window we can size

The second register is where honesty about grading earns its place. A quasiparticle wavepacket moving at group velocity v in an emergent-Lorentz window has an internal oscillation — a zitterbewegung, a spinor phase — that is a clock, and that clock runs slow. This is graded structured, not exact: it holds only in the window, and the window has an edge that band curvature sets.

Corollary 1 (Dilation window). *In the emergent-Lorentz window the internal clock of a wavepacket at group velocity v dilates by $\gamma_{\text{eff}} = (1 - v^2/c_{\text{eff}}^2)^{-1/2}$ — 1.15 at $v = 0.5 c_{\text{eff}}$ and 2.29 at $0.9 c_{\text{eff}}$. Writing the corrected dispersion $\varepsilon(k) = \hbar c_{\text{eff}} k [1 - (k/\Lambda_{\text{curv}})^2 + \dots]$, the fractional error in γ_{eff} from band curvature stays below one percent for $k \lesssim 0.1 \Lambda_{\text{curv}}$. For graphene $\Lambda_{\text{curv}} \approx 1.5 \times 10^9 \text{ m}^{-1}$, so the clean window is $k \lesssim 1.5 \times 10^8 \text{ m}^{-1}$ (Appendices A and C).*

The numbers are the point of the corollary. Dilation is a real, testable effect where the dispersion is linear to within a percent, and a computable artefact of band structure where it is not; the correction $\propto (k/\Lambda_{\text{curv}})^2$ is the leading even term of the band’s departure from linearity, derived in Appendix A in the standard $k \cdot p$ formalism [45]. Two adjacent effects sit in the same register and inherit the same window condition. Zitterbewegung itself has been imaged directly, in a trapped-ion Dirac simulator [23] and in a spin-orbit-coupled condensate [39], so the internal clock is not hypothetical. And the velocity that defines the window is not even constant: the graphene v_F that fixes c_{eff} runs by a factor of three as the Dirac point is approached [18], so “the” dilation factor is a slowly varying function of where in the band the wavepacket sits. Structured, corrected, and window-limited — but real, and the window is wide enough to measure in.

Concretely, a dilation measurement compares the internal-oscillation frequency of a moving wavepacket against a stationary reference of the same species, both timed by a clock external to the emergent world, and asks whether the moving one is slow by γ_{eff} . The window condition is what makes that comparison clean: while the packet is built from components with $k \lesssim 0.1 \Lambda_{\text{curv}}$ the band is linear to within a percent and γ_{eff} is the whole effect, whereas at larger k the reading is contaminated by curvature at a level Appendix A computes — correctable, but not to be ignored.

4.4 Boost invariance fails, and the failure is the rest of the paper

The third register fails provably, and the failures are not marginal. Boost invariance is the first and the deepest: a crystal has a ground state and a lattice, and both select a rest frame, so there is no continuous symmetry carrying one inertial description into another. The

counterexample is constructive — the lattice Hamiltonian is manifestly not invariant under a Lorentz boost with c_{eff} — and it is exactly the fact that makes the crystal an ether. That is not a defect to apologize for; it is the object Section 5 measures, because an ether is detectable and the paper computes the experiment that detects it. Relativity of simultaneity fails with it, beyond the low-energy window: the notion of a boost-dependent simultaneity slicing requires the boost symmetry that the lattice lacks, and outside the linear window there is not even an approximate one. The universality of the energy-momentum relation fails for a separate reason: a crystal has many bands, each with its own dispersion and its own c_{eff} , so there is no single invariant speed for all excitations the way c is single for all massless fields in the vacuum [17]. A magnon, a phonon, and a Dirac electron in the same material do not share a light cone. The census makes this concrete in a single table: co-existing excitations of one crystal can differ in cone speed by orders of magnitude, so “the” speed of light of a material is a category error — there are as many as there are bands. These three are graded exact-as-failures: each has a proof or an explicit counterexample in Appendix A, and together they draw the boundary of how far the census’s speeds are allowed to be taken as a relativity. The crystal keeps causality and, in a window, kinematics; it does not keep the principle of relativity, and the next section turns that loss into a measurement.

5 Running Michelson–Morley where you already know it fails

The lattice is an ether. Boost invariance failed provably in Section 4, which means the crystal has a rest frame, which means — unlike the vacuum of 1887 — there is really something for an interferometer to find. The honest way to use that fact is to invert the historical experiment: Michelson and Morley looked for an ether and, correctly, found nothing; we build the experiment inside a medium we know has a preferred frame and ask not whether we detect it but what detecting it measures.

Proposition 3 (The in-crystal Michelson–Morley). *A quasiparticle interferometer that compares cone speeds along and transverse to a medium drift v_d registers a first-order anisotropy $\delta c/c \simeq v_d/c_{\text{eff}}$ — large, and the inverse of the historical null. Yet an interferometer built entirely of the emergent quasiparticles nulls: its arms contract by the acoustic factor $\gamma_{\text{eff}} = (1 - v_d^2/c_{\text{eff}}^2)^{-1/2}$, exactly the factor that shifts its timing. The lattice rest frame is therefore detectable only with a clock and ruler external to the emergent world, and a positive result measures the quality of the emergent window, not anything about fundamental relativity.*

The magnitude is the first surprise and it is genuine.

Transfers exactly

front-speed causality — the Lieb–Robinson velocity c_{LR} is a hard front limit; influence outside the cone decays exponentially.

effective causal order — operations spacelike-separated in the c_{LR} metric commute up to exponential tails.

massless-cone kinematics — where the dispersion is linear, Dirac kinematics replay; **Klein tunneling measured in graphene** is the flagship [32, 67].

Transfers with a stated correction

internal-clock time dilation — a quasi-particle wavepacket at $v < c_{\text{eff}}$ in an emergent-Lorentz window runs slow by $\gamma_{\text{eff}} = (1 - v^2/c_{\text{eff}}^2)^{-1/2}$.

leading correction $\propto (k/\Lambda_{\text{curv}})^2$ from band curvature (Appendix A); the effect is clean for $k \lesssim 0.1 \Lambda_{\text{curv}}$.

zitterbewegung and velocity renormalization sit here too — real, but window-limited [18, 23].

Fails provably

boost invariance — the lattice/ground state selects a frame: the crystal is an ether, and an in-crystal Michelson–Morley detects it (Proposition 3).

relativity of simultaneity beyond the low-energy window.

universal E - p relation — dispersion is band-specific; there is no single invariant speed across all excitations of one crystal.

Figure 2: The transfer theorem (Proposition 2): which special-relativistic kinematic facts replay inside a crystal, and which cannot. The left column transfers exactly because it follows from locality alone (the Lieb–Robinson structure), with Klein tunneling the measured flagship; the middle column transfers only inside the emergent-Lorentz window, carrying a band-curvature correction the paper computes; the right column fails as a matter of proof, because a lattice has a rest frame that a relativistic vacuum does not. Each entry is argued at its stated order in Appendix A.

In a flowing or driven medium the cone speed along the drift is $c_{\text{eff}} + v_d$ one way and $c_{\text{eff}} - v_d$ the other, so the along-versus-transverse anisotropy is $\delta c/c \simeq v_d/c_{\text{eff}}$ — first order in the drift, not the second-order $v^2/c^2 \sim 10^{-8}$ that made the original signal so punishing to chase. At a tenth of the cone speed the anisotropy is ten percent (Appendix C). Figure 3 is the schematic: a source, a splitter, an along-drift arm and a transverse arm, and a first-order fringe shift. As a device this is not exotic — ultrasonic measurements already resolve a field-angle dependence of the sound velocity in the Weyl semimetal TaAs, a real anisotropic cone speed in exactly the material class the census uses for its Weyl row [37], and the general emergent-metric formalism that licenses the whole construction is standard analogue gravity [4].

The design is platform-agnostic, and three realizations are within reach. Cold-atom quench interferometry can compare correlation-cone slopes along two lattice axes in a drifting condensate; magnon interferometry in a YIG film can race spin-wave packets along and across an imposed drift; and an electronic Fabry–Pérot cavity in graphene can compare Dirac-cone transit times across a biased junction. Each measures the same quantity — the along-versus-transverse cone anisotropy — and each supplies its own external clock and ruler, the laboratory apparatus that, as we are about to see, is precisely what a fully-emergent version would lack.

The second surprise is the correction, and it is where the honesty of the inversion lives. Barceló and Jannes proved the exact result: a Michelson–Morley interferometer built out of the same quasiparticles whose signals it times cannot detect uniform motion through the medium, because the binding fields of its own arms obey the acoustic metric and the arms contract by exactly the acoustic Lorentz factor $\gamma_{\text{eff}} = (1 - v_d^2/c_{\text{eff}}^2)^{-1/2}$ needed to cancel the timing shift [3]. The contraction is real — a dynamical effect of the field equations, not an appearance — and it is graded exact, because it is a theo-

rem. Its consequence is precise: internal observers, confined to apparatus built from the emergent world, run the same null Michelson and Morley did, even though their medium has a rest frame. The lattice frame becomes visible only when the clock and ruler are external to the emergent physics — a laboratory clock, a crystallographic length — which in a real crystal they always are. This is why a positive in-crystal result carries the modest meaning we assign it: it does not bear on whether fundamental relativity is exact, only on how good the crystal’s emergent window is, because a window with a first-order drift anisotropy is a window with a detectable frame. Run where it is known to fail, the experiment is a window-quality meter, and that is the only thing we claim it is.

The moving-medium half of the design already has measured cousins, which fix the scale of a real drift signal even though none is a Michelson–Morley. A stirred sodium condensate has a critical velocity $v_c \approx 1.6$ mm/s against a Bogoliubov sound speed $c \approx 6.2$ mm/s — a factor of four apart, because the dissipation is set by vortex-pair nucleation, not by the Landau criterion alone [50]. Exciton-polariton fluids show the same velocity-dependent onset of dissipation, suppressed scattering below a critical flow and a Cerenkov-like wake or oblique solitons above it [1, 2]. These are the drift regimes an in-crystal Michelson–Morley would run in, and they make the governing distinction concrete: an absolute laboratory time coexists with an emergent quasiparticle time [29], and only an apparatus that can read the former ever detects the drift.

5.1 One dataset, three verdicts, and a point no theory explains

The transfer theorem makes predictions about cone shapes; one published dataset is sharp enough to confront. We take Richerme et al.’s trapped-ion α -scan —

eight error-barred light-cone exponents, printed in the text, not buried in a supplement — and set them against the long-range Lieb–Robinson thresholds, declaring the convention we compare under and reporting what we find.

Proposition 4 (The confrontation). *Re-analyzed against the long-range Lieb–Robinson thresholds, the trapped-ion light-cone exponents $\zeta(\alpha)$ of Richerme et al. [51], with couplings $\propto 1/r^\alpha$ and boundary fit $r \sim t^\zeta$, satisfy three statements at once: (i) all four ranges $\alpha \in \{0.63, 0.83, 1.00, 1.19\}$ lie below every threshold, where a linear cone $\zeta = 1$ is not guaranteed and super-linear growth is required, which the data show, violating the nearest-neighbour bound $v_{\text{LR}} = 12eJ_{\text{max}}$ for $\alpha < 1$; (ii) the data cannot discriminate between the three thresholds, which agree throughout this range; (iii) the XY point at $\alpha = 1.19$ ($\zeta = 1.67 \pm 0.08$, super-linear despite $\alpha > 1$) is explained by no known theory.*

The convention is Richerme’s own: the light-cone boundary is the fixed contour $C_{ij} = 0.04 C_{\text{max}}^{ij}$, fit as $r \sim t^\zeta$, for both an Ising and an XY Hamiltonian [51]. The measured exponents are $\zeta_{\text{Ising}} = 1.70 \pm 0.11$, 1.55 ± 0.07 , 1.57 ± 0.07 , 0.97 ± 0.17 and $\zeta_{\text{XY}} = 1.5 \pm 0.3$, 2.8 ± 0.2 , 2.2 ± 0.2 , 1.67 ± 0.08 at $\alpha = 0.63, 0.83, 1.00, 1.19$ respectively [51], plotted in Figure 4 against the free ($\alpha > 2$), Frobenius ($\alpha > 5/2$), and worst-case ($\alpha > 3$) thresholds [22, 34, 61, 62].

We use this dataset because its eight exponents and their error bars are printed in the main text, requiring no digitization, and because it matches the proposition’s own fallback language exactly [52]. A short-range alternative exists: Cheneau’s Bose–Hubbard cone is even more self-contained, carrying both its measured velocities and its own closed-form v_{max} prediction in one table [12], but it tests the bosonic short-range bound rather than the long-range thresholds Proposition 2 is framed around, so we keep it as the secondary angle and confront the ion data here.

The first verdict is a clean win for the theorems in their qualitative form. Every α measured falls below every threshold, in the regime where no finite-velocity linear cone is guaranteed and the theory instead requires faster-than-linear growth; the exponents are super-linear ($\zeta > 1$) at seven of the eight points, and the nearest-neighbour bound $v_{\text{LR}} = 12eJ_{\text{max}}$ is violated outright for $\alpha < 1$ [51]. Read against a naive short-range bound the data are decisive: the cone is genuinely non-local. The second verdict is the honest limit on that win. The three thresholds — worst-case, typical-state, and free — all lie above $\alpha = 2$, and Richerme’s range stops at $\alpha = 1.19$, so the data sit where all three theories predict the same non-linear behaviour and cannot say which is being tested [22, 61]. Discriminating between them would need an α -scan crossing a threshold, from perhaps 1 to 4, which this experiment does not provide. The third verdict is the one that resists resolution. The XY point at $\alpha = 1.19$




Figure 3: The in-crystal Michelson–Morley (Proposition 3). A quasiparticle interferometer compares cone speeds along and transverse to a medium drift v_d ; the first-order anisotropy $\delta c/c \simeq v_d/c_{\text{eff}}$ is linear in the drift and therefore large — the opposite of the vanishing signal Michelson and Morley chased. The subtlety Barceló and Jannes [3] make exact is that an interferometer built *entirely* of the emergent quasiparticles nulls, because its arms contract by the same acoustic $\gamma_{\text{eff}} = (1 - v_d^2/c_{\text{eff}}^2)^{-1/2}$ that shifts the timing; the crystal experiment detects the lattice rest frame only because the experimenter’s clock and ruler are external to the emergent world. Running the experiment where it is known to “fail” measures the emergent window’s quality, not fundamental relativity. Drift magnitudes in Appendix C.

grows super-linearly, $\zeta = 1.67 \pm 0.08$, even though $\alpha > 1$ and a naive counting would call the interaction effectively short-ranged; the original authors already flagged that no known analytic theory predicts it, and numerics confirm it persists to larger systems [51]. The confrontation names this tension rather than dissolving it.

We hedge the novelty deliberately. To our knowledge this explicit numeric confrontation of Richerme’s eight exponents against a specific threshold prediction has not been published: the theory papers cite the ion data as motivation without re-analyzing its numbers, and the most recent work in the area does the same, citing Richerme and Jurcevic only as prior evidence of deviations from light-cone propagation [26]. That is “not found in a live search,” not a proof of absence, and we state it as such. The tabulation, thresholds, and verdict logic are the paper’s own seeded computation (Appendix C); the dataset is Richerme’s, and the analysis convention is declared above so the comparison can be reproduced or disputed.

6 The gap between the crystal’s window and the vacuum’s is the argument

The crystal is the one system in which we know, rather than suppose, that relativity is emergent. That makes




Figure 4: The confrontation (Proposition 4). The measured light-cone exponents $\zeta(\alpha)$ of the trapped-ion long-range quench of Richerme et al. [51] — eight error-barred numbers, Ising and XY — against the long-range Lieb–Robinson threshold theory. All four interaction ranges ($\alpha = 0.63$ – 1.19) fall in the shaded regime below every threshold ($\alpha > 2$ free, $> 5/2$ typical-state, > 3 worst-case; [22, 34, 61, 62]), where no finite-velocity linear cone ($\zeta = 1$) is guaranteed and the theorems require super-linear growth — which the data show, and where the nearest-neighbour bound $v_{LR} = 12eJ_{\max}$ is violated. Two honest limits follow: the data cannot discriminate between the three thresholds (all agree here), and the XY point at $\alpha = 1.19$ ($\zeta = 1.67 \pm 0.08$, super-linear despite $\alpha > 1$) is explained by no known theory. To our knowledge this explicit numeric confrontation has not been published. Values in Appendix C.

it a calibrated instrument for a question fundamental physics can only ask indirectly: if our own vacuum’s Lorentz symmetry were emergent from some deeper structure with a preferred frame, how good would its window have to be, and how good is it? The crystal supplies a measured floor for “how good an emergent window can look,” and the comparison with the vacuum’s Lorentz-violation bounds is not close.

Proposition 5 (What the ether teaches; graded structured). *Read as a template, the crystal’s emergent-Lorentz window has a size, a breaking pattern, and preferred-frame observables. Its best window quality is of order $\delta c/c \sim 10^{-2}$. The best measured vacuum electron Lorentz-violation bound is $\sim 10^{-16}$, against a natural size $\sim 10^{-22}$. A naive Planck-frame quantum gravity would generically induce percent-level, $\sim 10^{-2}$, violations — the crystal’s own level — unless fine-tuned by some twenty orders of magnitude or protected by a symmetry. Our vacuum’s window is therefore at least ~ 14 orders better than any crystal’s, and that gap is the argument, graded structured: a fact an emergent-spacetime proposal must explain, not a coincidence it may assume away.*

The numbers are worth setting side by side because

their distance is the content. The best crystal window we can point to sits in the $\delta c/c \sim 10^{-2}$ class — a first-order drift anisotropy is a ten-percent effect, and even the cleanest Dirac window runs its own v_F by a factor of three near the node. The vacuum is another matter. Liberati’s review gives the best measured constraint on a rotationally-invariant dimension-four electron Lorentz-violation coefficient as of order 10^{-16} , while the “natural” size such a coefficient would take from a term suppressed only by m_e/M_{Pl} is of order 10^{-22} [40]. Both numbers matter and they must not be conflated: 10^{-16} is a very strong bound, and it still leaves six orders of room above the natural expectation. The photon sector is bounded even more tightly by astrophysical baselines — a 31 GeV photon from a gamma-ray burst at redshift 0.903 arriving with no energy-dependent delay constrains linear Lorentz violation to beyond the Planck energy [21] — and the full sector-by-sector compendium is Kostelecký and Russell’s living tables [33], with the framework’s history in Mattingly’s earlier review [44]. Against any of these, the crystal’s 10^{-2} window is crude by fourteen orders or more.

The reason that gap is an argument and not a curiosity is a naturalness result the paper imports as a cautionary parallel, not as a bound on crystals. Collins, Perez, Sudarsky, Urrutia, and Vucetich showed that naively combining Standard Model particle physics with a Planck-scale preferred frame generically produces Lorentz violation about twenty orders of magnitude larger than dimensional analysis suggests — percent-level, $\sim 10^{-2}$, in the low-energy theory — because loop corrections drag high-energy violation down into dimension-four operators with no natural suppression, unless the bare parameters are fine-tuned to implausible precision [15]. Polchinski defended the argument against a published attempt to evade it and identified supersymmetry as the one broadly recognized way to protect small violations from being radiatively enhanced [49]. The parallel to our census is exact in structure and must be stated carefully: the crystal is the “preferred-frame emergent world” made concrete, and its window sits precisely at the $\sim 10^{-2}$ level the naive quantum-gravity estimate predicts for an unprotected emergent Lorentz symmetry. The vacuum’s window is fourteen orders better. So either the vacuum is fine-tuned, or its Lorentz symmetry is protected by something the crystal’s is not — and a serious emergent-spacetime proposal owes an account of which. That is the template the crystal provides: not a model of the vacuum, but a measured example of how bad an unprotected emergent window looks, against which the vacuum’s suspiciously good one stands out.

We are explicit about the register of this last claim, because it is the paper’s most speculative. Reading the crystal window as a template for constraining fundamental physics is graded structured, not exact: the crystal and the vacuum share the logical form of an emergent Lorentz symmetry over a preferred frame, and they share

the naturalness estimate, but the crystal is not a model of the vacuum and we do not treat it as one. Volovik’s program does treat a condensed-matter ground state as a working model of an emergent gravitating vacuum — emergent Weyl fermions, emergent gauge fields, and an emergent metric all at once, with the cosmological constant nullified by the same mechanism [30, 64]. That is the strong form, and we cite it as the strong form we do not endorse. Our claim is narrower and, we think, harder to dispute: the crystal fixes the scale of an unprotected emergent window, and the vacuum beats it by fourteen orders, and that number is a datum any emergent-spacetime picture has to meet.

What the paper delivers, then, is four things at their stated registers. A census of c_{LR} across nine measured platforms spanning nine and a half orders of magnitude — the honest span, not the round eleven — under one extraction convention. A transfer theorem that sorts special-relativistic kinematics into what a crystal keeps exactly (causality, causal order, and massless-cone physics, with Klein tunneling the flagship already measured on three platforms), what it keeps with a computed correction (internal-clock dilation, clean where $k \lesssim 0.1 \Lambda_{\text{curv}}$), and what it loses by proof (boost invariance, and with it the principle of relativity). An in-crystal Michelson–Morley whose drift signal is real and first-order yet nulls for a fully-emergent apparatus, so that a positive result grades the window and says nothing about fundamental relativity. And a confrontation of Richerme’s α -scan that the theorems win qualitatively, cannot resolve between competing thresholds, and leave with one point unexplained. Each claim states where it would break: the census by a measurement in the empty 10^{-2} – 10^2 m/s band, the dilation corollary by curvature beyond its window, the ether inversion by an apparatus that is secretly external to the emergent world, the confrontation by an α -scan that crosses a threshold. Every crystal has its own speed of light; what it does not have, and our vacuum apparently does, is a reason for that speed to be as good as light.

A The transfer theorem at its stated order

This appendix argues Proposition 2 register by register, at the order claimed in Section 4, and derives the band-curvature correction of Corollary 1.

A.1 Transfers exactly: proofs from locality

Front-speed causality. This is Theorem 2.1 restated. The exponential-tail form $\| [A_x(t), B_y] \| \leq C e^{-d(x,y) - c_{\text{LR}} t} / \xi$ gives, for any pair of strictly local observables, a commutator that is exponentially small once $d(x,y) > c_{\text{LR}} t$; the front speed c_{LR} is a property of the couplings alone

[8, 41]. No emergent Lorentz symmetry is used, which is why the entry is graded exact rather than window-limited.

Effective causal order. Let two operations act on regions X and Y at times separated so that X and Y are spacelike in the c_{LR} metric, $d(X,Y) > c_{\text{LR}} \Delta t$. Then $\| [A_X(t), B_Y] \|$ is exponentially small, so the two operations commute up to that tail and their order is immaterial to any observable to the same accuracy. The set of such relations is a partial order — reflexive, antisymmetric up to the tail, and transitive because commutator bounds compose — so a light-cone causal order exists inside the crystal exactly as in Minkowski space [8]. The generation of correlations and topological order across the cone obeys the same bound, so the order is not merely kinematic bookkeeping but constrains what dynamics can build [8].

Massless-cone kinematics. Where the low-energy dispersion is linear, $\varepsilon(k) = \hbar c_{\text{eff}} k$, the effective Hamiltonian is the massless Dirac operator $\hat{H}_0 = \hbar c_{\text{eff}} \boldsymbol{\sigma} \cdot \mathbf{k}$ with c_{eff} in the role of c [9, 32]. Every kinematic consequence of that operator — chirality, pseudospin–momentum locking, and the perfect normal-incidence transmission of Klein tunneling — follows algebraically and is therefore exact within the linear window. The measured flagship is the phase-resolved Klein signature of Young and Kim [67]; the scope condition (linearity of the band) is the entire qualification and is checked, not assumed, for each material.

A.2 Transfers with a correction: the dilation window

Inside the window a wavepacket at group velocity $v = c_{\text{eff}} \beta$ carries an internal phase — the spinor precession that appears macroscopically as zitterbewegung [23, 39]. Treating that phase as a clock and boosting with the emergent metric of velocity c_{eff} gives the proper-time relation $d\tau = dt / \gamma_{\text{eff}}$ with $\gamma_{\text{eff}} = (1 - \beta^2)^{-1/2}$, evaluated in Corollary 1 at $\gamma_{\text{eff}} = 1.005, 1.048, 1.155, 1.400, 2.294$ for $\beta = 0.1, 0.3, 0.5, 0.7, 0.9$ (Appendix C). This is exact only for a strictly linear band; a real band curves, and the curvature is the correction.

Expand the dispersion about the node in the standard $k \cdot p$ form [45]. Isotropy and inversion symmetry forbid a linear-in- k correction to the slope, so the leading departure is even:

$$\varepsilon(k) = \hbar c_{\text{eff}} k \left[1 - \left(\frac{k}{\Lambda_{\text{curv}}} \right)^2 + \dots \right], \quad (1)$$

which defines the band-curvature scale Λ_{curv} as the wavenumber at which the cubic term reaches the linear one. The group velocity is

$$v_g(k) = \frac{1}{\hbar} \frac{d\varepsilon}{dk} = c_{\text{eff}} \left[1 - 3 \left(\frac{k}{\Lambda_{\text{curv}}} \right)^2 + \dots \right], \quad (2)$$

so the effective cone speed the wavepacket actually experiences is reduced from c_{eff} by $3(k/\Lambda_{\text{curv}})^2$. Propagating this into $\gamma_{\text{eff}}(\beta)$ with $\beta = v_g/c_{\text{eff}}$, the fractional error incurred by using the uncorrected γ_{eff} is, to leading order,

$$\frac{\delta\gamma_{\text{eff}}}{\gamma_{\text{eff}}} = \gamma_{\text{eff}}^2 \beta^2 \cdot 3 \left(\frac{k}{\Lambda_{\text{curv}}} \right)^2 + \mathcal{O} \left[\left(\frac{k}{\Lambda_{\text{curv}}} \right)^4 \right], \quad (3)$$

which for wavepackets that are not ultra-relativistic in the emergent sense is of order $(k/\Lambda_{\text{curv}})^2$. Requiring this below one percent gives $k \lesssim 0.1 \Lambda_{\text{curv}}$. For graphene the curvature scale is set by the crossover from the linear Dirac cone to the full tight-binding band, $\Lambda_{\text{curv}} \approx 1.5 \times 10^9 \text{ m}^{-1}$, so the clean dilation window is $k \lesssim 1.5 \times 10^8 \text{ m}^{-1}$ (Appendix C). The velocity renormalization of Elias et al. [18] enters the same window: since $c_{\text{eff}} = v_F$ runs with density, the dilation factor is a slowly varying function of the operating point, and a precise dilation measurement must fix the density as well as k .

A.3 Fails provably: counterexamples

Boost invariance. The lattice Hamiltonian $H = \sum_{\langle ij \rangle} t_{ij} c_i^\dagger c_j + \dots$ is invariant under the lattice space group and under time translation, but not under any boost $x \rightarrow x - ut$, $t \rightarrow t$ with a finite u : the hopping graph and the ground state both single out the rest frame in which the ions are at rest. There is no continuous one-parameter family of symmetries connecting inertial descriptions, so boost invariance fails by construction — the explicit counterexample is the Hamiltonian itself. This is the ether: the failure is not a looseness of the bound but the existence of a preferred frame, which Section 5 measures.

Relativity of simultaneity. A boost-dependent simultaneity slicing is defined only where a boost symmetry exists to define it. Within the linear window one may construct an approximate emergent slicing, but it inherits the window's edge from Eq. (1); beyond the window, and for any excitation whose band is not linear, there is not even an approximate boost, so simultaneity has no emergent meaning to be relative.

Universality of $\varepsilon(k)$ across bands. A crystal carries many bands, each with its own dispersion and its own low-energy c_{eff} ; a magnon, an acoustic phonon, and a Dirac electron in one material have different cone speeds, as the census's own spread across co-existing excitations shows. There is thus no single invariant speed for all excitations, unlike the vacuum where c is common to all massless fields. At small interaction-range exponent the situation is worse than non-universal: quasi-locality itself can break down and no bounded cone exists at all [17], the provable end of the register.

B How each census number was extracted

The census of Table 1 uses one convention, and the value of a compilation is that the convention is applied visibly and its failures are declared. This appendix states the convention and then walks the nine rows, giving for each the source number, the extraction step, and the caveat that keeps the row honest.

The convention. For every platform we report the cone-front group velocity — the slope of the measured or dispersion-derived light cone — at the relevant filling or state, in metres per second. Where a platform tunes (by interaction range, filling, or dispersion regime) we report the measured range rather than a single value; where a source reports a dimensionless cone we convert once, with the conversion named in the row and here. Three conversions are the verifier's, not the source's, and are labelled “derived”; one platform (superconducting qubits) is not a spatial crystal at all and is flagged wherever it appears.

Bose–Hubbard cold atoms. Cheneau et al. [11] report the dimensionless propagation velocity $v\hbar/(Ja_{\text{lat}}) = 5.0(2), 5.6(5), 5.0(2)$ for $U/J = 5.0, 7.0, 9.0$, with $J/\hbar = 691(23), 522(14), 422(11) \text{ s}^{-1}$ and lattice period $a_{\text{lat}} = 532 \text{ nm}$. Multiplying, the physical doublon–holon cone is $\approx 1.84, 1.56, 1.12 \text{ mm/s}$, so the row is $1.1\text{--}1.8 \times 10^{-3} \text{ m/s}$. Caveat: this is a quasiparticle-pair group velocity at a specific quench, and the rigorous bound beneath it is not the clean 1972 cone but the bosonic almost-linear $(t \log^2 t)$, state-restricted result [20, 35, 66].

Trapped-ion spin chains. Jurcevic et al. [31] extract the maximum magnon group velocity from the fitted long-range dispersion of 7–15 $^{40}\text{Ca}^+$ spins; for the shortest range it matches the nearest-neighbour model and it is consistent with divergence as the range grows. Richerme et al. [51] give the light-cone exponents on 11 $^{171}\text{Yb}^+$ ions used in the confrontation. The metres-per-second row, $2 \times 10^{-3}\text{--}1 \times 10^{-2}$, is the group-velocity class these chains occupy at $J_{\text{max}} \sim 400 \text{ Hz}$ over micron-scale ion spacings. Caveat: $v_g^{\text{max}}(\alpha)$ is figure-extracted in the source [31], and the cone speed is not a fixed material constant — confinement can suppress it substantially [57].

Rydberg arrays. Lienhard et al. [42] observe the finite-speed delay in correlation build-up in a short-range Ising antiferromagnet of up to 36 spins, matched to a short-time expansion; Bernien et al. [5] establish the 51-atom platform. The row $3 \times 10^{-3}\text{--}3 \times 10^{-2} \text{ m/s}$ is the blockade-radius-times-Rabi-rate class expected for these arrays. Caveat: the sources report the qualitative delay rather than a single printed velocity, so the row is a

platform-class estimate, wider than a single quoted number.

Magnons (YIG). The spin-wave group velocity of yttrium iron garnet spans the dipolar (magnetostatic) regime at long wavelength through the exchange regime at short wavelength [55]; Böttcher et al. [6] measure the exchange stiffness and report velocity ratios in gallium-substituted films. The row 2×10^2 – 2×10^4 m/s brackets the two regimes. Caveat: the review is paywalled and reports no single velocity, and the Ga:YIG paper reports ratios, not absolute values, so the row is a cross-literature range with the regime named, not a single-source number.

Acoustic phonons (Si). MacCabe et al. [43] state the longitudinal sound velocity of crystalline silicon, $v_l = 8.433$ km/s, in the appendix of a nanomechanics paper. The row is the point 8.4×10^3 m/s. Caveat: this is a standard materials constant restated for the paper’s own purposes, not a new measurement; it anchors the phonon band of the census at a value no one disputes.

Superconducting qubits. Braumüller et al. [7] build a 3×3 hard-core Bose–Hubbard lattice of transmons with $J/2\pi = 8.1$ MHz and measure information propagation via out-of-time-ordered correlators on a Manhattan-distance metric. The row 5×10^3 – 1.5×10^4 m/s is $J/2\pi$ times an assumed chip pitch of a few millimetres. Caveat: this is the one row that is not a crystal velocity. The device is a wired circuit, its cone is dimensionless, and the conversion is explicit and unavoidable; the row is marked with a dagger in Table 1 and carries the same flag in Figure 1. It shows the platform’s ordinal place, not a velocity on the same footing as the others.

Spinons (KCuF₃). Lake et al. [36] give the two-spinon continuum’s lower boundary $\omega_l(k) = (\pi/2)J|\sin k|$ with $J = 33.5$ meV, so the small- k spinon velocity is $(\pi/2)J$ in energy–wavevector units. Converting with the c -axis chain spacing ≈ 3.9 Å gives $\approx 3 \times 10^4$ m/s. Caveat: derived. The velocity in metres per second is the verifier’s conversion using a crystallographic spacing not printed in the source; the source prints the exchange and the dispersion boundary, and the $\approx 98\%$ sum-rule agreement of the two- and four-spinon continuum supports the dispersion, but the m/s figure is ours.

Weyl/Dirac semimetals. The node Fermi velocity of TaAs-class semimetals, 1×10^5 – 1×10^6 m/s, is carried from the sibling paper’s own verified pool rather than re-derived here [16]. Caveat: this is a pre-authorized site-internal source; the primary Weyl papers report node separations and Fermi-arc structure but not a single

printed v_F in metres per second, which is why the value is taken from the series pool by design.

Graphene electrons. Novoselov et al. [48] give the headline Dirac velocity $\approx 1 \times 10^6$ m/s ($\approx c/300$); Elias et al. [18] measure it running up to $\approx 3 \times 10^6$ m/s near the neutrality point through interaction-driven renormalization. The row is the range 1 – 3×10^6 m/s. Caveat: the row is deliberately a range because there is no single graphene v_F ; the low-density renormalized value and the conventional value are both correct at their densities and must not be quoted as one constant — the same running that sets the top of this row sets the operating-point dependence of the dilation window in Appendix A.

The span. The extremes are the Bose–Hubbard 1.1×10^{-3} m/s and the graphene 3×10^6 m/s, a ratio of $10^{9.44}$: nine and a half orders of magnitude, the number Section 3 quotes in place of the rounder eleven, computed in Appendix C.

C The seeded computation

The census, the confrontation, the dilation window, the in-crystal Michelson–Morley magnitudes, and the closing gap all come from one script, `sim_lr.py`, whose source and outputs (`sim_lr_output.txt/.json`) ship with the paper. The computation is analytic; the seed recorded in the source (20260919) is for provenance only.

The census and its honest span

The script collects the measured cone speeds of Section 3 and reports the span (Fig. 1). The verified extremes are the 1.1–1.8 mm/s Bose–Hubbard doublon–holon cone [11] and the 1 – 3×10^6 m/s graphene Dirac cone [18, 48], a ratio of $10^{9.44}$ — nine and a half orders of magnitude, which the paper quotes in place of the rounder “eleven” the thesis first reached for. The engineered simulators (mm/s–cm/s) and the natural condensed-matter carriers (10^2 – 10^6 m/s) sit in two bands with a gap between $\sim 3 \times 10^{-2}$ and $\sim 2 \times 10^2$ m/s that no present platform fills.

The confrontation, the dilation window, the drift, and the gap

For the confrontation (Fig. 4) the script tabulates the eight Richerme exponents $\zeta(\alpha)$ against the linear-cone value $\zeta = 1$ and the long-range thresholds $\alpha > 2, 5/2, 3$; all four $\alpha \in [0.63, 1.19]$ lie below every threshold, the cones are super-linear, and the XY point at $\alpha = 1.19$ ($\zeta = 1.67 \pm 0.08$) exceeds what any threshold theory explains. For the dilation corollary the script evaluates $\gamma_{\text{eff}} = (1 - v^2/c_{\text{eff}}^2)^{-1/2}$ (e.g. 1.15 at $v = 0.5 c_{\text{eff}}$, 2.29 at $0.9 c_{\text{eff}}$) and the band-curvature scale $\Lambda_{\text{curv}} \approx$

$1.5 \times 10^9 \text{ m}^{-1}$ for graphene, so the correction $(k/\Lambda_{\text{curv}})^2$ stays below one percent for $k \lesssim 1.5 \times 10^8 \text{ m}^{-1}$ — the window where the dilation is a clean, testable effect (Appendix A). The in-crystal Michelson–Morley drift anisotropy is first order, $\delta c/c \simeq v_d/c_{\text{eff}}$ (Fig. 3); the acoustic Lorentz–FitzGerald contraction that nulls a fully-emergent interferometer is the exact result of Barceló and Jannes [3]. Finally the gap: the crystal’s emergent-window quality is of order 10^{-2} , the best vacuum electron Lorentz-violation bound is $\sim 10^{-16}$ [40], and the naive Planck-frame expectation of Collins et al. [15] sits at the crystal’s 10^{-2} level unless the vacuum is fine-tuned by some fourteen to twenty orders — the gap the paper reads as its argument.

References

- [1] A. Amo, J. Lefrère, S. Pigeon, C. Adrados, C. Ciuti, I. Carusotto, R. Houdré, E. Giacobino, et al. (2009). Observation of Superfluidity of Polaritons in Semiconductor Microcavities. *Nature Physics*, vol. 5, pp. 805–810. <https://arxiv.org/abs/0812.2748>.
- [2] A. Amo, S. Pigeon, D. Sanvitto, V. G. Sala, R. Hivet, I. Carusotto, F. Pisanello, G. Lemenager, et al. (2011). Polariton Superfluids Reveal Quantum Hydrodynamic Solitons. *Science*, vol. 332, pp. 1167–1170. <https://arxiv.org/abs/1101.2530>.
- [3] C. Barceló and G. Jannes (2008). A Real Lorentz–FitzGerald Contraction. *Foundations of Physics*, vol. 38, pp. 191–199. <https://arxiv.org/abs/0705.4652>.
- [4] C. Barceló, S. Liberati, and M. Visser (2011). Analogue Gravity. *Living Reviews in Relativity*, vol. 14, article 3. <https://arxiv.org/abs/gr-qc/0505065>.
- [5] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, et al. (2017). Probing Many-Body Dynamics on a 51-Atom Quantum Simulator. *Nature*, vol. 551, pp. 579–584. <https://arxiv.org/abs/1707.04344>.
- [6] T. Böttcher, M. Ruhwedel, K. O. Levchenko, Q. Wang, H. L. Chumak, M. A. Popov, I. V. Zavislyak, C. Dubs, et al. (2021). Fast Long-Wavelength Exchange Spin Waves in Partially-Compensated Ga:YIG. *Applied Physics Letters* (2022). <https://arxiv.org/abs/2112.11348>.
- [7] J. Braumüller, A. H. Karamlou, Y. Yanay, B. Kannan, D. Kim, M. Kjaergaard, A. Melville, B. M. Niedzielski, et al. (2022). Probing Quantum Information Propagation with Out-of-Time-Ordered Correlators. *Nature Physics*, vol. 18, pp. 172–178. <https://arxiv.org/abs/2102.11751>.
- [8] S. Bravyi, M. B. Hastings, and F. Verstraete (2006). Lieb–Robinson Bounds and the Generation of Correlations and Topological Quantum Order. *Physical Review Letters*, vol. 97, article 050401. <https://arxiv.org/abs/quant-ph/0603121>.
- [9] A. H. Castro Neto, F. Guinea, N. M. R. Peres, K. S. Novoselov, and A. K. Geim (2009). The Electronic Properties of Graphene. *Reviews of Modern Physics*, vol. 81, pp. 109–162. <https://arxiv.org/abs/0709.1163>.
- [10] M. Cheneau (2022). Experimental Tests of Lieb–Robinson Bounds. In “The Physics and Mathematics of Elliott Lieb: The 90th Anniversary Volume I,” Chapter 10, pp. 225–245. <https://arxiv.org/abs/2206.15126>.
- [11] M. Cheneau, P. Barmettler, D. Poletti, M. Endres, P. Schauß, T. Fukuhara, C. Gross, I. Bloch, et al. (2012). Light-Cone-Like Spreading of Correlations in a Quantum Many-Body System. *Nature*, vol. 481, pp. 484–487. <https://arxiv.org/abs/1111.0776>.
- [12] M. Cheneau, P. Barmettler, D. Poletti, M. Endres, P. Schauß, T. Fukuhara, C. Gross, I. Bloch, et al. (2012). Light-Cone-Like Spreading of Correlations in a Quantum Many-Body System. *Nature*, vol. 481, pp. 484–487. <https://arxiv.org/abs/1111.0776>.
- [13] C.-F. Chen and A. Lucas (2019). Finite Speed of Quantum Scrambling with Long Range Interactions. *Physical Review Letters*, vol. 123, article 250605. <https://arxiv.org/abs/1907.07637>.
- [14] C.-F. Chen, A. Lucas, and C. Yin (2023). Speed Limits and Locality in Many-Body Quantum Dynamics. *Reports on Progress in Physics*, vol. 86, article 116001. <https://arxiv.org/abs/2303.07386>.
- [15] J. Collins, A. Perez, D. Sudarsky, L. Urrutia, and H. Vucetich (2004). Lorentz Invariance and Quantum Gravity: An Additional Fine-Tuning Problem? *Physical Review Letters*, vol. 93, article 191301. <https://arxiv.org/abs/gr-qc/0403053>.
- [16] B. Corbeel (2026). A Strained Crystal Is a Designer Spacetime. *The Crossed Fields Papers*, No. 4, Absolute Digital Publishers. <https://absolutedigitalpublishers.com/papers/crossed-fields-04-designer-spacetime-crystal.pdf>.
- [17] J. Eisert, M. van den Worm, S. R. Manmana, and M. Kastner (2013). Breakdown of Quasilocality in Long-Range Quantum Lattice Models. *Physical Review Letters*, vol. 111, article 260401. <https://arxiv.org/abs/1309.2308>.

- [18] D. C. Elias, R. V. Gorbachev, A. S. Mayorov, S. V. Morozov, A. A. Zhukov, P. Blake, L. A. Ponomarenko, I. V. Grigorieva, et al. (2011). Dirac Cones Reshaped by Interaction Effects in Suspended Graphene. *Nature Physics*, vol. 7, pp. 701–704. <https://arxiv.org/abs/1104.1396>.
- [19] D. V. Else, F. Machado, C. Nayak, and N. Y. Yao (2020). Improved Lieb-Robinson Bound for Many-Body Hamiltonians with Power-Law Interactions. *Physical Review A*, vol. 101, article 022333. <https://arxiv.org/abs/1809.06369>.
- [20] J. Faupin, M. Lemm, and I. M. Sigal (2022). Maximal Speed for Macroscopic Particle Transport in the Bose-Hubbard Model. *Physical Review Letters*, vol. 128, article 150602. <https://arxiv.org/abs/2110.04313>.
- [21] A. A. Abdo et al. (Fermi GBM/LAT Collaboration) (2009). Testing Einstein’s Special Relativity with Fermi’s Short Hard Gamma-Ray Burst GRB 090510. *Nature*, vol. 462, pp. 331–334. <https://arxiv.org/abs/0908.1832>.
- [22] M. Foss-Feig, Z.-X. Gong, C. W. Clark, and A. V. Gorshkov (2015). Nearly Linear Light Cones in Long-Range Interacting Quantum Systems. *Physical Review Letters*, vol. 114, article 157201. <https://arxiv.org/abs/1410.3466>.
- [23] R. Gerritsma, G. Kirchmair, F. Zähringer, E. Solano, R. Blatt, and C. F. Roos (2010). Quantum Simulation of the Dirac Equation. *Nature*, vol. 463, pp. 68–71. <https://arxiv.org/abs/0909.0674>.
- [24] R. Gerritsma, B. Lanyon, G. Kirchmair, F. Zähringer, C. Hempel, J. Casanova, J. J. García-Ripoll, E. Solano, et al. (2011). Quantum Simulation of the Klein Paradox with Trapped Ions. *Physical Review Letters*, vol. 106, article 060503. <https://arxiv.org/abs/1007.3683>.
- [25] Z.-X. Gong, M. Foss-Feig, S. Michalakis, and A. V. Gorshkov (2014). Persistence of Locality in Systems with Power-Law Interactions. *Physical Review Letters*, vol. 113, article 030602. <https://arxiv.org/abs/1401.6174>.
- [26] C.-M. Halati, A. Sheikhan, G. Morigi, C. Kollath, and S. B. Jäger (2025). From Light-Cone to Supersonic Propagation of Correlations by Competing Short- and Long-Range Couplings. *arXiv:2503.13306*. <https://arxiv.org/abs/2503.13306>.
- [27] M. B. Hastings (2010). Locality in Quantum Systems. Lecture notes, Les Houches Summer School. *arXiv:1008.5137*. <https://arxiv.org/abs/1008.5137>.
- [28] M. B. Hastings and T. Koma (2006). Spectral Gap and Exponential Decay of Correlations. *Communications in Mathematical Physics*, vol. 265, pp. 781–804. <https://arxiv.org/abs/math-ph/0507008>.
- [29] G. Jannes (2015). Condensed Matter Lessons About the Origin of Time. *Foundations of Physics*, vol. 45, pp. 279–294. <https://arxiv.org/abs/0904.3627>.
- [30] G. Jannes and G. E. Volovik (2012). The Cosmological Constant: A Lesson from the Effective Gravity of Topological Weyl Media. *JETP Letters*, vol. 96, pp. 215–221. <https://arxiv.org/abs/1108.5086>.
- [31] P. Jurcevic, B. P. Lanyon, P. Hauke, C. Hempel, P. Zoller, R. Blatt, and C. F. Roos (2014). Quasiparticle Engineering and Entanglement Propagation in a Quantum Many-Body System. *Nature*, vol. 511, pp. 202–205. <https://arxiv.org/abs/1401.5387>.
- [32] M. I. Katsnelson, K. S. Novoselov, and A. K. Geim (2006). Chiral Tunnelling and the Klein Paradox in Graphene. *Nature Physics*, vol. 2, pp. 620–625. <https://arxiv.org/abs/cond-mat/0604323>.
- [33] V. A. Kostelecký and N. Russell (2011). Data Tables for Lorentz and CPT Violation. *Reviews of Modern Physics*, vol. 83, pp. 11–31. <https://arxiv.org/abs/0801.0287>.
- [34] T. Kuwahara and K. Saito (2020). Strictly Linear Light Cones in Long-Range Interacting Systems of Arbitrary Dimensions. *Physical Review X*, vol. 10, article 031010. <https://arxiv.org/abs/1910.14477>.
- [35] T. Kuwahara and K. Saito (2021). Lieb-Robinson Bound and Almost-Linear Light Cone in Interacting Boson Systems. *Physical Review Letters*, vol. 127, article 070403. <https://arxiv.org/abs/2103.11592>.
- [36] B. Lake, D. A. Tennant, J.-S. Caux, T. Barthel, U. Schollwöck, S. E. Nagler, and C. D. Frost (2013). Multispinon Continua at Zero and Finite Temperature in a Near-Ideal Heisenberg Chain. *Physical Review Letters*, vol. 111, article 137205. <https://arxiv.org/abs/1307.4071>.
- [37] F. Laliberté, F. Bélanger, N. L. Nair, J. G. Analytis, M.-E. Boulanger, M. Dion, L. Taillefer, and J. A. Quilliam (2020). Field-Angle Dependence of Sound Velocity in the Weyl Semimetal TaAs. *Physical Review B*, vol. 102, article 125104. <https://arxiv.org/abs/1909.04270>.

- [38] T. Langen, R. Geiger, M. Kuhnert, B. Rauer, and J. Schmiedmayer (2013). Local Emergence of Thermal Correlations in an Isolated Quantum Many-Body System. *Nature Physics*, vol. 9, pp. 640–643. <https://arxiv.org/abs/1305.3708>.
- [39] L. J. LeBlanc, M. C. Beeler, K. Jiménez-García, A. R. Perry, S. Sugawa, R. A. Williams, and I. B. Spielman (2013). Direct Observation of Zitterbewegung in a Bose-Einstein Condensate. *New Journal of Physics*, vol. 15, article 073011. <https://arxiv.org/abs/1303.0914>.
- [40] S. Liberati (2013). Tests of Lorentz Invariance: A 2013 Update. *Classical and Quantum Gravity*, vol. 30, article 133001. <https://arxiv.org/abs/1304.5795>.
- [41] E. H. Lieb and D. W. Robinson (1972). The Finite Group Velocity of Quantum Spin Systems. *Communications in Mathematical Physics*, vol. 28, pp. 251–257. <https://www.researchgate.net/publication/38328043>. DOI: 10.1007/BF01645779.
- [42] V. Lienhard, S. de Léséleuc, D. Barredo, T. Lahaye, A. Browaeys, M. Schuler, L.-P. Henry, and A. M. Läuchli (2018). Observing the Space- and Time-Dependent Growth of Correlations in Dynamically Tuned Synthetic Ising Antiferromagnets. *Physical Review X*, vol. 8, article 021070. <https://arxiv.org/abs/1711.01185>.
- [43] G. S. MacCabe, H. Ren, J. Luo, J. D. Cohen, H. Zhou, A. Sipahigil, M. Mirhosseini, and O. Painter (2019). Phononic Bandgap Nano-Acoustic Cavity with Ultralong Phonon Lifetime. *Science*, vol. 370, p. 840 (2020). <https://arxiv.org/abs/1901.04129>.
- [44] D. Mattingly (2005). Modern Tests of Lorentz Invariance. *Living Reviews in Relativity*, vol. 8, article 5. <https://arxiv.org/abs/gr-qc/0502097>.
- [45] E. McCann and M. Koshino (2013). The Electronic Properties of Bilayer Graphene. *Reports on Progress in Physics*, vol. 76, article 056503. <https://arxiv.org/abs/1205.6953>.
- [46] B. Nachtergaele and R. Sims (2006). Lieb-Robinson Bounds and the Exponential Clustering Theorem. *Communications in Mathematical Physics*, vol. 265, pp. 119–130. <https://arxiv.org/abs/math-ph/0506030>.
- [47] B. Nachtergaele and R. Sims (2010). Lieb-Robinson Bounds in Quantum Many-Body Physics. In “Entropy and the Quantum” (Contemporary Mathematics, vol. 529, American Mathematical Society), pp. 141–176. <https://arxiv.org/abs/1004.2086>.
- [48] K. S. Novoselov, A. K. Geim, S. V. Morozov, D. Jiang, M. I. Katsnelson, I. V. Grigorieva, S. V. Dubonos, and A. A. Firsov (2005). Two-Dimensional Gas of Massless Dirac Fermions in Graphene. *Nature*, vol. 438, pp. 197–200. <https://arxiv.org/abs/cond-mat/0509330>.
- [49] J. Polchinski (2011). Comment on “Small Lorentz Violations in Quantum Gravity: Do They Lead to Unacceptably Large Effects?” *Classical and Quantum Gravity*, vol. 29, article 088001 (2012). <https://arxiv.org/abs/1106.6346>.
- [50] C. Raman, M. Köhl, R. Onofrio, D. S. Durfee, C. E. Kuklewicz, Z. Hadzibabic, and W. Ketterle (1999). Evidence for a Critical Velocity in a Bose-Einstein Condensed Gas. *Physical Review Letters*, vol. 83, p. 2502. <https://arxiv.org/abs/cond-mat/9909109>.
- [51] P. Richerme, Z.-X. Gong, A. Lee, C. Senko, J. Smith, M. Foss-Feig, S. Michalakakis, A. V. Gorshkov, et al. (2014). Non-local Propagation of Correlations in Long-Range Interacting Quantum Systems. *Nature*, vol. 511, pp. 198–201. <https://arxiv.org/abs/1401.5088>.
- [52] P. Richerme, Z.-X. Gong, A. Lee, C. Senko, J. Smith, M. Foss-Feig, S. Michalakakis, A. V. Gorshkov, et al. (2014). Non-local Propagation of Correlations in Long-Range Interacting Quantum Systems. *Nature*, vol. 511, pp. 198–201. <https://arxiv.org/abs/1401.5088>.
- [53] T. Salger, S. Kling, C. Grossert, and M. Weitz (2011). Klein Tunneling of a Quasirelativistic Bose-Einstein Condensate in an Optical Lattice. *Physical Review Letters*, vol. 107, article 240401. <https://arxiv.org/abs/1108.4447>.
- [54] N. Schuch, S. K. Harrison, T. J. Osborne, and J. Eisert (2011). Information Propagation for Interacting Particle Systems. *Physical Review A*, vol. 84, article 032309. <https://arxiv.org/abs/1010.4576>.
- [55] A. A. Serga, A. V. Chumak, and B. Hillebrands (2010). YIG Magnonics. *Journal of Physics D: Applied Physics*, vol. 43, article 264002. <https://ui.adsabs.harvard.edu/abs/2010JPhD...43z4002S/abstract>. DOI: 10.1088/0022-3727/43/26/264002.
- [56] N. Stander, B. Huard, and D. Goldhaber-Gordon (2009). Evidence for Klein Tunneling in Graphene p-n Junctions. *Physical Review Letters*, vol. 102, article 026807. <https://arxiv.org/abs/0806.2319>.

- [57] W. L. Tan, P. Becker, F. Liu, G. Pagano, K. S. Collins, A. De, L. Feng, H. B. Kaplan, et al. (2021). Observation of Domain Wall Confinement and Dynamics in a Quantum Simulator. *Nature Physics*, vol. 17, pp. 742–747. <https://arxiv.org/abs/1912.11117>.
- [58] L. Tarruell, D. Greif, T. Uehlinger, G. Jotzu, and T. Esslinger (2012). Creating, Moving and Merging Dirac Points with a Fermi Gas in a Tunable Honeycomb Lattice. *Nature*, vol. 483, pp. 302–305. <https://arxiv.org/abs/1111.5020>.
- [59] M. C. Tran, A. Y. Guo, Y. Su, J. R. Garrison, Z. Eldredge, M. Foss-Feig, A. M. Childs, and A. V. Gorshkov (2019). Locality and Digital Quantum Simulation of Power-Law Interactions. *Physical Review X*, vol. 9, article 031006. <https://arxiv.org/abs/1808.05225>.
- [60] M. C. Tran, C.-F. Chen, A. Ehrenberg, A. Y. Guo, A. Deshpande, Y. Hong, Z.-X. Gong, A. V. Gorshkov, et al. (2020). Hierarchy of Linear Light Cones with Long-Range Interactions. *Physical Review X*, vol. 10, article 031009. <https://arxiv.org/abs/2001.11509>.
- [61] M. C. Tran, C.-F. Chen, A. Ehrenberg, A. Y. Guo, A. Deshpande, Y. Hong, Z.-X. Gong, A. V. Gorshkov, et al. (2020). Hierarchy of Linear Light Cones with Long-Range Interactions. *Physical Review X*, vol. 10, article 031009. <https://arxiv.org/abs/2001.11509>.
- [62] M. C. Tran, A. Y. Guo, C. L. Baldwin, A. Ehrenberg, A. V. Gorshkov, and A. Lucas (2021). The Lieb-Robinson Light Cone for Power-Law Interactions. *Physical Review Letters*, vol. 127, article 160401. <https://arxiv.org/abs/2103.15828>.
- [63] G. E. Volovik (2001). Superfluid Analogies of Cosmological Phenomena. *Physics Reports*, vol. 351, pp. 195–348. <https://arxiv.org/abs/gr-qc/0005091>.
- [64] G. E. Volovik (2001). Superfluid Analogies of Cosmological Phenomena. *Physics Reports*, vol. 351, pp. 195–348. <https://arxiv.org/abs/gr-qc/0005091>.
- [65] Z. Wang and K. R. A. Hazzard (2020). Tightening the Lieb-Robinson Bound in Locally Interacting Systems. *PRX Quantum*, vol. 1, article 010303. <https://arxiv.org/abs/1908.03997>.
- [66] C. Yin and A. Lucas (2022). Finite Speed of Quantum Information in Models of Interacting Bosons at Finite Density. *Physical Review X*, vol. 12, article 021039. <https://arxiv.org/abs/2106.09726>.
- [67] A. F. Young and P. Kim (2009). Quantum Interference and Klein Tunneling in Graphene Heterojunctions. *Nature Physics*, vol. 5, pp. 222–226. <https://arxiv.org/abs/0808.0855>.
- [68] J. Zhang, G. Pagano, P. W. Hess, A. Kypriandis, P. Becker, H. Kaplan, A. V. Gorshkov, Z.-X. Gong, et al. (2017). Observation of a Many-Body Dynamical Phase Transition with a 53-Qubit Quantum Simulator. *Nature*, vol. 551, pp. 601–604. <https://arxiv.org/abs/1708.01044>.