Synchronising noisy clocks in a gravitational field: locking, phase slips and a covariant rate connection

Ioannis Kokkinis*

Draft, September 28, 2026

Abstract

We apply known results on the synchronisation of coupled phase oscillators to clocks in a gravitational field and study them quantitatively. Identical-rate clocks with independent phase noise interact only at random coincidence events on the links of a lattice. Placed in a lapse gradient $\lambda(z) = 1 + gz$, plain Kuramoto-type coupling either locks the rates or slips. The locking condition for an open column of height L, $\rho \equiv \omega g L^2/(8K) \leq 1$, is the Strogatz–Mirollo / Ermentrout–Kopell chain condition specialised to a linear gradient. We state it as a gravitational locking length $\ell_{\text{lock}} = \sqrt{8K/(\omega g)}$. Below ℓ_{lock} the redshift is erased (frequency-gradient ratio 0.002 ± 0.002 at $\rho = 0.72$). Above it the column splits into slipping clusters (ratio 0.884 at $\rho = 7.2$), and a gauge-independent bond-stability measure falls from 0.849 (flat space) to 0.420 ± 0.004. With noise the threshold moves below 1, to $\rho_{\text{eff}} = 0.84 \pm 0.01$ for our run length, and lower still for longer runs. This is standard noise-activated slipping. A per-link integrated rate offset A_{ij} with $\dot{A}_{ij} = \omega(\lambda_j - \lambda_i)$ is a pure-gauge connection. With exact rates it restores the full redshift and flat-space coherence, and it also cancels random rate disorder, but only by a change of variables. The new question is what happens when each link must *measure* its rate offset. With noisy estimates the connection error random-walks and acquires holonomy. The mean redshift survives (ratio 1.000), but phase coherence has a finite lifetime and is lost once the accumulated rms error reaches about one radian. We motivate the problem by a decoherence bound on independent clock jitter in relational-time models, and discuss time-distribution networks.

1 Introduction

Networks of clocks are synchronised by exchanging phase information along links. When the clocks sit at different gravitational potentials, their proper rates differ by the lapse, $\lambda = 1 + \Phi/c^2$, and this difference is physical: it is measured by optical clocks separated by centimetres in height [1] and corrected for in satellite navigation [2]. A synchronisation rule that pulls neighbouring clocks towards each other's phase therefore competes with gravity. If it wins, it erases a real redshift; if it loses, the phases slip.

This paper studies that competition in a minimal stochastic model: identical noisy phase oscillators coupled only at random coincidence events, in the spirit of noisy lattice Kuramoto models [3, 4]. We state plainly what is known. The locking condition for a chain with open ends (the maximum partial sum of the frequency offsets must not exceed the coupling) is due to Strogatz and Mirollo [6] and to Ermentrout and Kopell [5]. Noise-activated phase slips below that threshold are Adler/Kramers physics [8, 9]. A link phase lag that is the difference of two site variables is pure gauge and can be removed by a change of variables [22]. The paper is therefore a gravitational-physics application and quantitative study of known synchronisation results. What is new is: (a) the gravitational framing, in which the chain condition becomes a locking length $\ell_{\text{lock}} = \sqrt{8K/(\omega g)}$: plain coupling erases the redshift below it and slips above it; (b) a systematic map of the noise-induced shift of the threshold in a 3D lattice (Fig. 2); (c) the case of a *measured* rate connection. When each link estimates its rate offset with noise, the connection acquires holonomy, and we quantify how phase coherence and redshift fidelity degrade (Sec. 5.1). The exact-connection results (full redshift, immunity to rate disorder) follow from the change of variables and are included as checks. We also reproduce, as checks of the model rather than as new results, the dimension dependence of phase order [11] and the effect of long-range links [12, 13].

The motivation comes from relational or emergent-spacetime models in which local clocks supply the quantum phase of matter (Sec. 2).

Table 1: Branch visibility for a Dirac walk with independent per-site tick jitter (research log, round 18).

μ	σ	width	t	V (sim.)	$e^{-\mu^2\sigma^2 t}$
0.1	0.5	20	200	0.604	0.607
0.1	0.5	60	200	0.603	0.607
0.1	1.0	30	100	0.367	0.368
0.1	1.0	30	200	0.130	0.135
0.2	0.5	30	200	0.135	0.135

There, independent clock jitter would decohere matter waves far too fast, so the clocks must be synchronised; and because the same clocks encode gravity as rate differences, the synchronisation must respect those differences. The results below do not depend on that motivation, and Sec. 8 discusses them in the context of engineered time-distribution networks.

2 Motivation: coherent local clocks

Suppose that the phase of a massive particle advances by one mass quantum μ per local tick, and that the tick interval has independent relative jitter σ at each place. Each path then accumulates independent phase noise, and the ensemble visibility between two branches decays as

$$V(t) = e^{-\mu^2\sigma^2 t} \quad \text{(per step)}, \qquad \Gamma = \left(\frac{mc^2}{\hbar}\right)^2 \sigma^2\, t_{\text{tick}} \tag{1}$$

in physical units. We tested Eq. (1) with a 1+1 dimensional Dirac quantum walk carrying site-local mass noise, two branches 900 sites apart, 60 realisations (Table 1). The visibility follows the prediction and does not depend on packet width. Coarse-graining does not help: counting N Poisson ticks gives $\sigma^2_{\text{eff}} = 1/N$ per interval $N t_{\text{tick}}$, so Γ is unchanged (Poisson clocks have time-diffusion constant t_{tick} at every scale).

Observed interference requires $\Gamma t_{\text{exp}} \lesssim 1$, hence, for Poisson ticks ($\sigma = 1$), $t_{\text{tick}} \lesssim \hbar^2/(m^2c^4 t_{\text{exp}})$. For caesium atom interferometry (133 u, $t_{\text{exp}} \sim 1$ s) this gives $t_{\text{tick}} \lesssim 2.8 \times 10^{-53}$ s, about nine orders of magnitude below the Planck time (5.4×10^{-44} s). For the molecules of Fein *et al.* [14] (mass $\approx$ 26 777 Da, coherence time $\tau = 7.5$ ms), it gives $t_{\text{tick}} \lesssim 9.3 \times 10^{-56}$ s, about twelve orders below the Planck time. These are order-of-magnitude estimates.

Key assumption. The bound assumes that the jitter experienced by the two interferometer arms is *independent*, i.e. the arms are separated by more than the correlation length of the clock noise. If both arms read one common clock, the noise is common-mode and cancels in the interference term. The bound therefore excludes independent random place-clocks, not clocks that are correlated over the size of the interferometer. This is the sense in which the clocks must be synchronised.

Related intrinsic- or gravitationally-induced decoherence mechanisms have a long history: Milburn's intrinsic decoherence [15], time treated as a stochastic variable [16], conformal fluctuations of spacetime [17], the decoherence floor from a stochastic gravitational-wave background [18], and decoherence by time dilation acting on internal degrees of freedom [19]. Eq. (1) is a simple instance of this class, specialised to tick jitter.

Two requirements follow for any model in which clocks carry the phase: (a) the clocks must be synchronised by some local interaction, which is only possible over long distances in sufficiently high dimension (Sec. 7); (b) the synchronisation must not erase the gravitational rate differences that the same clocks encode. Requirement (b) is the subject of Secs. 3–6. The phase difference picked up by matter at different heights is observed directly in neutron interferometry [20].

3 Model

Each site i of a graph carries a clock phase θ_i. In one time step,

$$\theta_i \leftarrow \theta_i + \omega\lambda_i, \tag{2}$$

$$\theta_i \leftarrow \theta_i + \varepsilon\, \xi_{ij} \sin(\theta_j - \theta_i - A_{ij}), \qquad \theta_j \leftarrow \theta_j - \varepsilon\, \xi_{ij} \sin(\theta_j - \theta_i - A_{ij}), \tag{3}$$

$$\theta_i \leftarrow \theta_i + \sqrt{2Q}\, \eta_i, \tag{4}$$

where $\xi_{ij} \in \{0, 1\}$ fires independently on each link with probability p per step (a "coincidence"), η_i is standard Gaussian noise of strength Q (we write d for the spatial dimension), and the kick is antisymmetric so that each firing conserves $\theta_i + \theta_j$. The mean coupling per link is $K = p\varepsilon$. On lattices the three axes are updated in sequence within a step. *Plain* coupling has $A_{ij} = 0$. *Covariant* coupling gives each link a variable that integrates the expected rate offset,

$$\dot{A}_{ij} = \omega(\lambda_j - \lambda_i), \qquad A_{ij}(0) = 0. \tag{5}$$

We use three measures, each averaged over the second half of the run. (1) The global order

$m = |\langle e^{i\theta_i}\rangle|$, used in flat space and with disorder; for covariant coupling with disorder, θ_i is replaced by $\theta_i - \omega\lambda_i t$. (2) The *proper-time fidelity* $F = |\langle e^{i(\theta_i - \omega\lambda_i t)}\rangle|$ in the gravity runs. F measures agreement with the phase each clock would have at its proper rate. It is *not* a coherence measure for plain coupling, because a coherent but rate-locked column also has low F. (3) The gauge-independent *bond stability* S. For each nearest-neighbour bond the phase difference $\theta_j - \theta_i$ is unwrapped step by step, a straight line (its mean beat frequency) is subtracted, and $S_b = |\langle e^{i\,\mathrm{residual}}\rangle_t|$; S is the mean over bonds. S uses only the clock phases, is unchanged when any site-dependent term linear in t is added to θ_i (a change of rates or of rate frame), and needs no reference to the lapse or the connection. Locked or rate-following bonds give S at its flat-space value, while 2π slips and wandering offsets lower it. Statistics are mean $\pm$ standard error over independent seeds where stated; the other entries are single realisations. The parameters are listed with each result.

4 Plain coupling in a lapse gradient

4.1 Locking criterion

Consider a column of L layers at heights $z_k = k - (L+1)/2$, $k = 1, \ldots, L$, with $\lambda_k = 1 + g z_k$, open at both ends and uniform in the transverse directions. Averaging over firings and dropping the noise, the mean-field dynamics of layer k is

$$\begin{aligned} \dot{\theta}_k &= \omega\lambda_k + J_{k+\frac{1}{2}} - J_{k-\frac{1}{2}}, \\ J_{k+\frac{1}{2}} &= K \sin(\theta_{k+1} - \theta_k), \end{aligned} \tag{6}$$

with $J_{1/2} = J_{L+1/2} = 0$ at the open ends. Transverse bonds carry no current in a uniform state and drop out. A locked state has $\dot{\theta}_k = \Omega$ for all k. Summing over k gives $\Omega = \omega\langle\lambda\rangle = \omega$. Summing up to layer k gives the bond currents

$$J_{k+\frac{1}{2}} = \sum_{i=1}^{k} (\Omega - \omega\lambda_i) = -\omega g \sum_{i=1}^{k} z_i. \tag{7}$$

$|J|$ is largest at the middle bond, $k = L/2$, where $\sum_{i=1}^{L/2} z_i = -L^2/8$ exactly for even L (in the continuum, $J(z) = \frac{\omega g}{2}(L^2/4 - z^2)$, with the same maximum). A locked state requires $|J| \le K$ on every bond, because $|\sin| \le 1$. Hence

$$\rho \equiv \frac{\omega g L^2}{8K} \le 1 \quad \Longleftrightarrow \quad L \le \ell_{\mathrm{lock}} = \sqrt{\frac{8K}{\omega g}}. \tag{8}$$

When $\rho < 1$, the solution with every bond phase difference in $(-\pi/2, \pi/2)$ is linearly stable: its Jacobian is minus a weighted graph Laplacian with positive weights $K\cos(\theta_{k+1} - \theta_k)$. For $\rho > 1$ no locked state exists and the column breaks into clusters separated by slipping bonds. Equation (8) is the known chain condition, $\max_k |\sum_{i\le k}(\omega_i - \bar{\omega})| \le K$ [6, 5], evaluated for a linear gradient. The cluster breakup is the frequency-plateau phenomenon of Ref. [5]. Only the gravitational reading, a locking length ℓ_{lock}, is ours. In the locked state all layers tick at the common rate Ω, so the redshift is erased.

4.2 Simulations

We simulated a $12 \times 12 \times 24$ cubic lattice, periodic in x, y and open in z, with $p = 0.5$, $\varepsilon = 0.4$ ($K = 0.2$), $Q = 0.3K$, $\omega = 1$, 6000 steps, and an ordered start. Layer frequencies $\bar{\Omega}(z)$ were measured over the second half of the run, and the *frequency-gradient ratio* is the fitted slope of $\bar{\Omega}(z)$ divided by ωg (1 = full redshift, 0 = locked). A rerun of the original script (research log, round 20) with the same seed reproduces the logged values exactly: ratios -0.002 and 0.883 (plain) and 1.001, 1.000 (covariant). Table 2 gives the mean over four further seeds (script `covsim.py`); Fig. 1 shows the layer profiles.

For $L = 24$, $K = 0.2$, $\omega = 1$, Eq. (8) gives $\rho = 0.72$ at $g = 0.002$ ($\ell_{\mathrm{lock}} = 28.3 > 24$) and $\rho = 7.2$ at $g = 0.02$ ($\ell_{\mathrm{lock}} = 8.9 < 24$). The simulations agree. At $g = 0.002$ the ratio is 0.002 ± 0.002: the column is locked and the redshift is gone. The column stays coherent: $S = 0.784 \pm 0.001$, somewhat below the flat-space 0.849 ± 0.001. We did not separate the cause of this reduction; softer bonds under tension ($K\cos\Delta\theta$) are one candidate. At $g = 0.02$ the ratio is 0.884. The interior layers follow the lapse through slipping bonds, and a cluster of about four layers at each end locks to a common frequency (Fig. 1). This is a genuine loss of coherence: $S = 0.420 \pm 0.004$. The fidelity $F \approx 0.10$ is low in *both* plain cases and does not distinguish them. At $g = 0.002$ it is low because the rates are locked, at $g = 0.02$ because of slips.

To test Eq. (8) beyond these two points, we scanned ρ (Fig. 2, script `lockscan.py`). In a noise-free column with every link firing every step ($\varepsilon = K = 0.2$), heights $L = 12$, 24 and 48 all stay locked at $\rho = 0.98$ (ratio ≤ 0.002) and all slip at $\rho = 1.02$ (ratios 0.098, 0.059, 0.047). This confirms $\rho_c = 1$ and the L^2 scaling. In the noisy 3D

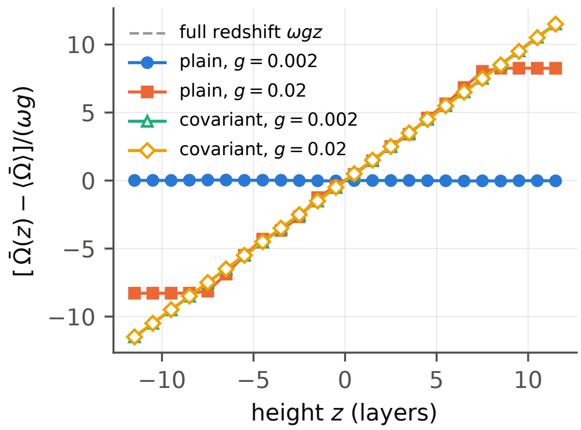

Figure 1: Layer frequency relative to the column mean, in units of ωg, for a $12 \times 12 \times 24$ lattice in the lapse $\lambda = 1 + gz$. Dashed: full redshift. Plain coupling locks at $g = 0.002$ ($\rho = 0.72$) and forms end clusters with a slipping interior at $g = 0.02$ ($\rho = 7.2$). Covariant coupling (open symbols; the two g values overlap) follows the redshift exactly.

Table 2: Lapse-gradient runs ($12 \times 12 \times 24$, $K = 0.2$, $Q = 0.3K$, 6000 steps), mean ± s.e. over 4 seeds. Ratio: fitted frequency gradient$/(\omega g)$. F: proper-time fidelity. S: gauge-independent bond stability. Entries without an error have s.e. < 0.0005.

	g	ρ	ratio	F	S
flat	0	0	–	0.916	0.849(1)
plain	0.002	0.72	0.002(2)	0.103(1)	0.784(1)
plain	0.02	7.2	0.884	0.102(2)	0.420(4)
cov.	0.002	–	1.001(1)	0.916	0.849(1)
cov.	0.02	–	1.000	0.916	0.849(1)

model we scanned $\rho = 0.5$–1.0 with six seeds per point. The ratio is 0.001 ± 0.001 for $\rho \leq 0.75$, 0.011 ± 0.009 at 0.80 (one seed of six slipped), 0.103 ± 0.006 at 0.85 and 0.223 ± 0.002 at 1.0. Take the effective threshold as the first grid value where the ratio exceeds 0.01. It is $\rho_{\text{eff}} = 0.842 \pm 0.008$ (s.e., six seeds; grid spacing 0.05, so it lies between 0.80 and 0.85). The shift depends on observation time, as expected for activated escape [9, 8]. At $\rho = 0.80$, runs four times longer (24 000 steps) slip in all four seeds tested (ratios 0.015–0.028), against one of four in 6000 steps. ρ_{eff} therefore characterises a finite observation window, not a sharp transition.

Plain synchronisation is thus incompatible with a lapse gradient at every scale. Below ℓ_{lock} it removes a physical redshift; above it, it replaces the redshift by random phase slips, i.e. by the kind of uncorrelated phase noise that Sec. 2 excludes. The slips show up as a fall in the gauge-

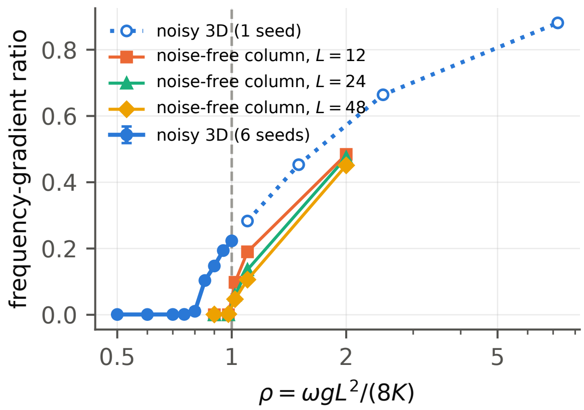

Figure 2: Frequency-gradient ratio versus $\rho = \omega g L^2/(8K)$ for plain coupling. Noise-free columns of three heights collapse onto a threshold at $\rho = 1$. The noisy 3D lattice (filled: six seeds, mean ± s.e.; open: single seed) slips earlier, from $\rho \approx 0.84$ for 6000-step runs.

independent S, not only in F.

5 Covariant coupling

With Eq. (5), the link compares hands after allowing for the offset the two clocks are expected to accumulate. Since $A_{ij}(t) = \chi_j - \chi_i$ with $\chi_i = \omega \int_0^t \lambda_i \, dt'$, the connection is *pure gauge*: its holonomy around every closed loop vanishes. As Torres-Hugas *et al.* show for general phase-lag networks [22], such a connection can be removed at the nodes. The substitution $\theta_i = \chi_i + \psi_i$ maps the dynamics exactly onto the flat-space model for ψ_i. Exact redshift and flat-space coherence therefore follow from a change of variables $\theta \to \theta - \omega \int \lambda \, dt$, given that each link knows the *exact* rate difference. The simulations confirm the implementation: ratios 1.001 ± 0.001 and 1.000, with F and S equal to their flat-space values (Table 2). They add no independent evidence.

A_{ij} enters the coupling as a lattice gauge field enters a hopping term, with the lapse as its time component, $A_0 = \omega\lambda$. The Sakaguchi–Kuramoto phase lag [21] is the constant, uniform special case. The gravitational connection has no Aharonov–Bohm-type loop phase. Loop phases (and with them frustration) appear only when the connection is *not* the difference of site variables. This is exactly what happens when it is measured.

5.1 Measured connection

In practice each link must estimate $\lambda_j - \lambda_i$. We drive the integrator by a noisy estimate in two ways, (i) white estimate noise and (ii) a static calibration error, at $g = 0.02$ ($\rho = 7.2$, where plain coupling slips), with three seeds per point:

$$\text{(i)} \quad A_{ij} \mathrel{+}= \omega(\lambda_j - \lambda_i) + \sigma_{\text{est}}\zeta_{ij}(t),$$
$$\text{(ii)} \quad A_{ij} \mathrel{+}= \omega(\lambda_j - \lambda_i + \delta_{ij}),$$

with $\zeta_{ij}(t)$ white Gaussian noise, independent per link and step, and $\delta_{ij} \sim \mathcal{N}(0, \sigma_{\text{cal}}^2)$ fixed per link. The integrator has no restoring term, so its error grows without bound: as $\sigma_{\text{est}}\sqrt{t}$ (random walk) or as $\sigma_{\text{cal}}t$. Independent errors on the links of a plaquette have nonzero sum, so the connection acquires holonomy, i.e. frustration [22], in addition to a gauge part that shifts the clocks away from proper time.

The results (Fig. 3; 6000 steps, errors are s.e. ≤ 0.003) are:

- The mean redshift is always preserved: the ratio is 1.000 ± 0.000 in every case.
- Estimate noise, with accumulated rms error $\sigma_{\text{est}}\sqrt{T}$ at the end of the run. At $\sigma_{\text{est}} = 10^{-3}$ (0.08 rad) there is no measurable effect ($F = 0.915$, $S = 0.848$). At 3×10^{-3} (0.23 rad): $F = 0.907$, $S = 0.839$. At 10^{-2} (0.77 rad): $F = 0.825$, $S = 0.758$. At 3×10^{-2} (2.3 rad): $F = 0.019$, $S = 0.252$, i.e. coherence is lost.
- Calibration error, with accumulated error $\sigma_{\text{cal}}T$. At $\sigma_{\text{cal}} = 10^{-4}$ (0.6 rad): $F = 0.875$, $S = 0.805$. At 10^{-3} (6 rad): $F = 0.013$, $S = 0.286$. At 10^{-2}: $F = 0.016$, $S = 0.134$.

Coherence therefore has a finite lifetime, set by when the accumulated connection error reaches about one radian: $t_c \sim \sigma_{\text{est}}^{-2}$ or σ_{cal}^{-1}. The degradation is not a function of the accumulated error alone. At equal $\sigma_{\text{est}}^2 T$, a four times longer run is worse ($S = 0.18$ at $\sigma_{\text{est}} = 0.01$, $T = 24\,000$, against $S = 0.32$ at $\sigma_{\text{est}} = 0.02$, $T = 6000$; two seeds each), and the same holds at equal $\sigma_{\text{cal}}T$ ($S = 0.18$ against 0.32), consistent with slips that accumulate. The open-loop integrator of Eq. (5) is therefore not viable on its own. A practical scheme needs bounded-error rate estimation, e.g. estimates referenced to phase comparisons over long baselines, as in a proportional–integral servo. We have not simulated such a scheme, and whether one can keep both the redshift and flat-space coherence indefinitely remains open.

6 Quenched rate disorder

The same connection absorbs *any* static rate differences, not only gravitational ones. We drew $\lambda_i = 1 + \sigma\xi_i$ with Gaussian ξ_i on periodic L^3 lattices ($p = 0.5$, $\varepsilon = 0.4$, $Q = 0.3K$, 4000 steps). With plain coupling, the order decreases with L, and faster for larger σ. With covariant coupling it is independent of L (Table 3, Fig. 4). A rerun of the original single-seed script (research log, round 22) reproduced its values exactly. Table 3 gives the mean $\pm$ s.e. over four further seeds. As in Sec. 5, the covariant result is a change of variables that assumes exactly known rates.

The decline under plain coupling is consistent with the known absence of phase order below four dimensions for locally coupled oscillators with random frequencies [7] (see [6] for chains, and [10] for the separate question of frequency entrainment). Our sizes are too small to establish an asymptotic law. With exactly known rates the disorder is pure gauge and is removed. With measured rates the considerations of Sec. 5.1 apply.

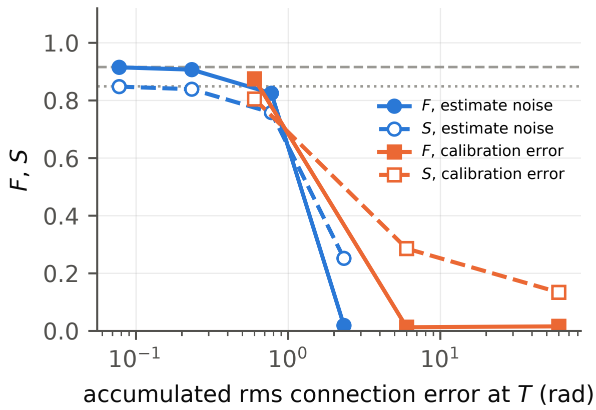

Figure 3: Measured connection at $g = 0.02$: proper-time fidelity F (filled) and bond stability S (open) versus the accumulated rms connection error at the end of a 6000-step run, for white estimate noise (circles) and static calibration error (squares). Three seeds per point. Dashed/dotted: exact-connection values of F and S.

Table 3: Order m on L^3 lattices with random rates $\lambda_i = 1 + \sigma\xi_i$ ($Q = 0.3K$, 4000 steps); mean $\pm$ s.e. over 4 seeds, digits in parentheses. Covariant entries have s.e. < 0.0005.

σ	coupling	$L = 8$	16	24
0.08	plain	0.917(1)	0.908(1)	0.895(2)
0.08	covariant	0.922	0.919	0.919
0.15	plain	0.901(2)	0.875(6)	0.830(8)
0.15	covariant	0.922	0.919	0.919

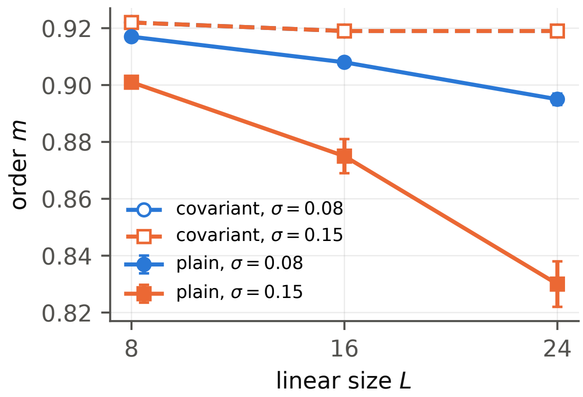

Figure 4: Order versus linear size with quenched random rates. Plain coupling (filled) loses order as L grows. Covariant coupling (open; the two σ values coincide) does not. Four seeds, mean $\pm$ s.e.

7 Dimension and long links

For completeness, we record the dimension dependence of the flat-space model with identical rates. These results are consistent with known theory and are not claimed as new. On hypercubic lattices ($p = 0.5$, $\varepsilon = 0.4$, $Q/K = 0.7$, ordered start), the order decays towards zero in 1D ($m = 0.187 \to 0.023$ for $L = 64 \to 4096$), decays slowly in 2D ($0.689 \to 0.570$ for $L = 16 \to 128$, roughly $L^{-0.09}$), and is size-independent in 3D ($0.818 \to 0.810$ for $L = 8 \to 24$) (Fig. 5a). With identical rates the model is a noisy XY dynamics, and this is the pattern the Mermin–Wagner theorem implies: no long-range phase order in $d \leq 2$ at finite noise Q [11].

The same holds on Poisson-random graphs where each point links to its six nearest neighbours (mean degree ≈ 7 in both 2D and 3D; $\varepsilon = 0.12$, $Q/K = 0.7$). In 2D the order falls from 0.869 to 0.799 as N goes from 256 to 16384; in 3D it stays at 0.885–0.896 for $N = 512$–21952. The static phase-difference variance of the Gaussian (spin-wave) approximation behaves the same way: it grows linearly with distance in 1D ($1.0 \to 1008$ for $r = 1 \to 1024$), logarithmically in 2D ($0.50 \to 1.62$ for $r = 1 \to 32$), and saturates in 3D ($0.33 \to 0.50$).

Adding a small density of random long-range links to 2D random graphs ($Q/K = 1$) restores size-independent order (Fig. 5b). With no long links, $m = 0.775$, 0.702, 0.703 at $N = 1024$, 4096, 16384. With long links numbering 2% of N, $m = 0.775$, 0.774, 0.766. With 10%, $m = 0.826$, 0.817, 0.814. This matches known results on synchronisation in small-world networks [12] and on the growth of estimation error with graph dimension in relative-measurement networks [13].

8 Time-distribution networks

The practical content is modest, and we state it plainly. Precision time distribution (IEEE 1588/PTP [23], White Rabbit [24], fibre-linked optical clock networks [25], GNSS [2]) already separates *rate* from *phase*: servos estimate and compensate frequency offset and drift, and GNSS satellite clocks are offset in advance for their relativistic rate difference. Covariant coupling is a continuous, fully decentralised version of this, with one integrator per link and no master clock. Our results make two points about it. First, rate compensation is not optional: naive neighbour averaging either suppresses true rate differences (below ℓ_{lock}) or produces slips (above it), and Eq. (8) gives the crossover length for a linear gradient. Second, with *exact* rate compensation, coherence becomes independent of static rate structure, whether systematic or random. With *measured* rates, however, a free-running per-link integrator accumulates error and loses phase within a time $\sim \sigma_{\text{est}}^{-2}$ (Sec. 5.1). This is why practical servos close the loop on phase. For mesh design, nearest-neighbour coupling on planar networks cannot hold one global phase at large size, while 3D meshes or a few percent of long links can, a known result [12, 13] that serves as a simple design rule. We make no claim of a new technology.

9 Discussion

The results are of three kinds. (i) Known, restated for gravity: plain coupling in a lapse gradient locks over lengths below $\ell_{\text{lock}} = \sqrt{8K/(\omega g)}$, the chain condition of Refs. [6, 5], and slips above it. Noise-free columns confirm $\rho_c = 1$ and the L^2 scaling. (ii) Quantified: noise moves the effective threshold to $\rho_{\text{eff}} = 0.84 \pm 0.01$ for 3000-step measurement windows, and further down for longer windows (Kramers/Adler). The exact connection is pure gauge, so exact redshift and disorder immunity follow from a change of variables. (iii) New and open: a connection built from *measured* rates acquires holonomy and drifts. It keeps the mean redshift but loses phase coherence once its accumulated error reaches about one radian.

The limitations are as follows. The system sizes are small and the seed counts modest. The threshold map uses a 0.05 grid and one run length.

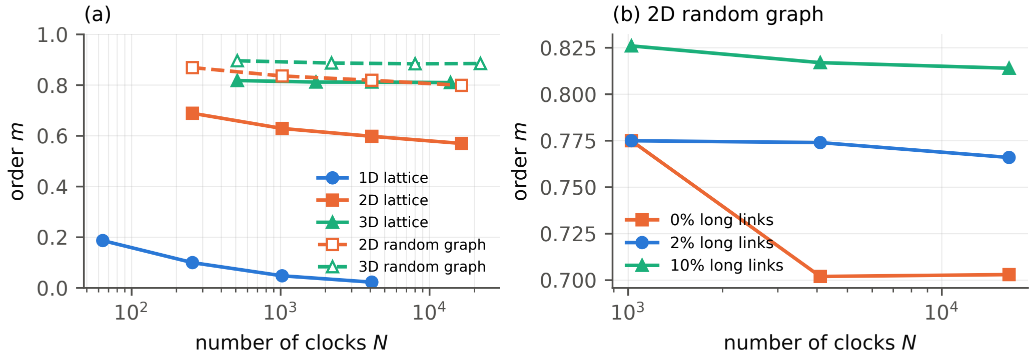

Figure 5: (a) Flat-space order versus number of clocks on hypercubic lattices (filled) and 6-nearest-neighbour random graphs (open). (b) 2D random graphs with a fraction of random long-range links (percent of N).

The dimension and long-link data are single realisations from the research log. The field is static and linear. We did not simulate a closed-loop (bounded-error) connection. For the emergent-spacetime motivation, the coupling rule (5) is a hypothesis, and Sec. 5.1 shows that its physical realisation would need an error-correcting mechanism.

Acknowledgements

Computations, analysis and drafting were carried out with AI assistance (Claude, Anthropic). All numerical results come from the author's research log or from re-runs of the accompanying scripts.

References

[1] C. W. Chou, D. B. Hume, T. Rosenband, D. J. Wineland, "Optical clocks and relativity," Science **329**, 1630–1633 (2010).

[2] N. Ashby, Living Rev. Relativ. **6**, 1 (2003).

[3] Y. Kuramoto, *Chemical Oscillations, Waves, and Turbulence* (Springer, Berlin, 1984).

[4] J. A. Acebrón, L. L. Bonilla, C. J. Pérez Vicente, F. Ritort, R. Spigler, Rev. Mod. Phys. **77**, 137 (2005).

[5] G. B. Ermentrout, N. Kopell, SIAM J. Math. Anal. **15**, 215 (1984), doi:10.1137/0515019.

[6] S. H. Strogatz, R. E. Mirollo, Physica D **31**, 143–168 (1988).

[7] H. Hong, H. Park, M. Y. Choi, Phys. Rev. E **72**, 036217 (2005).

[8] R. Adler, Proc. IRE **34**, 351 (1946).

[9] H. A. Kramers, Physica **7**, 284 (1940).

[10] H. Hong, H. Chaté, H. Park, L.-H. Tang, Phys. Rev. Lett. **99**, 184101 (2007).

[11] N. D. Mermin, H. Wagner, Phys. Rev. Lett. **17**, 1133 (1966).

[12] H. Hong, M. Y. Choi, B. J. Kim, Phys. Rev. E **65**, 026139 (2002).

[13] P. Barooah, J. P. Hespanha, "Error scaling laws for linear optimal estimation from relative measurements," IEEE Trans. Inf. Theory **55**, 5661–5673 (2009).

[14] Y. Y. Fein *et al.*, Nat. Phys. **15**, 1242 (2019).

[15] G. J. Milburn, Phys. Rev. A **44**, 5401 (1991).

[16] R. Bonifacio, Nuovo Cimento B **114**, 473 (1999).

[17] W. L. Power, I. C. Percival, Proc. R. Soc. A **456**, 955 (2000).

[18] B. Lamine, R. Hervé, A. Lambrecht, S. Reynaud, Phys. Rev. Lett. **96**, 050405 (2006).

[19] I. Pikovski, M. Zych, F. Costa, Č. Brukner, "Universal decoherence due to gravitational time dilation," Nat. Phys. **11**, 668–672 (2015).

[20] R. Colella, A. W. Overhauser, S. A. Werner, "Observation of gravitationally induced quantum interference," Phys. Rev. Lett. **34**, 1472–1474 (1975).

[21] H. Sakaguchi, Y. Kuramoto, Prog. Theor. Phys. **76**, 576 (1986).

[22] L. Torres-Hugas, J. Duch, S. Gómez, A. Arenas, "Cycle holonomy induces higher-order constraints and controls remote synchronization transitions via twisted Laplacian spectra," arXiv:2604.19682 (2026).

[23] IEEE Std 1588-2019, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems* (2020).

[24] M. Lipiński, T. Włostowski, J. Serrano, P. Alvarez, in *Proc. IEEE ISPCS* (2011).

[25] C. Lisdat *et al.*, Nat. Commun. **7**, 12443 (2016).