Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Multi-Walker Forward-Jump Random Walks on the Cycle with Laziness and Teleportation: Meeting, Collision, and Coalescence

Nassamon Bootwisas
Rajamangala University of Technology Suvarnabhumi
Pasin Marupanthorn
quantfilab@gmail.com
Maejo University
Rachanai Kaikeaw
Dhonburi Rajabhat University

Research Article

Keywords: Forward-jump random walk, random walks on cycles, multi-walker Markov chains, meeting times, collision and coalescence

Posted Date: June 1st, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9833736/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Multi-Walker Forward-Jump Random Walks on the Cycle with Laziness and Teleportation: Meeting, Collision, and Coalescence

Nassamon Bootwisas[1], Pasin Marupanthorn[2*], Rachanai Kaikeaw[3]

[1]Division of Computational Science, Rajamangala University of Technology Suvarnabhumi, Ayutthaya, 13000, Thailand.
[2]Department of Mathematics, Faculty of Science, Maejo University, Chiang Mai, 50290, Thailand.
[3]Faculty of Science and Technology, Dhonburi Rajabhat University, Bangkok, 10600, Thailand.

*Corresponding author(s). E-mail(s): quantfilab@gmail.com;

Abstract

This paper develops a quantitative theory for encounter-driven phenomena in systems of multiple walkers on a finite cycle. The underlying single-particle dynamics is a translation-invariant Markov chain that combines forward jumps, laziness, and uniform teleportation (FJRW–LT). By lifting the process to the joint configuration space, pairwise meetings, multi-particle collisions, and coalescence events are formulated as first-arrival problems for diagonal subsets of the product group. The diagonal-shift symmetry of the model shows that the essential dependence is on relative separations rather than absolute locations. For independent walkers, the product transition kernel remains circulant on the finite abelian group, which permits Fourier diagonalization and yields an explicit finite Fourier-series representation of the time-$\boldsymbol{n}$ arriving distribution. In the two-walker case, a one-dimensional displacement chain gives closed-form Fourier-sum formulas for the mean and variance of the meeting time for arbitrary initial offsets. For three or more walkers, the paper derives computable bounds and well-mixed occupancy approximations for first-collision and phase-wise coalescence times. Teleportation is shown to act as a Doeblin-type refresh mechanism, leading to explicit total-variation contraction and geometric or exponential tail estimates that distinguish teleport-dominant, diffusion-dominant, and crossover regimes. The framework is further extended to correlated common-shock couplings and to congestion-motivated interaction rules, including blocking, bouncing, and pushing. Monte

Carlo experiments support the theoretical formulas, bounds, and approximations across representative parameter regimes.

Keywords: Forward-jump random walk , random walks on cycles , multi-walker Markov chains , meeting times , collision and coalescence

1 Introduction

Random walks on graphs constitute a fundamental class of Markov chains used in diffusion modeling, randomized algorithms, network analysis, and stochastic process theory. Their behavior is commonly characterized through hitting times, which describe first arrival to prescribed sets, and mixing times, which measure convergence to stationarity. Standard treatments develop these notions together with spectral methods, coupling arguments, and electrical-network interpretations; see, for example, Lovász [1], Levin and Peres [2], and Doyle and Snell [3]. The cycle graph C_n is a particularly important benchmark because its group structure often permits closed-form analysis and because it captures ring-like geometries arising in communication, scheduling, and transport systems; explicit cycle formulas are developed, for instance, in [4].

Recent studies have extended classical nearest-neighbor cycle walks to models with longer-range directed motion. Kaikeaw and Marupanthorn [5] introduced a forward-jump walk on C_n in which each step advances in a fixed direction by at most m vertices, and derived exact distributional and moment formulas for hitting times using combinatorial methods. Two standard modifications are often introduced to improve ergodicity and analytical tractability. The first is *laziness*, which inserts a self-loop with probability $p > 0$ and thereby removes periodicity in many discrete-time chains [2]. The second is *teleportation*, or reset/restart, in which the walker jumps to a location sampled from a prescribed distribution v. Teleportation is central to PageRank, where the linear system

$$(I - \alpha P)x = (1 - \alpha)v, \qquad 0 < \alpha < 1,$$

incorporates a refresh component that ensures uniqueness and improves numerical robustness [6, 7]. Localized choices of v lead to personalized ranking and random-walk-with-restart models in network analysis [8–10]. For cycle graphs, Bootwisas, Marupanthorn, and Kaikeaw [11] studied a forward-jump walk with laziness and uniform teleportation (FJRW–LT). Translation invariance gives a circulant transition kernel, allowing diagonalization by the discrete Fourier basis and yielding explicit spectral, mixing, and hitting-time formulas, including geometric convergence and exponential hitting-time tails induced by teleportation.

For systems with several walkers on the same graph, the central first-passage observables are encounter-based. The *meeting time* records the first co-location of a designated pair of walkers, the *collision time* records the first meeting among any pair in a population, and the *coalescence time* records the time until a single cluster remains when walkers merge upon meeting. These quantities are closely related to single-walk mixing and hitting properties. Oliveira [12], for example, identifies regimes in which coalescence is controlled by the two-walker mean meeting time $m(G)$ and in

which the first meeting among k walkers occurs on the scale $m(G)/\binom{k}{2}$. Bounds for lazy chains in terms of relaxation and hitting parameters are obtained in [13], while near-tight comparisons between meeting and coalescence times across graph classes are developed in [14]. Multi-walker dynamics also accelerate exploration and cover-time behavior [15]. Conversely, discrete-time periodicity may prevent meetings from certain initial states; product-chain criteria and the stabilizing role of aperiodicity modifications are discussed in [16].

The motivation for studying multi-walker encounters is not limited to abstract probability. In distributed systems, agents often exchange information only when they interact, as in randomized gossip, token-based aggregation, and encounter-driven coordination [17–19]. In well-mixed networks, simultaneous mobility can generate substantial speedups, as in models of talkative particles [20]; mobility can also accelerate gossip-type convergence even when only part of the population moves [21]. Encounter times therefore provide natural latency and risk metrics for robot swarms, warehouse fleets, traffic deconfliction, epidemiological contact processes, distributed update propagation, and other first-passage systems; see also [22–24].

Although single-walker forward-jump dynamics on cycles are analytically tractable, considerably less is explicit for multi-walker forward-jump systems, where the principal observables are diagonal hitting times. General multi-particle theory often emphasizes rapidly mixing graphs, whereas cycles are highly structured and may exhibit slow diffusion, strong correlations, and synchronization effects unless refresh mechanisms such as laziness or teleportation are present. The objective of this paper is to provide a tractable quantitative framework for multi-walker forward-jump dynamics on a finite cycle with laziness and teleportation. Building on the FJRW–LT kernel of [11], the model is lifted to L walkers on the product group $(\mathbb{Z}_d)^L$. Translation invariance is then used to obtain a full Fourier diagonalization and explicit arriving distributions [25]. Relative-coordinate reductions yield closed-form Fourier-series expressions for the mean and variance of the two-walker meeting time, while bounds and well-mixed approximations are developed for first collision and phase-wise coalescence in larger populations. A recurring theme is that teleportation supplies a Doeblin-type refresh component, giving total-variation contraction and exponential tail control for collision and coalescence observables [2, 26]. The framework also accommodates common-shock couplings and interaction rules such as blocking, bouncing, and pushing, which model congestion and contention in encounter-driven systems.

The remainder of the paper is organized as follows. Section 2 defines the single-walker FJRW–LT kernel, introduces L-walker couplings on $(\mathbb{Z}_d)^L$, and formalizes meeting, collision, and coalescence targets as diagonal hitting sets. Section 3 develops the independent-walker baseline, including Fourier diagonalization of the product chain, explicit arriving distributions, and a displacement-chain reduction for the two-walker meeting problem. Section 4 isolates the effect of teleportation through exact mixture identities and tensorized bounds, leading to collision-probability lower bounds, geometric-type tails, and regime classifications. Section 5 introduces interacting kernels and formulates perturbative and coupling-based principles for transferring baseline predictions when conflicts are rare. Section 6 presents numerical validation across representative values of (d, L) and the model parameters. Section 7 outlines

application contexts in ring-like distributed systems, and Section 8 summarizes the main contributions and future directions.

2 Model Setup and Notation

This section recalls the single-walker Forward-Jump Random Walk with Laziness and Teleportation (FJRW–LT) on the cycle, following [11], and then extends the construction to an L-walker process on the corresponding product space. The product-space formulation is useful because meetings, collisions, and coalescence events can be represented as hitting problems for diagonal subsets of $(\mathbb{Z}_d)^L$, namely sets in which two or more coordinates coincide. Two coupling classes are considered: the baseline product coupling, in which updates are independent across walkers and time, and common-shock couplings, in which selected sources of randomness are shared across walkers, such as a common teleport clock or a common increment. These variants make it possible to distinguish independent encounter formation from encounters generated by correlated restarts or partially synchronous motion.

2.1 State space and single-walker kernel

Fix an integer $d \geq 2$ and write $\mathbb{Z}_d = \{0, 1, \ldots, d-1\}$ for the cyclic group under addition modulo d. A single walker occupies a state $R_n \in \mathbb{Z}_d$ at time n. The FJRW–LT dynamics is specified by a step distribution μ on $\mathbb{Z}_d$, written $\mu(x) = p_x$ with $\sum_{x\in\mathbb{Z}_d} p_x = 1$, together with a laziness parameter $s \in [0, 1]$ and a teleportation parameter $\alpha \in [0, 1)$.

Conditioned on $R_{n-1} = r$, the walker updates as follows. With probability α, it teleports to a fresh location by drawing $U_n \sim \mathrm{Unif}(\mathbb{Z}_d)$ and setting $R_n = U_n$. With the remaining probability $1 - \alpha$, no teleportation occurs and the walker instead performs a lazy increment step: it stays at r with probability s, and with probability $1 - s$ it draws an increment $X_n \sim \mu$ and moves to $R_n = r + X_n \pmod d$. Unless stated otherwise, all random choices are independent over time.

Let P be the increment convolution kernel $P(r, u) = \mu(u - r)$, let I denote the identity kernel, and let J be the uniform kernel $J(r, u) = 1/d$. The one-step transition kernel is

$$P^{(\alpha,s)} = \alpha J + (1 - \alpha)\big(sI + (1 - s)P\big).$$

This kernel is translation-invariant (circulant) on $\mathbb{Z}_d$, hence diagonalizable in the discrete Fourier basis [2]. Writing $\omega = \exp(2\pi i/d)$ and $\phi_k(r) = \omega^{kr}$ for $k = 0, 1, \ldots, d-1$, define the Fourier transform of μ by

$$\lambda_k := \sum_{x\in\mathbb{Z}_d} p_x\, \omega^{kx} = \mathbb{E}\big[\omega^{kX_1}\big], \qquad k = 0, 1, \ldots, d-1,$$

so $\lambda_0 = 1$. The lazy-increment eigenvalues are

$$\widetilde{\lambda}_k := s + (1 - s)\lambda_k,$$

and the corresponding FJRW–LT eigenvalues are

$$\Lambda_0 := 1, \qquad \Lambda_k := (1-\alpha)\widetilde{\lambda}_k = (1-\alpha)\big(s + (1-s)\lambda_k\big), \qquad k = 1, \dots, d-1. \quad (1)$$

Most quantitative statements below will be expressed in terms of $\{\Lambda_k\}_{k\neq 0}$.

Since $P^{(\alpha,s)}$ is doubly stochastic, the uniform distribution $\pi(r) = 1/d$ is stationary. If $\alpha > 0$, then the chain is automatically irreducible and aperiodic; if $\alpha = 0$, irreducibility depends on the support of μ, while aperiodicity holds whenever $s > 0$.

2.2 Assumptions and well-posedness

Many statements below rely on basic ergodicity (irreducibility and aperiodicity) of the underlying Markov chains. The following simple sufficient conditions and the specific non-degeneracy assumption used in the two-walker displacement-chain results are recorded for later reference.

Assumption 1 (Ergodicity and non-degeneracy) *Throughout the paper, unless stated otherwise, the following assumptions are imposed:*

(A1) ***Single-walker ergodicity.*** *Either $\alpha > 0$, in which case $P^{(\alpha,s)}$ is automatically irreducible and aperiodic on $\mathbb{Z}_d$ (Doeblin minorization), or $\alpha = 0$ and the increment walk is irreducible, e.g.* $\gcd(\mathrm{supp}(\mu), d) = 1$*, together with $s > 0$ to ensure aperiodicity.*

(A2) ***Displacement-chain irreducibility (two-walker theory).*** *For the displacement chain (D_n) in Section 3.2, we assume it is irreducible on $\mathbb{Z}_d$, equivalently that*

$$\Gamma_k \neq 1 \quad \text{for all } k \in \{1, \dots, d-1\}, \qquad \Gamma_k := |\Lambda_k|^2.$$

This excludes degenerate cases where relative motion is confined to a proper subgroup (e.g. perfectly synchronized deterministic increments with $\alpha = 0$).

Remark 1 When $\alpha > 0$, condition (A2) automatically holds and (D_n) is aperiodic because $\mathbb{P}(D_n = D_{n-1}) \geq \alpha^2/d > 0$. When $\alpha = 0$, (A2) can be checked via the step law of the difference increment $X^{(1)} - X^{(2)}$ and standard gcd criteria on $\mathbb{Z}_d$.

2.3 Multi-walker processes

Fix $L \geq 2$ and consider L walkers on the same cycle. The configuration at time n is

$$\mathbf{R}_n = \big(R_n^{(1)}, R_n^{(2)}, \dots, R_n^{(L)}\big) \in (\mathbb{Z}_d)^L,$$

and $\mathbf{x} = (x_1, \dots, x_L)$ denotes a generic element of $(\mathbb{Z}_d)^L$.

In the baseline model, each coordinate evolves as a single-walker FJRW–LT chain, and all random choices are independent across walkers and time. Thus, conditional on

$\mathbf{R}_{n-1} = \mathbf{x}$,

$$\mathbb{P}\big(\mathbf{R}_n = \mathbf{y} \mid \mathbf{R}_{n-1} = \mathbf{x}\big) = \prod_{\ell=1}^{L} P^{(\alpha,s)}(x_\ell, y_\ell), \qquad \mathbf{x}, \mathbf{y} \in (\mathbb{Z}_d)^L,$$

and the transition kernel is the product

$$\mathbf{P}_{\text{ind}} := \big(P^{(\alpha,s)}\big)^{\otimes L}.$$

This L-dimensional chain is invariant under diagonal shifts (adding the same $a \in \mathbb{Z}_d$ to every coordinate), and the uniform product measure $\boldsymbol{\pi}(\mathbf{x}) = d^{-L}$ is stationary.

A basic tool is the pairwise difference process: for $i \neq j$,

$$D_n^{(i,j)} := R_n^{(i)} - R_n^{(j)} \pmod{d} \in \mathbb{Z}_d.$$

In the independent-update model, the meeting event $\{R_n^{(i)} = R_n^{(j)}\}$ is equivalent to $\{D_n^{(i,j)} = 0\}$, reducing the two-walker meeting problem to a one-dimensional chain. This reduction is used repeatedly throughout the analysis.

To represent correlated updates, variants are also considered in which walkers share part or all of the randomness at a given time step. In the first variant, a single teleport decision is drawn at each time n: one Bernoulli variable B_n is sampled with $\mathbb{P}(B_n = 1) = \alpha$. If $B_n = 1$, then all walkers teleport at the same time, each choosing an independent destination $U_n^{(\ell)}$ drawn uniformly from $\mathbb{Z}_d$; if $B_n = 0$, then each walker carries out its own lazy-increment step independently.

In the second variant, a shared increment shock is introduced: at time n, one increment $X_n \sim \mu$ is drawn and applied to every walker that makes an increment move, while teleportation decisions and laziness decisions remain walker-specific.

A fully shared coupling is also considered, in which all randomness at each time step is common across walkers; this produces maximal correlation among trajectories. These couplings can substantially alter collision statistics. A shared teleport clock induces synchronized relocations that can create bursts of encounters, whereas shared increments can limit relative dispersion and thereby promote repeated meetings.

2.4 Collision, meeting, and coalescence targets

Collisions correspond to configurations where at least two coordinates coincide. For each pair $1 \leq i < j \leq L$, define the pair-diagonal

$$\mathcal{C}_{i,j} := \big\{\mathbf{x} \in (\mathbb{Z}_d)^L : \; x_i = x_j\big\},$$

and let

$$\mathcal{C} := \bigcup_{1 \leq i < j \leq L} \mathcal{C}_{i,j}$$

be the collision set. Full coalescence corresponds to the main diagonal

$$\mathcal{D} := \big\{\mathbf{x} \in (\mathbb{Z}_d)^L : \; x_1 = x_2 = \cdots = x_L\big\},$$

so $\mathcal{D} \subseteq \mathcal{C}$ when $L \geq 2$.

Given $\mathbf{R}_0 = \mathbf{x}$, define the pair meeting time

$$T_{i,j} := \inf\{n \geq 0 : \ R_n^{(i)} = R_n^{(j)}\} = \inf\{n \geq 0 : \ \mathbf{R}_n \in \mathcal{C}_{i,j}\},$$

the first collision time

$$T_{\text{coll}} := \inf\{n \geq 0 : \ \mathbf{R}_n \in \mathcal{C}\} = \min_{1 \leq i < j \leq L} T_{i,j},$$

and the full coalescence time

$$T_{\text{coal}} := \inf\{n \geq 0 : \ \mathbf{R}_n \in \mathcal{D}\}.$$

Each is a stopping time with respect to the natural filtration of $\{\mathbf{R}_n\}$.

Two standard post-encounter conventions are distinguished. Under the *observation-only* convention, walkers continue to evolve after encounters, and the above stopping times are recorded without altering the dynamics. Under the *coalescing* convention, walkers merge upon meeting and subsequently move together, so that $\mathcal{D}$ becomes absorbing once all walkers have merged; this can be formalized by an absorbed chain frozen after hitting $\mathcal{D}$.

Finally, diagonal-shift invariance implies that many quantities depend only on relative separations rather than absolute locations. In particular, pair meeting laws depend on the initial offset $R_0^{(i)} - R_0^{(j)} \pmod d$, and for T_{coll} and T_{coal} one can often reduce dependence to the multiset of initial gaps along the cycle. These symmetries are used to reduce multi-walker first-passage questions to lower-dimensional problems whenever possible.

2.5 Notation summary

Table 1 collects the main symbols used throughout.

3 Baseline Theory: Independent Walkers

This section establishes the main spectral representation and the basic first-passage reductions for the independent-update model, where the joint kernel is the L-fold product

$$\mathbf{P}_{\text{ind}} = (P^{(\alpha,s)})^{\otimes L}.$$

Two simple structural facts drive what follows. First, the chain lives on the finite abelian group $(\mathbb{Z}_d)^L$ and is translation-invariant, so it admits a complete Fourier diagonalization with explicit eigenvalues (see, e.g., [25]). Second, encounter events (meetings and collisions) are most naturally expressed in relative coordinates; in particular, when $L = 2$ the meeting problem reduces to a one-dimensional displacement chain on $\mathbb{Z}_d$.

Symbol	Meaning
d	Cycle size; state space $\mathbb{Z}_d = \{0, 1, \ldots, d-1\}$
L	Number of walkers
μ	Increment law on $\mathbb{Z}_d$ with masses $p_x = \mu(x)$
s	Laziness probability (stay put given no teleport)
α	Teleportation probability (uniform reset)
$P^{(\alpha,s)}$	Single-walker FJRW–LT kernel
λ_k	Fourier transform of μ: $\lambda_k = \sum_x p_x\, \omega^{kx}$
Λ_k	Single-walker eigenvalues (1)
Γ_k	Displacement eigenvalues: $\Gamma_k = \lvert\Lambda_k\rvert^2$
$\mathbf{R}_n$	L-walker configuration at time n
$\mathcal{C}$	Collision set (some pair coincides)
$\mathcal{D}$	Full diagonal (all coincide)
$T_{i,j}$	Meeting time of walkers i and j
T_{coll}	First collision time among L walkers
T_{coal}	Coalescence time under merge-on-meet convention

Table 1 Core notation.

Throughout this section, $\omega = \exp(2\pi i/d)$ and the single-walker eigenvalues are recalled as

$$\Lambda_0 = 1, \qquad \Lambda_k = (1-\alpha)\big(s + (1-s)\lambda_k\big), \quad k = 1, \ldots, d-1, \qquad \lambda_k = \sum_{x \in \mathbb{Z}_d} p_x\, \omega^{kx}.$$

3.1 Fourier diagonalization on the product group

Let

$$G := (\mathbb{Z}_d)^L$$

be the abelian group under componentwise addition modulo d. For $\mathbf{x} = (x_1, \ldots, x_L) \in G$ and a multi-index $\mathbf{k} = (k_1, \ldots, k_L) \in G$, define the character

$$\Phi_{\mathbf{k}}(\mathbf{x}) := \omega^{\langle \mathbf{k}, \mathbf{x} \rangle}, \qquad \langle \mathbf{k}, \mathbf{x} \rangle := \sum_{\ell=1}^{L} k_\ell x_\ell \pmod{d}. \tag{2}$$

With the normalized inner product

$$\langle f, g \rangle_G := \frac{1}{d^L} \sum_{\mathbf{x} \in G} f(\mathbf{x})\, \overline{g(\mathbf{x})},$$

the family $\{\Phi_{\mathbf{k}} : \mathbf{k} \in G\}$ is an orthonormal basis of complex-valued functions on G.

The product kernel has the explicit form

$$\mathbf{P}_{\mathrm{ind}}(\mathbf{x}, \mathbf{y}) = \prod_{\ell=1}^{L} P^{(\alpha,s)}(x_\ell, y_\ell), \qquad \mathbf{x}, \mathbf{y} \in G.$$

Since each coordinate kernel depends only on differences, $\mathbf{P}_{\mathrm{ind}}$ depends only on $\mathbf{y}-\mathbf{x}$ (componentwise modulo d). Equivalently, for any $\mathbf{a} \in G$,

$$\mathbf{P}_{\mathrm{ind}}(\mathbf{x}+\mathbf{a}, \mathbf{y}+\mathbf{a}) = \mathbf{P}_{\mathrm{ind}}(\mathbf{x}, \mathbf{y}),$$

so $\mathbf{P}_{\mathrm{ind}}$ is translation-invariant on G and is diagonalized by the product Fourier basis.

Lemma 1 (Product Fourier diagonalization) *For every* $\mathbf{k} = (k_1, \ldots, k_L) \in (\mathbb{Z}_d)^L$,

$$(\mathbf{P}_{\mathrm{ind}}\Phi_{\mathbf{k}})(\mathbf{x}) = \mathbf{\Lambda}_{\mathbf{k}}\, \Phi_{\mathbf{k}}(\mathbf{x}), \qquad \mathbf{\Lambda}_{\mathbf{k}} := \prod_{\ell=1}^{L} \Lambda_{k_\ell}.$$

Proof Using the product form of the kernel and the factorization of $\Phi_{\mathbf{k}}$,

$$\Phi_{\mathbf{k}}(\mathbf{y}) = \prod_{\ell=1}^{L} \phi_{k_\ell}(y_\ell), \qquad \phi_k(x) := \omega^{kx},$$

we obtain

$$(\mathbf{P}_{\mathrm{ind}}\Phi_{\mathbf{k}})(\mathbf{x}) = \sum_{\mathbf{y}\in G} \prod_{\ell=1}^{L} P^{(\alpha,s)}(x_\ell, y_\ell) \prod_{\ell=1}^{L} \phi_{k_\ell}(y_\ell) = \prod_{\ell=1}^{L} \sum_{y_\ell \in \mathbb{Z}_d} P^{(\alpha,s)}(x_\ell, y_\ell)\phi_{k_\ell}(y_\ell).$$

Since $P^{(\alpha,s)}\phi_{k_\ell} = \Lambda_{k_\ell}\phi_{k_\ell}$, the ℓth factor equals $\Lambda_{k_\ell}\phi_{k_\ell}(x_\ell)$. Multiplying over ℓ yields the claim. □

Let

$$\mathbf{a}_n(\mathbf{r}) := \mathbb{P}(\mathbf{R}_n = \mathbf{r} \mid \mathbf{R}_0 = \mathbf{0}), \qquad \mathbf{r} \in G.$$

Because $\mathbf{P}_{\mathrm{ind}}$ is doubly stochastic, the stationary distribution is uniform on G:

$$\boldsymbol{\pi}(\mathbf{r}) = \frac{1}{d^L}.$$

Theorem 1 (Arriving distribution on $(\mathbb{Z}_d)^L$) *For all* $n \geq 1$ *and* $\mathbf{r} \in (\mathbb{Z}_d)^L$,

$$\mathbf{a}_n(\mathbf{r}) = \frac{1}{d^L} + \frac{1}{d^L} \sum_{\mathbf{k}\in(\mathbb{Z}_d)^L\setminus\{\mathbf{0}\}} \omega^{-\langle \mathbf{k},\mathbf{r}\rangle}\, \mathbf{\Lambda}_{\mathbf{k}}^{\,n}, \qquad \mathbf{\Lambda}_{\mathbf{k}} = \prod_{\ell=1}^{L} \Lambda_{k_\ell}. \tag{3}$$

Proof Define the Fourier transform

$$\widehat{\mathbf{a}}_n(\mathbf{k}) := \sum_{\mathbf{r}\in G} \mathbf{a}_n(\mathbf{r})\, \omega^{\langle \mathbf{k},\mathbf{r}\rangle} = \mathbb{E}[\Phi_{\mathbf{k}}(\mathbf{R}_n) \mid \mathbf{R}_0 = \mathbf{0}]\,.$$

Lemma 1 gives $\widehat{\mathbf{a}}_n(\mathbf{k}) = \mathbf{\Lambda}_{\mathbf{k}}^n$. Fourier inversion on G yields (3). □

An immediate Parseval consequence is the L^2 identity

$$\sum_{\mathbf{r}\in G}\left|\mathbf{a}_n(\mathbf{r})-\frac{1}{d^L}\right|^2=\frac{1}{d^L}\sum_{\mathbf{k}\neq\mathbf{0}}|\mathbf{\Lambda_k}|^{2n}. \tag{4}$$

As in the single-walker setting, (4) converts to total-variation bounds via Cauchy–Schwarz.

3.2 Two-walker reduction via relative displacement

The analysis now specializes to $L=2$, where the meeting problem can be converted into a one-dimensional hitting-time problem on $\mathbb{Z}_d$. Let $(R_n^{(1)},R_n^{(2)})$ be two independent walkers with kernel $P^{(\alpha,s)}$ and define the displacement

$$D_n:=R_n^{(1)}-R_n^{(2)}\ (\mathrm{mod}\ d)\in\mathbb{Z}_d.$$

The meeting time can be written as

$$T_{1,2}:=\min\{n\geq 1:\ R_n^{(1)}=R_n^{(2)}\}=\min\{n\geq 1:\ D_n=0\},$$

so $T_{1,2}$ is the hitting time of 0 for (D_n).

Lemma 2 (Markov property and translation invariance of (D_n)) *In the independent-update model, $(D_n)_{n\geq 0}$ is a time-homogeneous Markov chain on $\mathbb{Z}_d$. Moreover, its transition kernel is circulant (depends only on differences), so (D_n) is translation-invariant and has the uniform stationary distribution.*

Proof sketch The pair process is Markov on $(\mathbb{Z}_d)^2$ with kernel $P^{(\alpha,s)}\otimes P^{(\alpha,s)}$. Since D_n is a function of the pair state and the pair kernel is invariant under diagonal shifts, the induced process on differences is Markov and translation-invariant. Uniform stationarity follows from double stochasticity. □

To identify the spectrum, observe that

$$\omega^{k(x_1-x_2)}=\phi_k(x_1)\phi_{-k}(x_2)$$

is a product Fourier mode. Under $\mathbf{P}_{\mathrm{ind}}$, its eigenvalue is $\Lambda_k\Lambda_{-k}$; since the kernel has real entries, $\Lambda_{-k}=\overline{\Lambda_k}$, hence $\Lambda_k\Lambda_{-k}=|\Lambda_k|^2$.

Lemma 3 (Displacement eigenvalues) *Let $\Gamma_k:=|\Lambda_k|^2$ for $k=0,1,\dots,d-1$ (so $\Gamma_0=1$). Then the displacement chain (D_n) is diagonalized by $\phi_k(r)=\omega^{kr}$ with eigenvalues*

$$\Gamma_0=1,\qquad \Gamma_k=|\Lambda_k|^2,\quad k=1,\dots,d-1.$$

Proof Consider $f_k(x_1, x_2) = \phi_k(x_1)\phi_{-k}(x_2) = \omega^{k(x_1-x_2)}$ on $(\mathbb{Z}_d)^2$. By Lemma 1 with $L = 2$, $(P^{(\alpha,s)} \otimes P^{(\alpha,s)})f_k = \Lambda_k\Lambda_{-k}f_k$. Since f_k depends only on $x_1 - x_2$, it descends to an eigenfunction for (D_n) with eigenvalue $\Lambda_k\Lambda_{-k} = |\Lambda_k|^2$. □

Let $a_n^D(r) := \mathbb{P}(D_n = r \mid D_0 = 0)$. Fourier inversion gives the arriving distribution.

Theorem 2 (Arriving distribution for the displacement chain) *For all $n \geq 1$ and $r \in \mathbb{Z}_d$,*

$$a_n^D(r) = \frac{1}{d} + \frac{1}{d}\sum_{k=1}^{d-1} \omega^{-rk}\,\Gamma_k^n, \qquad \Gamma_k = |\Lambda_k|^2. \tag{5}$$

More generally, if $D_0 = \delta$, then

$$\mathbb{P}(D_n = r \mid D_0 = \delta) = \frac{1}{d} + \frac{1}{d}\sum_{k=1}^{d-1} \omega^{-(r-\delta)k}\,\Gamma_k^n.$$

Meeting-time moments are obtained by applying the renewal/generating-function method to the displacement chain, with Γ_k replacing the single-walker eigenvalues. Let

$$T_{1,2}^{(\delta)} := \min\{n \geq 1 :\ D_n = 0 \mid D_0 = \delta\}.$$

By translation invariance, this is equivalent to hitting $r_\delta := (-\delta) \pmod d$ starting from 0.

Theorem 3 (Mean meeting time for two independent walkers) *Assume the displacement chain is irreducible (equivalently $\Gamma_k \neq 1$ for all $k \neq 0$). Then for any initial displacement $\delta \in \mathbb{Z}_d$,*

$$\mathbb{E}\Big[T_{1,2}^{(\delta)}\Big] = d + \sum_{k=1}^{d-1}\big(1 - \omega^{-r_\delta k}\big)\frac{\Gamma_k}{1-\Gamma_k} = d + \sum_{k=1}^{d-1}\big(1 - \omega^{\delta k}\big)\frac{\Gamma_k}{1-\Gamma_k}, \qquad \Gamma_k = |\Lambda_k|^2. \tag{6}$$

Theorem 4 (Variance of the meeting time) *Under the same assumptions, let $\mu_\delta := \mathbb{E}[T_{1,2}^{(\delta)}]$ be given by* (6)*. Then*

$$\mathrm{Var}\Big(T_{1,2}^{(\delta)}\Big) = 2\sum_{k=1}^{d-1}\big(1-\omega^{\delta k}\big)\frac{\Gamma_k}{(1-\Gamma_k)^2} + \mu_\delta\left(d - 1 + \sum_{k=1}^{d-1}\big(1+\omega^{\delta k}\big)\frac{\Gamma_k}{1-\Gamma_k}\right), \qquad \Gamma_k = |\Lambda_k|^2. \tag{7}$$

Proof sketch for Theorems 3–4 Apply the renewal identity and generating-function argument to (D_n), using (5) for the displacement arriving distribution. The first-hit generating function is obtained by the standard quotient identity; differentiation at $z \uparrow 1$ yields the stated Fourier sums with Γ_k in place of Λ_k. □

Although (6)–(7) are written in complex form, the values are real because terms pair as k and $d - k$. If desired, one can rewrite $1 - \omega^{\delta k}$ using trigonometric identities to obtain an explicit cosine-series representation.

3.3 Numerical evaluation of the Fourier-sum formulas

The closed-form expressions in Theorems 1, 2, 3, and 4 are finite Fourier sums, so once the eigenvalues Λ_k (or Γ_k) are available they can be evaluated efficiently.

To obtain λ_k and Λ_k, let $\mathbf{p} = (p_0, \dots, p_{d-1})$ denote the increment probability vector. Then $(\lambda_k)_{k=0}^{d-1}$ is exactly the discrete Fourier transform (DFT) of $\mathbf{p}$. If μ is supported on only $m \ll d$ points (for instance, forward jumps of bounded range), each λ_k can be computed directly in $O(md)$ time. In the general case where μ is an arbitrary distribution on $\mathbb{Z}_d$, all λ_k can be computed via an FFT in $O(d \log d)$ time. From these, $\Lambda_k = (1-\alpha)\big(s + (1-s)\lambda_k\big)$ for $k \neq 0$, with $\Lambda_0 = 1$.

For the meeting-time sums, the mean and variance formulas contain terms such as $\Gamma_k/(1-\Gamma_k)$ and $\Gamma_k/(1-\Gamma_k)^2$. When Γ_k is close to 1—as in diffusion-dominant or crossover regimes—straightforward floating-point evaluation may lose precision due to cancellation. This can be mitigated by rewriting the sums in a real cosine form (Appendix B) to avoid complex-conjugate cancellation, by using the identities $\Gamma/(1-\Gamma) = -1 + 1/(1-\Gamma)$ and $\Gamma/(1-\Gamma)^2 = 1/(1-\Gamma)^2 - 1/(1-\Gamma)$, or, when d is large and α is small, by evaluating $1-\Gamma_k$ with `expm1`-style transformations after expressing $\Gamma_k = (1-\alpha)^2|\widetilde{\lambda}_k|^2$.

In terms of computational complexity, once $\{\lambda_k\}$ are computed, evaluating any single Fourier sum over $k = 1, \dots, d-1$ is $O(d)$. Thus, for fixed parameters (α, s, μ), one can compute $\mathbb{E}[T_{1,2}^{(\delta)}]$ for *all* δ either by $O(d^2)$ direct summation (loop over δ and k) or by applying an inverse FFT to the sequence of coefficients in the cosine-series representation to obtain all δ values simultaneously.

3.4 First collision among many walkers

The analysis now returns to general $L \geq 2$ and considers the first collision time

$$T_{\text{coll}} := \min\Big\{n \geq 1 :\ \exists\, i < j \text{ such that } R_n^{(i)} = R_n^{(j)}\Big\}.$$

Equivalently, define the collision (union-of-diagonals) set

$$\mathcal{C} := \bigcup_{1 \leq i < j \leq L} \{\mathbf{x} \in (\mathbb{Z}_d)^L :\ x_i = x_j\}, \qquad \text{so that} \qquad T_{\text{coll}} = \min\{n \geq 1 :\ \mathbf{R}_n \in \mathcal{C}\}. \tag{8}$$

Exact hitting-time formulas for unions of diagonals are typically not available, but the product structure yields simple bounds and accurate approximations once the configuration is close to uniform.

(i) Union and second-moment bounds

Let $A_{i,j}(n) := \{R_n^{(i)} = R_n^{(j)}\}$. Then $\{\mathbf{R}_n \in \mathcal{C}\} = \cup_{i<j} A_{i,j}(n)$ and, for any fixed n,

$$\mathbb{P}(\mathbf{R}_n \in \mathcal{C}) \leq \sum_{1 \leq i < j \leq L} \mathbb{P}\big(A_{i,j}(n)\big) = \binom{L}{2} \mathbb{P}\big(R_n^{(1)} = R_n^{(2)}\big), \tag{9}$$

where the final equality uses exchangeability. Moreover,

$$\mathbb{P}\big(R_n^{(1)} = R_n^{(2)}\big) = \mathbb{P}(D_n = 0) = \frac{1}{d} + \frac{1}{d}\sum_{k=1}^{d-1} \Gamma_k^n$$

by Theorem 2 (with $r = 0$).

A lower bound follows from inclusion–exclusion:

$$\mathbb{P}(\mathbf{R}_n \in \mathcal{C}) \geq \sum_{i<j} \mathbb{P}(A_{i,j}(n)) - \sum_{(i<j)\neq(k<\ell)} \mathbb{P}\big(A_{i,j}(n) \cap A_{k,\ell}(n)\big). \tag{10}$$

The intersection terms can be controlled by symmetry and become close to $1/d^2$ for disjoint pairs once the configuration is well mixed, recovering the classical occupancy scaling.

(ii) Block-tail bounds from mixing

Let

$$p_t := \inf_{\mathbf{x}} \mathbb{P}_{\mathbf{x}}(T_{\text{coll}} \leq t)$$

be the worst-case probability of producing a collision within t steps. A standard block argument gives

$$\mathbb{P}_{\mathbf{x}}(T_{\text{coll}} > kt) \leq (1 - p_t)^k \leq \exp(-kp_t), \tag{11}$$

so it suffices to obtain a usable lower bound on p_t.

Let $t_{\text{mix}}^{(L)}(\varepsilon)$ denote the total-variation mixing time of $\mathbf{P}_{\text{ind}}$ on G [2] and set

$$t_\star := t_{\text{mix}}^{(L)}\left(\frac{\boldsymbol{\pi}(\mathcal{C})}{2}\right).$$

Then for every start $\mathbf{x}$,

$$\mathbf{P}_{\text{ind}}^{t_\star}(\mathbf{x}, \mathcal{C}) \geq \boldsymbol{\pi}(\mathcal{C}) - \frac{\boldsymbol{\pi}(\mathcal{C})}{2} = \frac{\boldsymbol{\pi}(\mathcal{C})}{2}.$$

Since $\{\mathbf{R}_{t_\star} \in \mathcal{C}\} \subseteq \{T_{\text{coll}} \leq t_\star\}$, we obtain

$$p_{t_\star} \geq \frac{\boldsymbol{\pi}(\mathcal{C})}{2},$$

and therefore

$$\mathbb{P}_{\mathbf{x}}(T_{\text{coll}} > kt_\star) \leq \exp\left(-k\,\frac{\boldsymbol{\pi}(\mathcal{C})}{2}\right). \tag{12}$$

Under $\boldsymbol{\pi}$, $\boldsymbol{\pi}(\mathcal{C})$ is the probability that L i.i.d. uniform samples on $\mathbb{Z}_d$ contain a repeat:

$$\boldsymbol{\pi}(\mathcal{C}) = 1 - \frac{(d)_L}{d^L}, \qquad (d)_L := d(d-1)\cdots(d-L+1). \tag{13}$$

In the sparse regime $L = o(\sqrt{d})$, the standard approximation yields

$$\boldsymbol{\pi}(\mathcal{C}) = 1 - \exp\left(-\frac{L(L-1)}{2d} + o(1)\right) \approx \frac{L(L-1)}{2d}, \tag{14}$$

i.e., the familiar birthday-paradox scaling.

(iii) Well-mixed approximations in teleport-dominant regimes

When teleportation is present and α is not too small, the product chain typically refreshes rapidly, and collisions behave approximately as if each step produced an (almost) independent occupancy configuration. In that regime, T_{coll} is well approximated by a geometric random variable with success parameter $\boldsymbol{\pi}(\mathcal{C})$, leading to the heuristic

$$\mathbb{P}(T_{\text{coll}} > n) \approx (1 - \boldsymbol{\pi}(\mathcal{C}))^n, \qquad \mathbb{E}[T_{\text{coll}}] \approx \frac{1}{\boldsymbol{\pi}(\mathcal{C})}.$$

Combining with (14) suggests the scale

$$\mathbb{E}[T_{\text{coll}}] \approx \frac{2d}{L(L-1)} \qquad \text{when } L = o(\sqrt{d}). \tag{15}$$

The bound (12) provides a rigorous analogue of this picture, with an additional factor $t_\star$ capturing the time required to mix toward stationarity.

3.5 Coalescence time

Two related notions of global agreement are distinguished. The first treats $\mathcal{D}$ as a target set for the *non-coalescing* dynamics. The second modifies the dynamics by *merging* walkers after meetings, producing a coalescing-walker process. Both quantify a time to synchronization, but they reflect different application conventions.

(A) All-equal target (non-coalescing walkers). Let

$$\mathcal{D} := \{\mathbf{x} \in (\mathbb{Z}_d)^L : \ x_1 = x_2 = \cdots = x_L\}, \qquad T_{\text{all}} := \min\{n \geq 1 : \ \mathbf{R}_n \in \mathcal{D}\}.$$

This is a small-target hitting problem: $\boldsymbol{\pi}(\mathcal{D}) = d^{-(L-1)}$, so in a well-mixed regime T_{all} is typically large for $L \geq 3$.

(B) Coalescing walkers (merge-after-meeting). Consider a coalescing dynamics in which walkers that meet merge and subsequently move together. A convenient representation tracks a partition of $\{1, \ldots, L\}$ into clusters together with a position for each cluster; whenever two clusters occupy the same vertex, they merge. Let K_n be the number of clusters at time n ($K_0 = L$) and define

$$T_{\text{coal}} := \min\{n \geq 0 : \ K_n = 1\}.$$

A useful decomposition represents T_{coal} as a sum of successive merge times. Let τ_k denote the time required to reduce the number of clusters from k to $k-1$,

$$\tau_k := \min\{n \geq 1 : \ K_{t_{k-1}+n} \leq k-1\}, \quad \text{where } t_{k-1} := \sum_{j=k}^{L} \tau_j,$$

so that $T_{\mathrm{coal}} = \sum_{k=2}^{L} \tau_k$. Conditional on having k clusters, the next merge occurs exactly when the k cluster positions experience a collision.

Let $\mathcal{C}_k \subseteq (\mathbb{Z}_d)^k$ be the collision set for k labeled particles. Under the uniform product measure on $(\mathbb{Z}_d)^k$,

$$\pi_k(\mathcal{C}_k) = 1 - \frac{(d)_k}{d^k}.$$

Assume that between merges, clusters evolve as independent copies of $P^{(\alpha,s)}$. Let

$$t_\star^{(k)} := t_{\mathrm{mix}}^{(k)}\left(\frac{\pi_k(\mathcal{C}_k)}{2}\right)$$

be the k-fold product mixing time [2]. The same minorization argument as in (12) yields

$$\inf_{\mathbf{x}\in(\mathbb{Z}_d)^k} \mathbb{P}_{\mathbf{x}}\big(\text{a collision occurs within } t_\star^{(k)}\big) \geq \frac{\pi_k(\mathcal{C}_k)}{2},$$

and therefore

$$\mathbb{E}[\tau_k] \ \leq \ \frac{2\, t_\star^{(k)}}{\pi_k(\mathcal{C}_k)}. \tag{16}$$

Summing over k gives

$$\mathbb{E}[T_{\mathrm{coal}}] \leq \sum_{k=2}^{L} \frac{2\, t_\star^{(k)}}{\pi_k(\mathcal{C}_k)} = \sum_{k=2}^{L} \frac{2\, t_\star^{(k)}}{1-(d)_k/d^k}. \tag{17}$$

In teleport-dominant regimes (e.g., α bounded away from 0), the primary dependence on the dynamics enters through the refresh/mixing scale $t_\star^{(k)}$, while the occupancy factor $1/\pi_k(\mathcal{C}_k)$ captures the collision opportunity at stationarity. In the sparse regime $k = o(\sqrt{d})$, $\pi_k(\mathcal{C}_k) \approx k(k-1)/(2d)$, which suggests

$$\mathbb{E}[\tau_k] \ \lesssim \ \frac{4d}{k(k-1)}\, t_\star^{(k)}, \qquad \mathbb{E}[T_{\mathrm{coal}}] \ \lesssim \ 4d\Big(\sum_{k=2}^{L} \frac{1}{k(k-1)}\Big) \max_{2\leq k\leq L} t_\star^{(k)} = 4d\left(1-\frac{1}{L}\right) \max_k t_\star^{(k)}.$$

A simple lower bound follows from the fact that coalescence cannot occur before any pair meets: for any distinct i, j,

$$\mathbb{E}[T_{\mathrm{coal}}] \geq \mathbb{E}[T_{i,j}],$$

and for $L = 2$ the right-hand side is given explicitly by Theorem 3 (once the initial displacement is specified). More pronounced dependence on L is expected when L becomes comparable to $\sqrt{d}$, where the sparse approximation (14) no longer applies.

4 Teleportation and Collision Frequency

This section quantifies how teleportation changes collision statistics in the multi-walker model. Teleportation contributes a *uniform refresh* component, so each step is a mixture of (i) a full reset to the uniform law and (ii) a "no-teleport" move. For independent walkers, this Doeblin-type structure lifts to the L-fold product chain through an explicit tensor expansion, leading to concrete total-variation bounds that are uniform over initial states. These bounds translate directly into (a) explicit lower bounds on collision probabilities over suitable time blocks and (b) exponential tails for first-collision and phase-by-phase coalescence times. The resulting picture organizes collision behavior into three regimes: teleport-dominant (near-i.i.d. occupancy), diffusion-dominant (rare refresh, relative diffusion controls encounters), and an intermediate crossover where both effects act on comparable time scales.

4.1 Mixture identities and consequences for the product chain

For a single walker on $\mathbb{Z}_d$, write the FJRW–LT kernel as the Doeblin mixture

$$P^{(\alpha,s)} = \alpha\Pi + (1-\alpha)Q, \qquad \Pi(x,\cdot) = \pi(\cdot) = \frac{1}{d}, \qquad Q := sI + (1-s)P.$$

The uniform kernel Π is a projection ($\Pi^2 = \Pi$) and satisfies $\Pi Q = Q\Pi = \Pi$. Consequently, the n-step kernel admits the exact identity

$$\big(P^{(\alpha,s)}\big)^n = \Pi + (1-\alpha)^n\big(Q^n - \Pi\big), \qquad n \geq 0, \tag{18}$$

and therefore

$$\|P^{(\alpha,s)\,n}(x,\cdot) - \pi\|_{\mathrm{TV}} \leq (1-\alpha)^n, \qquad \text{for all } x \in \mathbb{Z}_d,$$

a standard uniform-ergodicity bound under Doeblin minorization (see, e.g., [2, 26]).

Identity (18) is next lifted to L independent walkers on $G = (\mathbb{Z}_d)^L$. The joint kernel is

$$\mathbf{P}_{\mathrm{ind}} := \big(P^{(\alpha,s)}\big)^{\otimes L}, \qquad \boldsymbol{\pi} := \pi^{\otimes L},$$

and $\mathbf{\Pi}$ denotes the projection kernel on G with every row equal to $\boldsymbol{\pi}$. Define also the "no-teleport" product kernel $\mathbf{Q} := Q^{\otimes L}$.

A convenient form of (18) is the time-n mixture

$$\big(P^{(\alpha,s)}\big)^n = (1-\rho_n)\Pi + \rho_n Q^n, \qquad \rho_n := (1-\alpha)^n,$$

which tensorizes into an explicit subset expansion.

Lemma 4 (Subset mixture expansion for the n-step product chain) *For every $n \geq 0$ and every start $\mathbf{x} = (x_1, \ldots, x_L) \in (\mathbb{Z}_d)^L$,*

$$\mathbf{P}_{\mathrm{ind}}^n(\mathbf{x}, \cdot) = \sum_{S \subseteq [L]} (1-\rho_n)^{L-|S|} \rho_n^{|S|} \left(\bigotimes_{\ell \in S} Q^n(x_\ell, \cdot) \right) \otimes \pi^{\otimes([L]\setminus S)}, \qquad \rho_n = (1-\alpha)^n. \quad (19)$$

Proof Since $\mathbf{P}_{\mathrm{ind}}^n = \big(P^{(\alpha,s)\,n}\big)^{\otimes L}$ and $P^{(\alpha,s)\,n} = (1-\rho_n)\Pi + \rho_n Q^n$, expanding the L-fold tensor product across coordinates yields (19), with S indexing the coordinates contributing the Q^n term. □

Although (19) is not a two-component mixture on G (because refresh occurs independently across coordinates), it still yields a clean worst-case total-variation bound. Indeed, each coordinate is exactly at stationarity with probability $1-\rho_n$, independently across ℓ, so

$$\sup_{\mathbf{x}} \big\| \mathbf{P}_{\mathrm{ind}}^n(\mathbf{x}, \cdot) - \boldsymbol{\pi} \big\|_{\mathrm{TV}} \leq 1 - (1-\rho_n)^L = 1 - \big(1 - (1-\alpha)^n\big)^L. \quad (20)$$

For small $(1-\alpha)^n$, the bound simplifies to

$$\sup_{\mathbf{x}} \big\| \mathbf{P}_{\mathrm{ind}}^n(\mathbf{x}, \cdot) - \boldsymbol{\pi} \big\|_{\mathrm{TV}} \leq L(1-\alpha)^n. \quad (21)$$

Under the shared-teleport-clock coupling from Section 2.3 (a single Bernoulli teleport decision common to all walkers, with independent uniform landing sites), the joint kernel is a genuine Doeblin mixture $\mathbf{P} = \alpha\boldsymbol{\Pi} + (1-\alpha)\mathbf{Q}$. In this case one recovers the dimension-free identity $\mathbf{P}^n = \boldsymbol{\Pi} + (1-\alpha)^n(\mathbf{Q}^n - \boldsymbol{\Pi})$ and hence $\sup_{\mathbf{x}} \|\mathbf{P}^n(\mathbf{x}, \cdot) - \boldsymbol{\pi}\|_{\mathrm{TV}} \leq (1-\alpha)^n$.

4.2 Teleportation-based collision bounds

Refresh control is next converted into collision bounds. Let $\mathcal{C} \subseteq (\mathbb{Z}_d)^L$ be the collision set (union of pairwise diagonals). Under the uniform product law,

$$\boldsymbol{\pi}(\mathcal{C}) = 1 - \frac{(d)_L}{d^L}, \qquad (d)_L := d(d-1)\cdots(d-L+1). \quad (22)$$

In the independent-update model, the event that *every* walker teleports at a given step has probability α^L. Conditional on this event, the next configuration is i.i.d. uniform across walkers, so

$$\inf_{\mathbf{x}} \mathbf{P}_{\mathrm{ind}}(\mathbf{x}, \mathcal{C}) \;\geq\; \alpha^L \boldsymbol{\pi}(\mathcal{C}) \;=\; \alpha^L \left(1 - \frac{(d)_L}{d^L}\right). \quad (23)$$

If $T_{\mathrm{coll}} := \min\{n \geq 1 : \mathbf{R}_n \in \mathcal{C}\}$ denotes the first-collision time and $p_{\mathrm{coll}} := \alpha^L \boldsymbol{\pi}(\mathcal{C})$, then the Markov property yields

$$\mathbb{P}_{\mathbf{x}}(T_{\mathrm{coll}} > n) \leq (1 - p_{\mathrm{coll}})^n \leq \exp\bigl(-p_{\mathrm{coll}}\, n\bigr), \qquad n \geq 0, \tag{24}$$

a conservative bound since it counts only collisions created by all-teleport steps.

A sharper bound uses total variation. For any $t \geq 1$ and any start $\mathbf{x}$,

$$\mathbf{P}^t_{\mathrm{ind}}(\mathbf{x}, \mathcal{C}) \geq \boldsymbol{\pi}(\mathcal{C}) - \bigl\|\mathbf{P}^t_{\mathrm{ind}}(\mathbf{x}, \cdot) - \boldsymbol{\pi}\bigr\|_{\mathrm{TV}},$$

so by (20),

$$\inf_{\mathbf{x}} \mathbf{P}^t_{\mathrm{ind}}(\mathbf{x}, \mathcal{C}) \geq \boldsymbol{\pi}(\mathcal{C}) - \Bigl[1 - \bigl(1 - (1-\alpha)^t\bigr)^L\Bigr]. \tag{25}$$

In particular, if t is chosen so that

$$1 - \bigl(1 - (1-\alpha)^t\bigr)^L \leq \frac{\boldsymbol{\pi}(\mathcal{C})}{2} \qquad \text{(e.g., it suffices that } L(1-\alpha)^t \leq \boldsymbol{\pi}(\mathcal{C})/2),$$

then

$$p_t := \inf_{\mathbf{x}} \mathbb{P}_{\mathbf{x}}(T_{\mathrm{coll}} \leq t) \geq \frac{\boldsymbol{\pi}(\mathcal{C})}{2}. \tag{26}$$

A standard block argument (e.g., [2]) then gives

$$\mathbb{P}_{\mathbf{x}}(T_{\mathrm{coll}} > kt) \leq \exp\left(-k\,\frac{\boldsymbol{\pi}(\mathcal{C})}{2}\right), \qquad k \geq 0, \tag{27}$$

with a block length that can be taken on the scale

$$t \;\asymp\; \frac{1}{\alpha}\left(\log L + \log\frac{1}{\boldsymbol{\pi}(\mathcal{C})}\right), \tag{28}$$

using $(1-\alpha)^t \approx e^{-\alpha t}$ and the sufficient condition $L(1-\alpha)^t \leq \boldsymbol{\pi}(\mathcal{C})/2$.

The same reasoning applies to coalescing dynamics. During the phase with k clusters, let $\mathcal{C}_k \subseteq (\mathbb{Z}_d)^k$ be the collision set for k labeled particles and

$$\pi_k(\mathcal{C}_k) = 1 - \frac{(d)_k}{d^k}.$$

Choosing a block length t_k so that $k(1-\alpha)^{t_k} \leq \pi_k(\mathcal{C}_k)/2$ yields

$$\mathbb{P}(\tau_k > m t_k) \leq \exp\left(-m\,\frac{\pi_k(\mathcal{C}_k)}{2}\right), \qquad m \geq 0, \tag{29}$$

and hence

$$\mathbb{E}[\tau_k] \;\leq\; \frac{2t_k}{\pi_k(\mathcal{C}_k)}, \qquad t_k \;\asymp\; \frac{1}{\alpha}\left(\log k + \log\frac{1}{\pi_k(\mathcal{C}_k)}\right). \tag{30}$$

This decomposition cleanly separates a refresh contribution (order $1/\alpha$) from an occupancy contribution (order $1/\pi_k(\mathcal{C}_k)$).

4.3 Selection of bounds and approximations

The collision and coalescence estimates developed in Sections 3.4 and 4.2 serve different analytical purposes. Their relative usefulness depends on whether the objective is a uniform worst-case guarantee or a sharper approximation in a specific parameter regime:

- **Worst-case, fully explicit tail bound.** The one-step all-teleport minorization in (24) is the most conservative estimate. It depends only on (d, L, α) and is uniform over all initial configurations, but it can be very loose when α^L is small.
- **Uniform tail bound after burn-in.** The total-variation block argument in (25)–(27) is typically sharper. One selects a block length t for which the total-variation term is at most, for example, $\boldsymbol{\pi}(\mathcal{C})/2$; the resulting tail then decays exponentially at a rate comparable to $\boldsymbol{\pi}(\mathcal{C})$.
- **Well-mixed occupancy approximation.** When the refresh mechanism is sufficiently strong, the configuration after burn-in is well approximated by independent occupancy draws from $\boldsymbol{\pi}$. In this regime, T_{coll} is approximated by a geometric random variable with success probability $\boldsymbol{\pi}(\mathcal{C})$, giving the scale $\mathbb{E}[T_{\mathrm{coll}}] \approx 1/\boldsymbol{\pi}(\mathcal{C})$. This approximation is most accurate in teleport-dominant regimes and least accurate at early times from highly clustered initial configurations.

For coalescing dynamics, the same hierarchy of estimates applies phase by phase after replacing L by the current number of clusters k and $\boldsymbol{\pi}(\mathcal{C})$ by $\pi_k(\mathcal{C}_k)$.

4.4 Regimes and scaling laws

Teleportation introduces a refresh time scale of order α^{-1}. In contrast, for bounded-step lazy increments, the characteristic diffusive mixing time along the cycle is of order d^2 (up to constants depending on μ and s). Comparing these two scales yields three collision regimes.

When α is large enough that refresh acts well before diffusion mixes, the joint law approaches $\boldsymbol{\pi}$ rapidly. After a burn-in of order t as in (28), collision events are well approximated by i.i.d. occupancy draws from $\boldsymbol{\pi}$. Under stationarity,

$$\boldsymbol{\pi}(\mathcal{C}) = 1 - \frac{(d)_L}{d^L} \approx 1 - \exp\!\left(-\frac{L(L-1)}{2d}\right),$$

and in the sparse regime $L = o(\sqrt{d})$ one has $\boldsymbol{\pi}(\mathcal{C}) \approx L(L-1)/(2d)$. Thus T_{coll} is close to a geometric time with success parameter $\boldsymbol{\pi}(\mathcal{C})$, giving the heuristic scale

$$\mathbb{E}[T_{\mathrm{coll}}] \;\approx\; \frac{1}{\boldsymbol{\pi}(\mathcal{C})} \;\approx\; \frac{2d}{L(L-1)}, \tag{31}$$

while (27) provides rigorous exponential tails. For coalescing dynamics, the same occupancy heuristic gives $\mathbb{E}[T_{\mathrm{coal}}] \approx \sum_{k=2}^{L} 1/\pi_k(\mathcal{C}_k) \approx 2d(1-1/L)$ in the sparse regime, with refresh overhead captured by t_k in (30).

When α is very small, long stretches of the evolution behave like the no-teleport chain Q, and collisions are driven mainly by relative diffusion. For $L = 2$, this reduces exactly to a displacement-chain hitting problem (Section 3.2), typically producing a diffusive scale of order d^2 for bounded-step increments. For general L, a crude benchmark is obtained by noting that if a typical pair meeting time is order d^2, then (heuristically) the first collision among $\binom{L}{2}$ pairs occurs on the order of d^2/L^2 when L is not too large and pair dependencies are weak.

Lemma 5 (Low-frequency expansion and the d^{-2} diffusion scale) *Assume the increment law μ is the reduction modulo d of an integer-valued law with finite second moment, with mean $m_1 = \mathbb{E}[X]$ and variance $\sigma^2 = \mathrm{Var}(X)$, and assume these moments do not depend on d (e.g. μ has bounded support independent of d). Let $t_k := 2\pi k/d$. Then as $d \to \infty$ with $k = o(d)$,*

$$\lambda_k = \mathbb{E}[e^{it_k X}] = 1 + im_1 t_k - \frac{\sigma^2}{2} t_k^2 + o(t_k^2), \qquad |\lambda_k|^2 = 1 - \sigma^2 t_k^2 + o(t_k^2).$$

Consequently, for the no-teleport eigenvalues $\widetilde{\lambda}_k = s + (1-s)\lambda_k$ and the displacement eigenvalues $\Gamma_k = |\Lambda_k|^2 = (1-\alpha)^2 |\widetilde{\lambda}_k|^2$,

$$1 - \Gamma_k = 2\alpha + (1-s)^2 \sigma^2 \, t_k^2 + o(\alpha) + o(t_k^2) \qquad (k = o(d)).$$

In particular, when $\alpha = 0$ the spectral gap of the displacement chain is of order d^{-2}, while for $\alpha > 0$ the small-k contraction is governed jointly by α and d^{-2}.

Remark 2 Lemma 5 is stated in a "large-d" form to justify the heuristic comparison between refresh (α^{-1}) and diffusion (d^2) time scales. Constants depend on (s, σ^2) and the regularity of the increment law near frequency 0.

A crossover occurs when refresh and diffusion contribute on comparable scales, heuristically when

$$\alpha \asymp d^{-2}. \tag{32}$$

This can be seen spectrally: for small Fourier modes in bounded-step regimes, $s + (1 - s)\lambda_k \approx 1 - c(k/d)^2$, and hence

$$\Lambda_k = (1-\alpha)\big(s + (1-s)\lambda_k\big) \approx 1 - \alpha - c\frac{k^2}{d^2}.$$

Thus contraction is governed by α and $1/d^2$ on equal footing. For times $n \ll \alpha^{-1}$, collision statistics resemble those of Q, while for $n \gg \alpha^{-1}$ repeated refresh events drive the configuration toward $\boldsymbol{\pi}$ and collision probabilities toward $\boldsymbol{\pi}(\mathcal{C})$. The inequalities (25)–(27) quantify this transition by explicitly linking collision probabilities to total-variation proximity to stationarity.

In summary, teleportation shapes collision frequency through two complementary mechanisms: it accelerates convergence of the joint law to the uniform product distribution (via (20)) and it supplies explicit minorization constants (either the one-step all-teleport event (23) or the block minorization (26)) that yield exponential tails for T_{coll} and for phase-by-phase merge times in coalescing dynamics.

5 Interaction Rules: Blocking, Bouncing, Pushing

The multi-walker FJRW–LT framework is extended by allowing *interactions* that modify how proposed moves are executed when walkers attempt to occupy the same site. Such rules arise naturally in multi-agent systems with crowding, exclusion, congestion, or queueing constraints. They also typically break the circulant structure of the independent model, so the Fourier diagonalization from Section 3 no longer applies directly. The objective of this section is twofold: first, to define a compact and reusable family of interacting kernels; and second, to formulate a perturbative viewpoint in which an interacting kernel is treated as a controlled deviation from the baseline product chain. This viewpoint enables the transfer of mixing and first-passage predictions when interaction events are sufficiently rare.

5.1 Interacting dynamics

Fix $L \geq 2$ and a configuration $\mathbf{x} = (x_1, \ldots, x_L) \in (\mathbb{Z}_d)^L$. Let

$$\mathcal{O}(\mathbf{x}) := \{x_1, \ldots, x_L\}, \qquad N_{\mathbf{x}}(u) := \big|\{\ell : \ x_\ell = u\}\big|, \quad u \in \mathbb{Z}_d,$$

denote the occupied sites (as a multiset) and the occupancy count at site u. Given the current configuration $\mathbf{x}$, each walker ℓ forms a proposal y_ℓ using the same local FJRW–LT mechanism as in the independent model. Specifically, with probability α the walker proposes to teleport by drawing y_ℓ uniformly from $\mathbb{Z}_d$. With probability $(1-\alpha)s$ it proposes to remain in place, so that $y_\ell = x_\ell$. With the remaining probability $(1-\alpha)(1-s)$ it proposes an increment move of the form $y_\ell = x_\ell + X_\ell \pmod d$, where $X_\ell \sim \mu$. Collecting these proposals yields the proposal vector $\mathbf{y} = (y_1, \ldots, y_L)$.

An *interaction rule* then maps $(\mathbf{x}, \mathbf{y})$ to the executed next configuration $\mathbf{x}'$. This rule may depend on the current occupancy profile $N_{\mathbf{x}}(\cdot)$ and, under synchronous updates, on the joint set of proposals $\mathbf{y}$.

Teleportation can interact with congestion in different ways, and we distinguish two conventions. Under an *unconstrained reset*, teleport destinations are sampled uniformly on $\mathbb{Z}_d$ without regard to current occupancy, so teleportation may create co-occupancy even when local moves are subject to exclusion. Under a *constrained reset*, teleportation must respect the interaction rule; for example, a teleport may be drawn uniformly from vacant sites or repeatedly resampled until it lands on a vacant site. Unless stated otherwise, we adopt the unconstrained convention because it preserves a global refresh component and enables uniform tail bounds via Doeblin-type arguments [2, 26].

Both timing conventions are allowed: (i) *synchronous* updates, where all proposals are generated from $\mathbf{x}$ and then resolved jointly; and (ii) *sequential* updates, where walkers are updated one-by-one (in a fixed or random order) with occupancies updated after each move. Sequential timing often better matches queue-like resolution and can reduce deadlocks under strict exclusion.

Three interacting families are now described. Once timing and tie-breaking conventions are fixed, each rule defines a Markov kernel $\mathbf{P}_{\mathrm{int}}(\mathbf{x}, \mathbf{x}')$ on $(\mathbb{Z}_d)^L$.

Blocking prevents moves into occupied sites. In the reject-to-stay variant (synchronous form),

$$x'_\ell = \begin{cases} y_\ell, & \text{if } y_\ell \notin \mathcal{O}(\mathbf{x}) \setminus \{x_\ell\}, \\ x_\ell, & \text{otherwise.} \end{cases}$$

A common alternative replaces rejection by rerolling: if y_ℓ is blocked, the walker resamples its increment (and, under constrained teleportation, its teleport destination) up to a maximum number of attempts M_{re}, defaulting to x_ℓ if no admissible site is found. Rerolling increases mobility relative to strict rejection, but induces state dependence because admissible proposals depend on $N_{\mathbf{x}}$.

Bouncing resolves attempted co-occupancy by reflecting or permuting motion so that walkers do not pass through each other. A basic reflection rule applies to increment proposals: if $y_\ell = x_\ell + X_\ell$ is occupied, replace the move by a reflected step

$$x'_\ell = x_\ell - X_\ell \pmod{d},$$

provided the reflected site is unoccupied; otherwise the walker stays put or triggers a secondary mechanism (e.g., rerolling). Under synchronous updates, a natural alternative is a swap-or-cycle rule: if proposals form directed cycles (e.g., two walkers propose to move into each other's current sites), execute the cycle permutation. This preserves motion while avoiding co-occupancy and parallels standard synchronous interacting-particle updates.

Pushing allows a walker to enter an occupied site by displacing the occupant, potentially triggering a cascade. In a forward-cascade variant, if walker ℓ proposes to enter an occupied site, the occupant is pushed forward (either to its own proposed destination or via a prescribed local displacement), and the push propagates until a vacant site is reached. On a finite cycle, cascades can wrap around, so a convention is required: under truncated pushing, a full wrap-around (no vacancy) cancels the cascade; under compressing pushing, multiple occupancy is allowed but an internal ordering is maintained, yielding a queueing-type model. A simpler single-step push displaces the occupant forward by one site if that site is vacant and otherwise rejects the move, capturing local pressure while retaining a partial exclusion constraint.

All rules above are well posed once timing and tie breaking are specified. In later sections we focus on permutation-invariant rules and on regimes with $\alpha > 0$ and unconstrained teleportation, where irreducibility is typically retained because teleportation provides access to broad regions of the configuration space.

5.2 Loss of translation invariance and a perturbation viewpoint

In the independent model, the joint kernel factorizes and depends on $\mathbf{x}$ only through coordinate-wise differences, yielding invariance under diagonal shifts and a circulant structure on $(\mathbb{Z}_d)^L$. Under interactions, executed moves depend on the global occupancy profile $N_{\mathbf{x}}(\cdot)$ and (under synchronous timing) on the joint proposal set, so in general

$$\mathbf{P}_{\mathrm{int}}(\mathbf{x}, \mathbf{x}') \text{ is not a function of } (\mathbf{x}' - \mathbf{x}) \text{ alone,}$$

and Fourier diagonalization on the product group no longer applies directly.

Accordingly, the circulant baseline is taken to be

$$\mathbf{P}_0 := \mathbf{P}_{\mathrm{ind}} = \big(P^{(\alpha,s)}\big)^{\otimes L}$$

as a reference and treat $\mathbf{P}_{\mathrm{int}}$ as a perturbation of $\mathbf{P}_0$. A robust measure of interaction strength is the worst-case one-step total-variation discrepancy

$$\Delta_{\mathrm{TV}} := \sup_{\mathbf{x}} \big\|\mathbf{P}_{\mathrm{int}}(\mathbf{x},\cdot) - \mathbf{P}_0(\mathbf{x},\cdot)\big\|_{\mathrm{TV}}. \tag{33}$$

Proposition 5 (A quantitative perturbation bound for reject-to-stay blocking) *Consider the* reject-to-stay blocking *rule under synchronous updates with* unconstrained *teleportation (as defined in Section 5.1). Let* $\mathbf{P}_{\mathrm{int}}$ *be the resulting interacting kernel and* $\mathbf{P}_0 = \mathbf{P}_{\mathrm{ind}}$ *the baseline product kernel. Write* $p_{\max} := \max_{x\in\mathbb{Z}_d} \mu(x)$. *Then the one-step perturbation size* (33) *satisfies*

$$\Delta_{\mathrm{TV}} \le \min\Big\{1,\ L(L-1)\Big(\frac{\alpha}{d} + (1-\alpha)(1-s)\,p_{\max}\Big)\Big\}.$$

Moreover, under the stationary product law $\boldsymbol{\pi}$ *(i.i.d. uniform positions),*

$$\mathbb{E}_{\boldsymbol{\pi}}\big[\big\|\mathbf{P}_{\mathrm{int}}(\mathbf{R}_0,\cdot) - \mathbf{P}_0(\mathbf{R}_0,\cdot)\big\|_{\mathrm{TV}}\big] \le L(L-1)\Big(\frac{\alpha}{d} + (1-\alpha)(1-s)\,\frac{1}{d}\Big).$$

Proof sketch Under reject-to-stay blocking, the executed next configuration differs from the baseline proposal configuration only if *some* walker proposes to move into a site occupied by another walker at the current time. Thus the mismatch probability under the natural coupling "baseline proposals + blocking resolver" upper-bounds the total-variation discrepancy. A union bound over ordered pairs of walkers gives

$$\mathbb{P}(\text{some proposal targets another's current site}) \le \sum_{i\neq j} \mathbb{P}(y_i = x_j).$$

If walker i teleports then $\mathbb{P}(y_i = x_j) = 1/d$. If it performs an increment move then $\mathbb{P}(y_i = x_j) = \mu(x_j - x_i) \le p_{\max}$. Combining these cases and summing over $i \neq j$ yields the stated bound. Averaging under $\boldsymbol{\pi}$ replaces $\mu(x_j - x_i)$ by its mean value $1/d$. □

Remark 3 The worst-case bound can be loose when μ is highly concentrated (e.g. deterministic unit steps), because an adversarial configuration can force many proposals to target occupied sites. The stationary-average bound is more indicative in sparse, well-mixed regimes where $x_j - x_i$ is close to uniform.

By maximal coupling, Δ_{TV} is the largest (over starting states) achievable probability that one interacting step differs from one baseline step under an optimal coupling. This interpretation links directly to conflict probabilities: for blocking/bouncing, discrepancies are driven primarily by proposed destinations that are occupied; for pushing, by proposed moves that trigger cascades. In sparse regimes $L \ll d$, such conflicts are rare, so Δ_{TV} is typically small.

When additional structure is available (e.g., reversibility), one may also compare kernels in operator norms on mean-zero functions (such as $\|\mathbf{P}_{\mathrm{int}} - \mathbf{P}_0\|_{2\to 2}$ under a

reference measure), but the total-variation/coupling viewpoint applies without such assumptions.

5.3 Approximation principles for mixing and collision times

This subsection outlines how baseline predictions for mixing and collision-related stopping times can be transferred to interacting chains. Couple $\mathbf{P}_{\mathrm{int}}$ and $\mathbf{P}_0$ step-by-step using maximal couplings: construct $(\mathbf{R}_n^{\mathrm{int}}, \mathbf{R}_n^0)_{n\geq 0}$ so that, conditional on $\mathbf{R}_{n-1}^{\mathrm{int}} = \mathbf{R}_{n-1}^0 = \mathbf{x}$, the mismatch probability equals $\|\mathbf{P}_{\mathrm{int}}(\mathbf{x},\cdot) - \mathbf{P}_0(\mathbf{x},\cdot)\|_{\mathrm{TV}}$. Let

$$\tau := \inf\{n \geq 1 : \mathbf{R}_n^{\mathrm{int}} \neq \mathbf{R}_n^0\}$$

be the first mismatch time. Then, for any path-dependent stopping time T (e.g., first collision), the coupled chains agree up to τ, and hence

$$\mathbb{P}(T_{\mathrm{int}} \neq T_0) \leq \mathbb{P}(\tau \leq T_0).$$

Thus, when Δ_{TV} is small and the relevant baseline time scale T_0 is not too large, one expects $\tau > T_0$ with high probability, so baseline collision-time predictions remain accurate up to a controllable error. Conversely, if an interaction systematically suppresses or amplifies conflicts, (33) provides a principled way to quantify how far encounter statistics can deviate from the baseline.

Teleportation strengthens these comparisons by repeatedly refreshing the configuration. Under unconstrained teleportation, interacting chains often retain uniform minorization events (e.g., an all-teleport step), implying uniform ergodicity and stabilizing mixing profiles under perturbations [26]. The same refresh mechanism supports block arguments for collision times: any uniform lower bound on collision probability over a fixed block yields geometric tails for first-collision times, and teleportation frequently supplies such bounds even when local interactions introduce strong state dependence.

Finally, in low-density regimes $L \ll d$, mean-field heuristics provide a simple effective-parameter picture. Under near-uniform occupancy, a typical proposed destination is occupied with probability $\asymp (L-1)/d$, so blocking-type rules primarily act by increasing effective laziness. For reject-to-stay blocking, one is led to the approximation

$$s_{\mathrm{eff}} \approx s + (1-s)\frac{L-1}{d},$$

and analogous effective-parameter heuristics can be formulated for bouncing (via reflections) and pushing (via short cascades). These approximations complement the coupling-based discrepancy control by indicating when baseline spectral predictions remain quantitatively reliable.

6 Numerical Validation

This section evaluates the theoretical quantities derived for the FJRW–LT model on the directed cycle $\mathbb{Z}_d$ by comparing closed-form expressions, rigorous bounds, and

occupancy-based approximations with Monte Carlo simulations. Unless otherwise stated, walkers are updated synchronously, the forward increment law is the unit step $\mu(1) = 1$, and the laziness parameter is fixed at $s = 0.30$. Teleportation occurs independently with probability α at each step, as specified in Section 2.1. The numerical experiments are designed to assess three aspects of the theory: the accuracy of the Fourier formulas in the two-walker displacement chain, the quality of occupancy approximations for many-walker first collisions, and the effect of interaction resolvers on collision and tail behavior.

The simulations are vectorized over independent trials. Whenever a finite horizon $n_{\max}$ is imposed, trajectories that do not hit the relevant target before the horizon are treated as right-censored observations, and censoring rates are monitored. In the reported parameter regimes, censoring is negligible unless explicitly indicated.

6.1 Simulation protocol and uncertainty quantification

All Monte Carlo experiments in this section simulate the discrete-time Markov dynamics exactly as defined in Section 2. For each parameter setting, the number of independent trials is reported explicitly, for example as N_{meet} or N_{coll}. Uncertainty is assessed in three complementary ways. First, sample means and variances are accompanied by standard errors across independent trials; when censoring is negligible, uncertainty may be displayed as $\pm 1.96\,\widehat{\mathrm{SE}}$ under a normal approximation. Second, for tail probabilities and survival curves, $\mathbb{P}(T > n)$ is estimated by the binomial proportion $\widehat{\mathbb{P}}(T > n)$ at each time n, with Wilson or Clopper–Pearson intervals available for pointwise confidence bands. Third, reproducibility is supported by fixing and reporting random-number-generator seeds, simulation horizons, burn-in choices, and censoring rules. If a finite horizon is used, the fraction of right-censored trials should also be reported.

6.2 Two-walker displacement chain

Lemmas 2–3 and Theorems 2, 3–4 are validated in the fixed small case $(d, \alpha, s) = (50, 0.05, 0.30)$. Two independent walkers are simulated with $(R_0^{(1)}, R_0^{(2)}) = (\delta, 0)$ so that the displacement $D_n := R_n^{(1)} - R_n^{(2)} \pmod d$ satisfies $D_0 = \delta$. The meeting time is $T_{1,2}^{(\delta)} := \min\{n \geq 1 : D_n = 0\}$.

For arriving distributions, $N_{\text{arrive}} = 1.2 \times 10^5$ replicates are used for each (δ, n) with $n \in \{1, 5, 10, 25\}$ and $\delta \in \{1, \lfloor d/4 \rfloor, \lfloor d/2 \rfloor\}$. Figure 1 compares the empirical distribution of D_n (markers) to the closed-form arriving distribution (5) from Theorem 2 (curve). The early-time panels show strong structure induced by the forward drift/laziness mixture, while later-time panels flatten toward uniformity as teleportation repeatedly refreshes the displacement. Across all (δ, n) shown, the Fourier-series prediction (5) captures both the location of the dominant peaks (short times) and the rate at which mass spreads (longer times).

For meeting-time moments, $\delta \in \{1, \ldots, \lfloor d/2 \rfloor\}$ is considered, and $\mathbb{E}[T_{1,2}^{(\delta)}]$ and $\mathrm{Var}(T_{1,2}^{(\delta)})$ are estimated from $N_{\text{meet}} = 5 \times 10^4$ replicates per δ, using a generous horizon $n_{\max} = \lceil 200 \max\{d^2, 1/\alpha\} \rceil$ to eliminate censoring. Figure 2 compares Monte Carlo (markers) to the closed-form expressions (6) and (7) (curves). Numerically, both mean

and variance rise from small δ and rapidly approach a plateau. This is consistent with the "mix-then-hit" picture: once the displacement mixes, the remaining time to hit 0 is governed primarily by the stationary rare-event scale rather than the precise initial separation, exactly the regime encoded by the spectral denominators in (6)–(7).

6.3 Many-walker first collision

First collisions are examined for $L \geq 2$ under the same local kernel ($\mu(1) = 1$, $s = 0.30$) at teleport rate $\alpha = 0.05$ on

$$d \in \{50, 100, 200\}, \qquad L \in \{4, 8, 16\} \cap \{1, \ldots, d\}.$$

The simulations record T_{coll} as defined in Section 2.4, using a randomly shifted clustered start $R_0^{(i)} = U + (i-1) \pmod{d}$ to avoid time-0 collisions while preserving translation symmetry. For each (d, L), $N_{\text{coll}} = 3 \times 10^4$ replicates are simulated.

The comparison relates Monte Carlo survival curves to the well-mixed geometric/occupancy approximation described in Section 3.4(iii), with stationary collision probability given by (13). Figure 3 shows that the approximation is least accurate at very small n (where the clustered initialization creates transient correlations), and becomes increasingly accurate once the process has effectively refreshed toward stationarity (consistent with the mixing-based perspective in (12) and the teleportation-based block bounds (27)). Across the grid of (d, L), the approximation captures the correct scale and the correct rate once early-time transients subside.

The composite Figure 4 collects the first-collision and coalescence diagnostics that were previously displayed separately. Panel (A) summarizes $\mathbb{E}[T_{\text{coll}}]$ as L varies for each $d \in \{50, 100, 200\}$. Two baselines from the theory section are overlaid: (i) the occupancy predictor built from (13) (Section 3.4(iii)), and (ii) the sparse-regime heuristic scaling (15). Numerically, the occupancy predictor tracks the Monte Carlo means closely on this grid, indicating that first collision is well-approximated as a near-stationary rare event after refresh. The sparse heuristic (15) remains accurate in trend and interpretable, especially when L is not too large relative to $\sqrt{d}$.

For the representative case $(d, L) = (50, 8)$ (clustered start), $\mathbb{P}(R_n \in \mathcal{C})$ is estimated at selected times n using $N_{\text{pn}} = 6 \times 10^4$ replicates, and the estimates are compared with the pairwise union bound (9). Pairwise equal-location probabilities are computed via the two-walker displacement spectrum using Theorem 2 (specialized to $r = 0$). As shown in panel (B) of Figure 4, the union bound is conservative and can be loose at early times (clustered start), but it correctly captures the relaxation toward the stationary collision level described by (13).

6.4 Coalescence and full-diagonal hitting

This subsection evaluates the coalescing (merge-on-meet) dynamics from Section 3.5 and the non-coalescing all-equal diagonal hitting observable. Throughout this subsection, the same kernel parameters are retained, with $\alpha = 0.05$.

The simulations record the successive merge times τ_k (Section 3.5) and compare Monte Carlo means to the occupancy-based scaling that appears in the phase bound

(16). Panel (C) of Figure 4 shows a robust qualitative feature: merges occur very quickly when many clusters remain (large k), and the final few merges dominate the total time. This is exactly what (16) predicts: the relevant collision probability decreases with smaller k, so the expected waiting time per phase increases. Quantitatively, Monte Carlo typically exceeds the pure occupancy prediction in the last few phases, consistent with residual dependence between the surviving cluster locations (i.e., the well-mixed assumption underlying the occupancy term becomes least accurate precisely when only a few clusters remain).

The analysis next examines $\mathbb{E}[T_{\text{coal}}]$ across $d \in \{50, 100, 200\}$ and $L \in \{4, 8, 16\}$. Panel (D) of Figure 4 compares Monte Carlo means to the phase-sum structure in (17). The observed scaling matches the qualitative predictions encoded there: increasing d slows coalescence (clusters must search a larger space to meet), while increasing L increases the number of required merge phases. Numerically, Monte Carlo is consistently above the simplest well-mixed baseline, which aligns with the deviation already visible in panel (C) for the final phases (the sum is dominated by those phases).

Finally, the non-coalescing all-equal hitting time T_{all} defined in Section 3.5(A) is simulated. Panel (E) of Figure 4 confirms the strong growth of $\mathbb{E}[T_{\text{all}}]$ with L and the fact that this observable is governed by an extreme small-target event (as discussed immediately after the definition of T_{all} in Section 3.5). In particular, the plot verifies that the dominant scale is captured by the occupancy picture, while constant-factor discrepancies arise from short-time dependence and imperfect step-to-step independence.

Figure 5 plots the one-time collision probability $\mathbb{P}(R_n \in \mathcal{C})$ as a function of time n for several teleport rates, comparing two teleport-coupling variants: *independent teleports* (each walker has its own teleport clock) and a *shared teleport clock* (a single global teleport indicator applied to all walkers). Across regimes, the Monte Carlo estimates (markers) exhibit a strong early-time transient from the clustered initialization: the collision probability starts near one (many walkers are close and quickly pile up under the local drift/laziness dynamics) and then relaxes downward toward the stationary collision level $\pi(\mathcal{C})$ (horizontal line), which is the uniform-occupancy benchmark implied by (13). Increasing α accelerates this relaxation, consistent with the interpretation that teleportation rapidly refreshes the configuration toward the uniform stationary law. The mixing-based lower baseline $\pi(\mathcal{C}) - \mathrm{TV}(\cdot, \pi)$ (dashed) becomes informative only once total-variation distance is small, while the pairwise union upper bound (dashed) is conservative, especially during the clustered transient.

Figure 6 validates the survival function of the first collision time T_{coll} across three representative teleport regimes (high/cross/low). The geometric (occupancy) approximation $(1-\pi(\mathcal{C}))^{\,n}$ (curve) corresponds to the "well-mixed rare-event" heuristic in which each step is treated as an (approximately) independent draw from stationarity. The Monte Carlo survival curves (markers) deviate most at very small n due to the clustered initialization and short-time correlations, but the decay rate becomes increasingly consistent with the geometric prediction once teleportation-driven refresh places the system near stationarity. The one-step and block-type exponential bounds remain conservative throughout, with visible looseness when early-time dependence dominates.

Figure 7 summarizes Monte Carlo estimates of $\mathbb{E}[T_{\text{coll}}]$ (markers) for $d \in \{50, 100, 200\}$ and $L \in \{4, 8, 16\}$ across the three regimes $\alpha \in \{1/(10d^2),\ 1/d^2,\ 0.05\}$. Two stationary baselines are overlaid: the occupancy predictor $1/\pi(\mathcal{C})$ and the sparse heuristic $2d/(L(L-1))$ (Section 3.4(iii)). A consistent pattern is that d and L primarily set the scale, while α has a comparatively milder effect on the mean on this grid—because first collision often occurs during the early transient before full mixing, especially for smaller L under clustered initialization. As L increases (or when teleportation is strong enough to rapidly refresh), the well-mixed occupancy benchmark becomes more appropriate, and the sparse heuristic is most accurate in the truly sparse range (roughly $L^2 \ll d$), where collisions are driven by pairwise encounters.

The same simulation framework extends to alternative interaction resolvers (blocking, bouncing, pushing) from Section 5. In that setting proposal-modification frequencies are additionally recorded and compared with baseline collision rates (e.g., via (9) and stationarity levels from (13)), which provides a direct quantitative proxy for interaction strength.

Figures 8–10 compare the baseline (no interaction) multi-walker FJRW–LT dynamics with three collision-resolution mechanisms—*blocking* (synchronous), *bouncing* (synchronous), and *pushing* (sequential)—implemented through a proposal–resolution split. At each step, a full proposal configuration is first generated; a resolver is then applied to edit only the conflicting components (e.g., converting moves to "stay put", reflecting, or applying cascading pushes). Results are shown for three representative teleport regimes, `low`/`cross`/`high`, over a grid of (d, L) values. In Fig. 8, the dashed horizontal reference denotes the well-mixed occupancy benchmark $1/\pi(\mathcal{C})$.

Figure 8 shows that interaction resolvers can inflate the mean first-collision time $\mathbb{E}[T_{\text{coll}}]$ by orders of magnitude relative to the baseline, most notably in the diffusion-dominant and crossover regimes. In `low` and `cross`, blocking and pushing strongly suppress co-occupancy generated by local drift/laziness moves, so collisions occur primarily through comparatively rare refresh events; accordingly, $\mathbb{E}[T_{\text{coll}}]$ can lie far above the occupancy scale $1/\pi(\mathcal{C})$ (dashed line). As teleportation increases toward `high`, repeated refresh re-randomizes positions and re-introduces collisions closer to the i.i.d. occupancy picture, narrowing the gap between interacting rules and the benchmark. Bouncing typically produces an intermediate effect: it reduces collisions relative to baseline, but—because it preserves more motion via reflections/permutations—it does not suppress encounters as aggressively as strict rejection (blocking) or long cascades (pushing), leaving $\mathbb{E}[T_{\text{coll}}]$ closer to the occupancy scale.

Figure 9 reports how often the resolver actively changes a proposed step, which empirically tracks how strongly the interacting kernel departs from the baseline product kernel. Modification rates are normalized to the corresponding baseline reference on the same (d, L, α) grid, so 100% represents the baseline level. Blocking and pushing exhibit the smallest normalized modification rates in sparser configurations (larger d and smaller L) and in regimes where refresh quickly disperses walkers, consistent with a "rare-interaction" perturbative picture. In contrast, bouncing can induce substantially more frequent edits—especially under synchronous timing—indicating that the resolution mechanism is invoked often enough that the interacting evolution is not merely a small perturbation of the baseline dynamics.

Figure 10 plots the empirical survival function $\mathbb{P}(T_{\mathrm{coll}} > t)$ for the same family of resolvers. Across methods, the tails decay approximately geometrically/exponentially at moderate and large t, consistent with teleportation-driven minorization events that enforce eventual mixing and yield exponential tail control. Interaction rules primarily act as a horizontal shift of these survival curves: blocking and pushing produce markedly heavier tails (collisions are delayed), while bouncing yields a smaller but still visible delay relative to baseline. The dashed reference curve provides a conservative geometric-rate prediction/bound based on the well-mixed occupancy heuristic and teleportation refresh; discrepancies at very small times reflect transient clustering and short-time dependence, with improved agreement after repeated refresh moves the system closer to stationarity.

7 Application Contexts

The preceding analysis is motivated by encounter-driven systems whose state space is naturally cyclic or ring-like. In such systems, several agents, packets, tokens, or tasks evolve on a ring under a combination of local motion, waiting, and occasional reinitialization. The FJRW–LT parameters have direct operational interpretations: the increment law μ describes local progression or scheduled movement, the laziness parameter s represents idling or backoff, and the teleportation parameter α represents refresh, reseeding, redeployment, or churn. The diagonal hitting times studied above therefore quantify operational events such as contention, rendezvous, aggregation, stabilization, and congestion.

7.1 Token rings and randomized mutual exclusion

A direct application is token-based mutual exclusion on a directed ring. Consider d processes arranged on $\mathbb{Z}_d$, where a circulating token grants permission to enter a critical section. The legal configuration contains exactly one token, whereas transient faults may create multiple tokens through duplicated messages, corrupted local states, or inconsistent recovery. Representing each token as a walker converts duplicate-token elimination into a coalescence problem. The first collision time T_{coll} describes the first contention event, and the coalescence time T_{coal} describes the time required to reduce the system to a single token when meeting tokens are merged. This interpretation is consistent with classical probabilistic analyses of token management and self-stabilizing mutual exclusion [27–30].

Within this interpretation, teleportation represents protocol-level recovery mechanisms such as timeout-driven reseeding, randomized reinitialization, or re-entry after local failure. Increasing α generally improves mixing and tail control, but may also increase short-run contention by injecting tokens into uniformly selected locations. The bounds in Sections 3.4 and 4.2 therefore provide quantitative tools for evaluating stabilization–contention trade-offs.

7.2 Rendezvous, gossip, and encounter-driven dissemination

In intermittently connected networks, communication may occur only when two agents occupy the same location or communication range. Under this constraint, the delay to transmit information between two designated agents is a meeting time, while dissemination to a population depends on repeated collision events. This setting is closely related to epidemic protocols, randomized rumor spreading, and gossip algorithms for distributed aggregation and averaging [17, 18, 31, 32].

The multi-walker FJRW–LT model provides a concise abstraction for such systems. Mobile devices, robots, or vehicles moving around a cyclic route are represented by walkers on $\mathbb{Z}_d$. When several walkers co-locate, a local broadcast or data exchange is performed. If $\inf_i(n) \in \{0, 1\}$ denotes whether agent i is informed, a dissemination time can be defined by

$$T_{\text{spread}} := \inf\{n \geq 0 : \ \inf_i(n) = 1 \text{ for all } i = 1, \ldots, L\}.$$

Meeting and collision estimates then describe the latency of the first contact, while coalescence-type quantities describe aggregation or consensus processes. Teleportation models redeployment, intermittent infrastructure assistance, or deliberate randomized reseeding, and the exponential tail bounds obtained above quantify the reliability benefits of such refresh mechanisms.

7.3 Congestion and constrained interaction rules

The interaction rules introduced in Section 5 are also relevant to congestion-controlled systems. In applications such as cyclic service routes, warehouse loops, processor rings, and overlay networks, two agents may be prohibited from occupying the same site or may be forced to resolve conflicts through rejection, reflection, or displacement. The blocking, bouncing, and pushing mechanisms formalize these operational constraints. Although such rules generally destroy exact product structure and Fourier diagonalization, the independent FJRW–LT model remains a useful analytically tractable baseline.

When the population is sparse relative to the cycle size, conflicts occur infrequently and the interacting dynamics can often be treated as perturbations of the independent model. In denser regimes, the empirical modification rate of the resolver provides a practical measure of how strongly the interacting kernel departs from the baseline. The numerical results in Section 6 indicate that interaction rules mainly shift collision-time distributions and survival curves by delaying or suppressing co-occupancy, whereas teleportation preserves eventual refresh and hence exponential tail control.

8 Conclusion

This paper has developed a tractable framework for multi-walker forward-jump random walks on a finite cycle with laziness and teleportation. The product-group formulation on $(\mathbb{Z}_d)^L$ represents meetings, collisions, and coalescence as diagonal hitting problems and uses diagonal-shift symmetry to reduce the role of absolute position. This perspective provides a unified method for studying encounter-driven observables in finite cyclic systems.

The principal baseline results are obtained through harmonic analysis on the finite abelian group $(\mathbb{Z}_d)^L$. For independent walkers, the L-fold product kernel is diagonalized explicitly, giving a finite Fourier-series formula for the time-n arriving distribution. In the two-walker case, the displacement-chain reduction yields closed-form Fourier-sum expressions for the mean and variance of the meeting time under arbitrary initial offsets. For $L \geq 3$, exact diagonal-hitting formulas are generally less tractable; accordingly, the paper develops computable bounds and occupancy-based approximations for first collision and for phase-wise coalescence.

Teleportation serves as the main mechanism organizing the qualitative behavior of the model. Through Doeblin-type mixture identities, it provides explicit total-variation contraction and hence geometric or exponential tail control for collision and coalescence observables. The resulting regime classification—teleport-dominant, diffusion-dominant, and crossover—clarifies how the parameters (α, s, μ) jointly determine encounter frequency and stabilization scales on the cycle.

The framework was also extended beyond the independent baseline by incorporating common-shock couplings and congestion-motivated interaction rules, including blocking, bouncing, and pushing. Because these mechanisms typically break exact circulant structure, perturbative and coupling-based interpretations were used to relate the interacting dynamics to the analytically tractable product chain. Monte Carlo experiments support the explicit Fourier formulas, the occupancy approximations, and the qualitative predictions for interaction-induced delays.

Several directions remain open. Further work could address heterogeneous walkers, nonuniform or state-dependent teleportation, sharper large-d asymptotics, refined crossover estimates, non-cyclic geometries, and systematic error bounds for interacting dynamics outside sparse-conflict regimes. Such extensions would broaden the applicability of the framework to encounter-driven systems in which meeting, collision, and coalescence times function as operational measures of latency, contention, and risk.

A Derivation of the meeting-time formulas

This appendix provides the generating-function/renewal details leading to Theorems 3 and 4. Throughout, fix d and parameters (α, s, μ) satisfying Assumption 1, and let (D_n) be the displacement chain with eigenvalues $\Gamma_k = |\Lambda_k|^2$.

A.1 Renewal identity and generating functions

For $\delta \in \mathbb{Z}_d$ define the return probabilities

$$a_n^{(\delta)} := \mathbb{P}(D_n = 0 \mid D_0 = \delta), \qquad n \geq 0,$$

and the (first-)hitting-time distribution

$$f_n^{(\delta)} := \mathbb{P}(T_{1,2}^{(\delta)} = n), \qquad n \geq 1,$$

where $T_{1,2}^{(\delta)} = \min\{n \geq 1 : D_n = 0\}$. By conditioning on the *first* time the chain hits 0, for every $n \geq 1$ and every δ we have the renewal relation

$$a_n^{(\delta)} = \sum_{t=1}^{n} f_t^{(\delta)}\, a_{n-t}^{(0)}. \tag{34}$$

Introduce generating functions

$$A_\delta(z) := \sum_{n\geq 0} a_n^{(\delta)} z^n, \qquad F_\delta(z) := \sum_{n\geq 1} f_n^{(\delta)} z^n.$$

Multiplying (34) by z^n and summing over $n \geq 1$ yields

$$A_\delta(z) = \mathbf{1}\{\delta = 0\} + F_\delta(z)\, A_0(z), \qquad \text{so in particular} \qquad F_\delta(z) = \frac{A_\delta(z) - \mathbf{1}\{\delta = 0\}}{A_0(z)}. \tag{35}$$

Because $a_n^{(\delta)} \to 1/d$ as $n \to \infty$, the power series $A_\delta(z)$ has a simple pole at $z = 1$. It is convenient to cancel this pole by defining the *regularized* series

$$B_\delta(z) := (1 - z)A_\delta(z). \tag{36}$$

Then $B_\delta(z)$ is analytic at $z = 1$, and (35) becomes

$$F_\delta(z) = \frac{B_\delta(z) - (1 - z)\mathbf{1}\{\delta = 0\}}{B_0(z)}. \tag{37}$$

For $\delta \neq 0$, this simplifies to $F_\delta(z) = B_\delta(z)/B_0(z)$.

A.2 Closed forms for B_δ and derivatives at $z = 1$

From Theorem 2 with $r = 0$ we have, for $n \geq 1$,

$$a_n^{(\delta)} = \frac{1}{d} + \frac{1}{d}\sum_{k=1}^{d-1} \omega^{\delta k} \Gamma_k^n, \qquad \omega := e^{2\pi i/d}.$$

Thus

$$\begin{aligned} A_\delta(z) &= \mathbf{1}\{\delta = 0\} + \sum_{n\geq 1} \frac{1}{d} z^n + \frac{1}{d}\sum_{k=1}^{d-1} \omega^{\delta k} \sum_{n\geq 1} (\Gamma_k z)^n \\ &= \mathbf{1}\{\delta = 0\} + \frac{z}{d(1 - z)} + \frac{1}{d}\sum_{k=1}^{d-1} \omega^{\delta k} \frac{z\Gamma_k}{1 - z\Gamma_k}. \end{aligned} \tag{38}$$

Multiplying by $(1-z)$ gives

$$B_\delta(z) = (1-z)\mathbf{1}\{\delta = 0\} + \frac{z}{d} + \frac{1}{d}\sum_{k=1}^{d-1} \omega^{\delta k}\, h_k(z), \qquad h_k(z) := (1-z)\frac{z\Gamma_k}{1-z\Gamma_k}. \quad (39)$$

A direct calculation gives $h_k(1) = 0$ and

$$h_k'(1) = -\frac{\Gamma_k}{1-\Gamma_k}, \qquad h_k''(1) = -\frac{2\Gamma_k}{(1-\Gamma_k)^2}. \quad (40)$$

Since $B_\delta(1) = 1/d$ for every δ, differentiating (39) at $z = 1$ yields

$$B_\delta'(1) = -\mathbf{1}\{\delta = 0\} + \frac{1}{d} - \frac{1}{d}\sum_{k=1}^{d-1} \omega^{\delta k}\frac{\Gamma_k}{1-\Gamma_k}, \quad (41)$$

$$B_\delta''(1) = -\frac{2}{d}\sum_{k=1}^{d-1} \omega^{\delta k}\frac{\Gamma_k}{(1-\Gamma_k)^2}. \quad (42)$$

A.3 Mean and variance

For $\delta \neq 0$ we have $F_\delta(z) = B_\delta(z)/B_0(z)$, hence

$$F_\delta'(1) = \frac{B_\delta'(1)B_0(1) - B_\delta(1)B_0'(1)}{B_0(1)^2} = d\big(B_\delta'(1) - B_0'(1)\big), \quad (43)$$

using $B_\delta(1) = B_0(1) = 1/d$. Substituting (41) into (43) gives

$$\mathbb{E}\Big[T_{1,2}^{(\delta)}\Big] = F_\delta'(1) = d + \sum_{k=1}^{d-1}\big(1-\omega^{\delta k}\big)\frac{\Gamma_k}{1-\Gamma_k},$$

which is Theorem 3. (For $\delta = 0$, $T_{1,2}^{(0)}$ is the first *return* time and $\mathbb{E}[T_{1,2}^{(0)}] = d$ by Kac's formula.)

To obtain the second moment, note that for any nonnegative integer-valued T with pgf $F(z) = \sum_{n\geq 1}\mathbb{P}(T=n)z^n$,

$$\mathbb{E}[T] = F'(1), \qquad \mathbb{E}[T(T-1)] = F''(1), \qquad \mathbb{E}[T^2] = F''(1) + F'(1).$$

Differentiating $F_\delta(z) = B_\delta(z)/B_0(z)$ twice and evaluating at $z = 1$ yields

$$F_\delta''(1) = \frac{B_\delta''(1)B_0(1) - B_\delta(1)B_0''(1)}{B_0(1)^2} - 2\,\frac{B_\delta'(1)B_0(1) - B_\delta(1)B_0'(1)}{B_0(1)^2}\cdot\frac{B_0'(1)}{B_0(1)}. \quad (44)$$

Substituting (41)–(42) into (44), and then using $\mathrm{Var}(T) = \mathbb{E}[T^2] - \mathbb{E}[T]^2$, yields, after algebraic simplification, the expression stated in Theorem 4. (The key simplifications

come from pairing complex conjugate modes k and $d-k$, and from the identity $\sum_{k=1}^{d-1} \omega^{\delta k} = -1$ for $\delta \neq 0$.)

B Cosine-series forms and numerical stability

The Fourier expressions in the main text are real-valued although they are written using complex exponentials. Pairing k with $d-k$ yields explicit cosine-series representations that are often more convenient numerically.

For $\delta \in \{0, 1, \ldots, d-1\}$ and any function $g(k)$ satisfying $g(d-k) = \overline{g(k)}$,

$$\sum_{k=1}^{d-1} \big(1 - \omega^{\delta k}\big) g(k) = 2 \sum_{k=1}^{\lfloor (d-1)/2 \rfloor} \big(1 - \cos(2\pi\delta k/d)\big)\, \Re g(k) + \mathbf{1}\{d \text{ even}\} \left(1 - (-1)^{\delta}\right) g(d/2).$$

Applying this with $g(k) = \Gamma_k/(1-\Gamma_k)$ gives a purely real form of Theorem 3; analogous manipulations apply to the variance formula. These real-valued forms reduce cancellation between conjugate terms and are often preferable when $\Gamma_k \approx 1$.

Declarations

Funding declaration: No funding was received for conducting this study.
Conflict of interest: The authors declare that they have no conflict of interest.

References

[1] Lovász, L.: Random walks on graphs: A survey. In: Combinatorics, Paul Erdős Is Eighty (Volume 2). Bolyai Society Mathematical Studies, vol. 2, pp. 1–46. János Bolyai Mathematical Society, Keszthely, Hungary (1993). https://users.wpi.edu/~yli15/courses/DS504Spring21/includes/RWSurvey.pdf

[2] Levin, D.A., Peres, Y.: Markov Chains and Mixing Times, 2nd edn. American Mathematical Society, Providence, RI (2017). With contributions by Elizabeth L. Wilmer

[3] Doyle, P.G., Snell, J.L.: Random Walks and Electric Networks. Carus Mathematical Monographs, vol. 22. Mathematical Association of America, Washington, DC (1984)

[4] Ikeda, Y., Fukai, Y., Mizoguchi, Y.: A property of random walks on a cycle graph. Pacific Journal of Mathematics for Industry **7**, 3 (2015) https://doi.org/10.1186/s40736-015-0015-3 . Article number 3

[5] Kaikeaw, R., Marupanthorn, P.: Forward jump random walk on a cycle graph and its hitting time. Science & Technology Asia **29**(1), 29–46 (2024) https://doi.org/10.14456/scitechasia.2024.3

[6] Brin, S., Page, L.: The anatomy of a large-scale hypertextual web search engine. Computer Networks and ISDN Systems **30**(1–7), 107–117 (1998) https://doi.org/10.1016/S0169-7552(98)00110-X

[7] Gleich, D.F.: Pagerank beyond the web. SIAM Review **57**(3), 321–363 (2015) https://doi.org/10.1137/140976649

[8] Haveliwala, T.H.: Topic-sensitive PageRank: A context-sensitive ranking algorithm for web search. IEEE Transactions on Knowledge and Data Engineering **15**(4), 784–796 (2003) https://doi.org/10.1109/TKDE.2003.1208999

[9] Andersen, R., Chung, F., Lang, K.: Local graph partitioning using PageRank vectors, 475–486 (2006)

[10] Tong, H., Faloutsos, C., Pan, J.-Y.: Random walk with restart: fast solutions and applications. Knowledge and Information Systems **14**, 327–346 (2008) https://doi.org/10.1007/s10115-007-0094-2

[11] Bootwisas, N., Marupanthorn, P., Kaikeaw, R.: Forward-jump random walks on a finite cycle with laziness and teleportation: Hitting times, mixing, and cutoff. SSRN Electronic Journal (2025) https://doi.org/10.2139/ssrn.5958799

[12] Oliveira, R.I.: Mean field conditions for coalescing random walks. The Annals of Probability **41**(5), 3420–3461 (2013) https://doi.org/10.1214/12-AOP813

[13] Oliveira, R.I., Peres, Y.: Random walks on graphs: new bounds on hitting, meeting, coalescing and returning. In: Proceedings of the Sixteenth Workshop on Analytic Algorithmics and Combinatorics (ANALCO 2019), pp. 119–126. SIAM, Philadelphia, PA (2019). https://doi.org/10.1137/1.9781611975505.13 . https://arxiv.org/abs/1807.06858

[14] Kanade, V., Mallmann-Trenn, F., Sauerwald, T.: On coalescence time in graphs: When is coalescing as fast as meeting? In: Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2019), pp. 956–965. SIAM, Philadelphia, PA (2019). https://doi.org/10.1137/1.9781611975482.59 . https://epubs.siam.org/doi/10.1137/1.9781611975482.59

[15] Rivera, N., Sauerwald, T., Sylvester, J.: Multiple random walks on graphs: mixing few to cover many. Combinatorics, Probability and Computing **32**(4), 594–637 (2023) https://doi.org/10.1017/S0963548322000372

[16] George, M., Patel, R., Bullo, F.: The Meeting Time of Multiple Random Walks (2018). https://arxiv.org/abs/1806.08843

[17] Kempe, D., Dobra, A., Gehrke, J.: Gossip-based computation of aggregate information. In: Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science (FOCS '03), pp. 482–491. IEEE Computer Society, Los

Alamitos, CA (2003). https://doi.org/10.1109/SFCS.2003.1238221

[18] Boyd, S., Ghosh, A., Prabhakar, B., Shah, D.: Randomized gossip algorithms. IEEE Transactions on Information Theory **52**(6), 2508–2530 (2006) https://doi.org/10.1109/TIT.2006.874516

[19] Oliva, G., Panzieri, S., Setola, R., Gasparri, A.: Gossip algorithm for multi-agent systems via random walk. Systems & Control Letters **128**, 34–40 (2019) https://doi.org/10.1016/j.sysconle.2019.04.009

[20] Cooper, C., Frieze, A., Radzik, T.: Multiple random walks in random regular graphs. SIAM Journal on Discrete Mathematics **23**(4), 1738–1761 (2009) https://doi.org/10.1137/080729542

[21] Sarwate, A.D., Dimakis, A.G.: The impact of mobility on gossip algorithms. IEEE Transactions on Information Theory **58**(3), 1731–1752 (2012) https://doi.org/10.1109/TIT.2011.2177753

[22] Redner, S.: A Guide to First-Passage Processes. Cambridge University Press, Cambridge (2001)

[23] Liggett, T.M.: Interacting Particle Systems. Classics in Mathematics. Springer, Berlin (2012). Reprint of the 1985 edition

[24] Aldous, D., Fill, J.A.: Reversible Markov Chains and Random Walks on Graphs. Available online. Unfinished monograph, recompiled 2014 (2002). https://www.stat.berkeley.edu/%7Ealdous/RWG/book.html

[25] Diaconis, P.: Group Representations in Probability and Statistics. IMS Lecture Notes–Monograph Series, vol. 11. Institute of Mathematical Statistics, Hayward, CA (1988)

[26] Meyn, S.P., Tweedie, R.L.: Markov Chains and Stochastic Stability, 2nd edn. Cambridge University Press, Cambridge (2009)

[27] Dijkstra, E.W.: Self-stabilizing systems in spite of distributed control. Communications of the ACM **17**(11), 643–644 (1974) https://doi.org/10.1145/361179.361202

[28] Dolev, S.: Self-Stabilization. MIT Press, Cambridge, MA (2000)

[29] Israeli, A., Jalfon, M.: Token management schemes and random walks yield self-stabilizing mutual exclusion. In: Proceedings of the Ninth Annual ACM Symposium on Principles of Distributed Computing (PODC '90), pp. 119–131. ACM, New York, NY (1990). https://doi.org/10.1145/93385.93409

[30] Herman, T.: Probabilistic self-stabilization. Information Processing Letters **35**(2), 63–67 (1990) https://doi.org/10.1016/0020-0190(90)90174-M

[31] Demers, A., Greene, D., Hauser, C., Irish, W., Larson, J.: Epidemic algorithms for replicated database maintenance. In: Proceedings of the Sixth Annual ACM Symposium on Principles of Distributed Computing (PODC '87), pp. 1–12. ACM, New York, NY (1987). https://doi.org/10.1145/41840.41841

[32] Karp, R.M., Schindelhauer, C., Shenker, S., Vöcking, B.: Randomized rumor spreading. In: Proceedings of the 41st Annual Symposium on Foundations of Computer Science (FOCS '00), pp. 565–574. IEEE Computer Society, Los Alamitos, CA (2000)

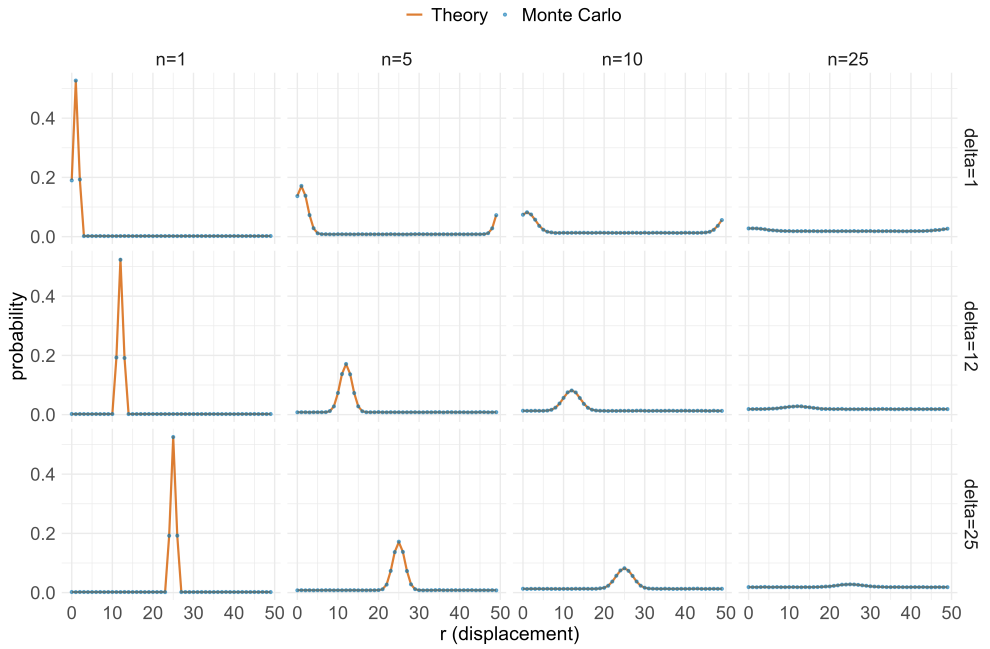


Fig. 1 Two-walker arriving distributions ($(d, \alpha, s) = (50, 0.05, 0.30)$). Distribution of the displacement D_n for $n \in \{1, 5, 10, 25\}$ and $\delta \in \{1, 12, 25\}$. Markers: Monte Carlo ($N_{\text{arrive}} = 1.2 \times 10^5$ per panel). Curve: (5) (Theorem 2).

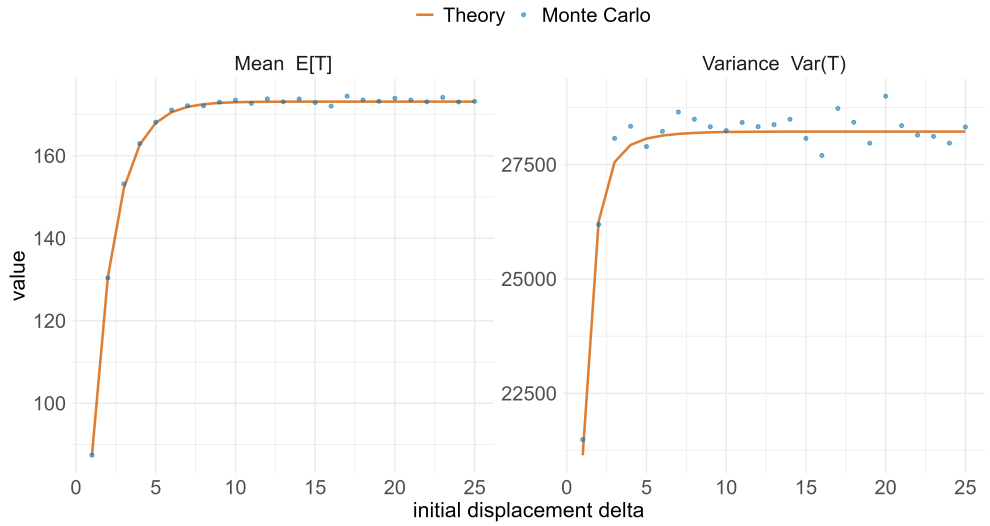


Fig. 2 Two-walker meeting-time moments ($(d, \alpha, s) = (50, 0.05, 0.30)$). Mean and variance of $T_{1,2}^{(\delta)}$ as a function of $\delta \in \{1, \ldots, 25\}$. Markers: Monte Carlo ($N_{\text{meet}} = 5 \times 10^4$ per δ). Curves: (6) and (7).

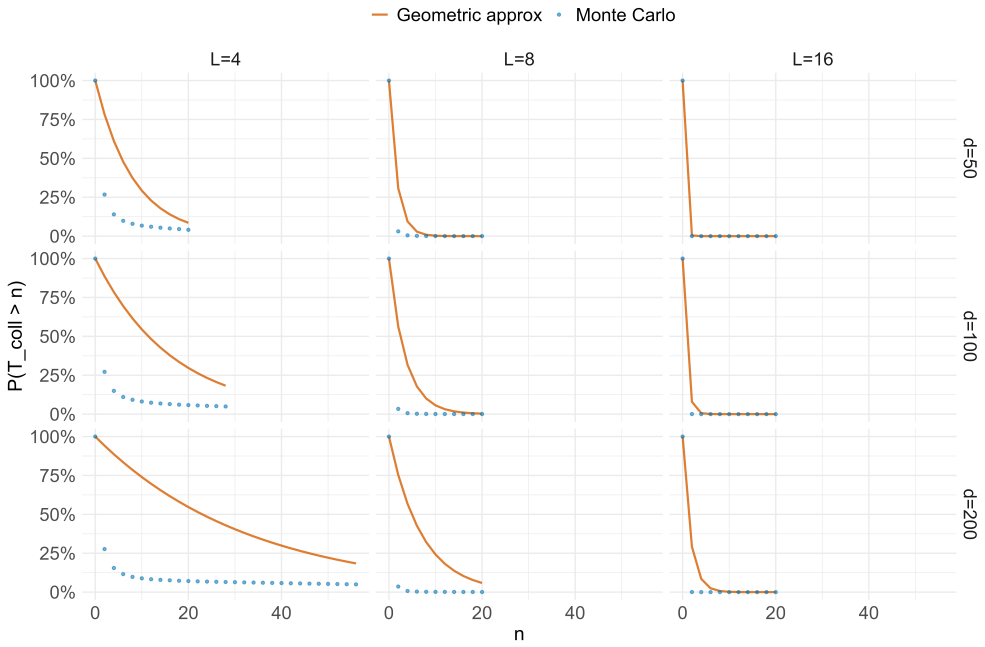


Fig. 3 Tail of T_{coll} across (d, L) at $\alpha = 0.05$, $s = 0.30$, $\mu(1) = 1$. Markers: Monte Carlo survival $\widehat{\mathbb{P}}(T_{\text{coll}} > n)$. Curve: occupancy-based geometric approximation from Section 3.4(iii) using (13).

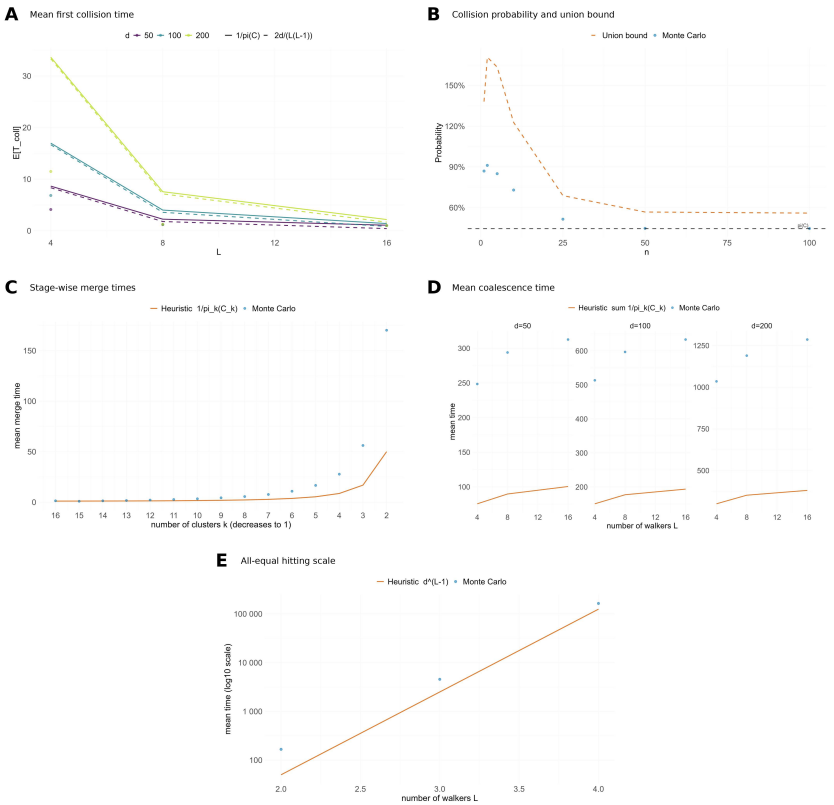


Fig. 4 Many-walker collision and coalescence diagnostics for the FJRW–LT model, redrawn in a consolidated landscape panel to improve readability. Panel (A) compares Monte Carlo estimates of $\mathbb{E}[T_{\mathrm{coll}}]$ with the occupancy predictor based on (13) and the sparse heuristic (15). Panel (B) compares $\mathbb{P}(R_n \in \mathcal{C})$ with the pairwise union bound (9) and the stationary level (13). Panels (C)–(E) summarize coalescing and full-diagonal hitting diagnostics: stage-wise merge times, total coalescence time, and the all-equal hitting scale.

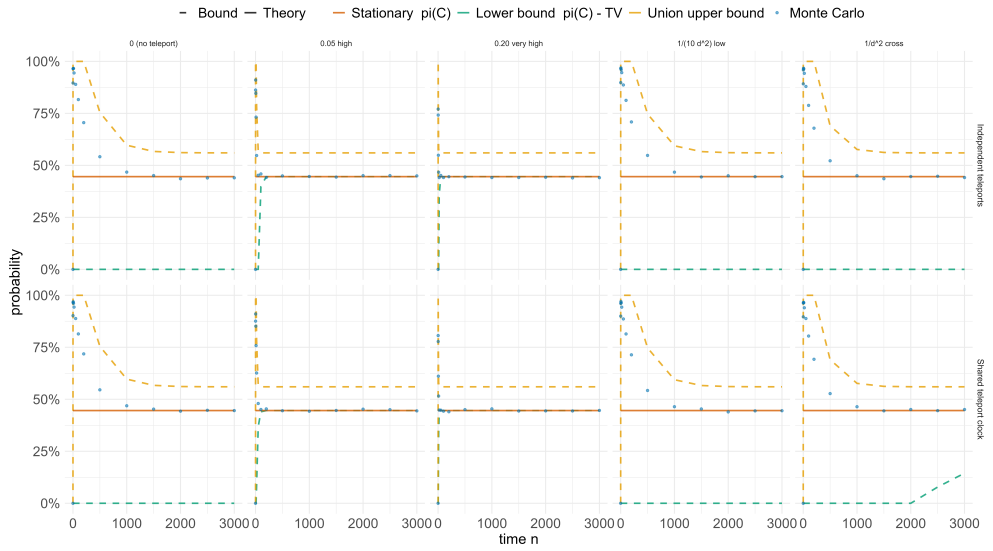


Fig. 5 Collision probability $\mathbb{P}(R_n \in \mathcal{C})$ vs. time n across teleport regimes, comparing independent teleports (top row) and a shared teleport clock (bottom row). Markers: Monte Carlo. Solid horizontal line: stationary collision probability $\pi(\mathcal{C})$ from (13), Dashed curves: mixing-based lower baseline and the pairwise union upper bound.

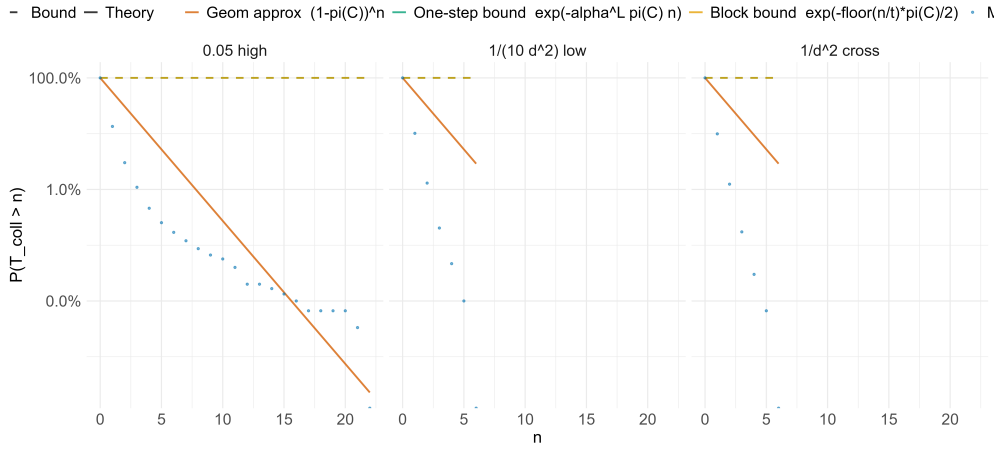


Fig. 6 Tail of T_{coll} across teleport regimes. Markers: Monte Carlo estimates of $\mathbb{P}(T_{\text{coll}} > n)$ (log scale). Curve: geometric/occupancy approximation $(1 - \pi(\mathcal{C}))^n$ based on (13), together with the conservative exponential bounds derived in the theoretical analysis

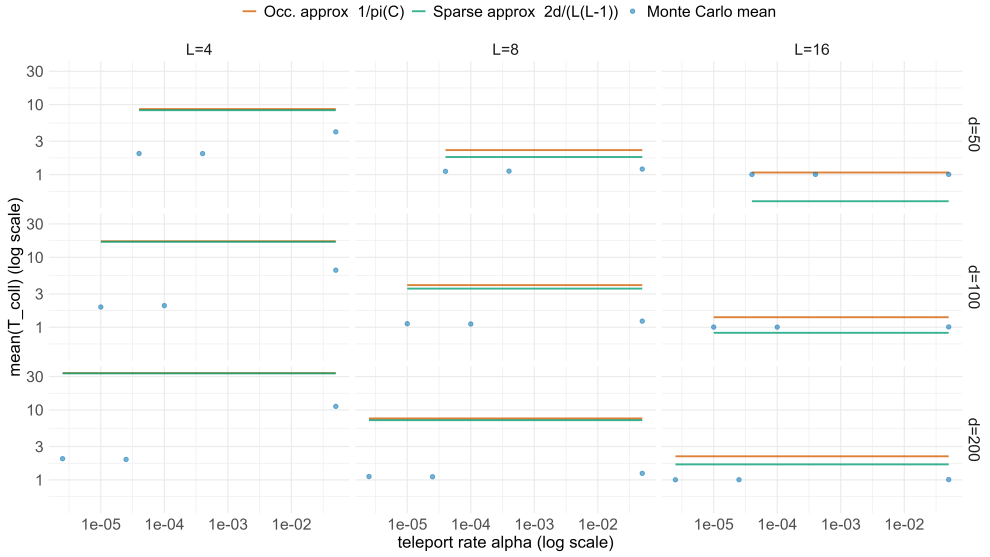


Fig. 7 Mean first collision time across teleport regimes. Markers: Monte Carlo estimates of $\mathbb{E}[T_{\text{coll}}]$ (log scale) for $d \in \{50, 100, 200\}$ and $L \in \{4, 8, 16\}$. Curves: stationary occupancy predictor $1/\pi(\mathcal{C})$ and sparse heuristic $2d/(L(L-1))$.

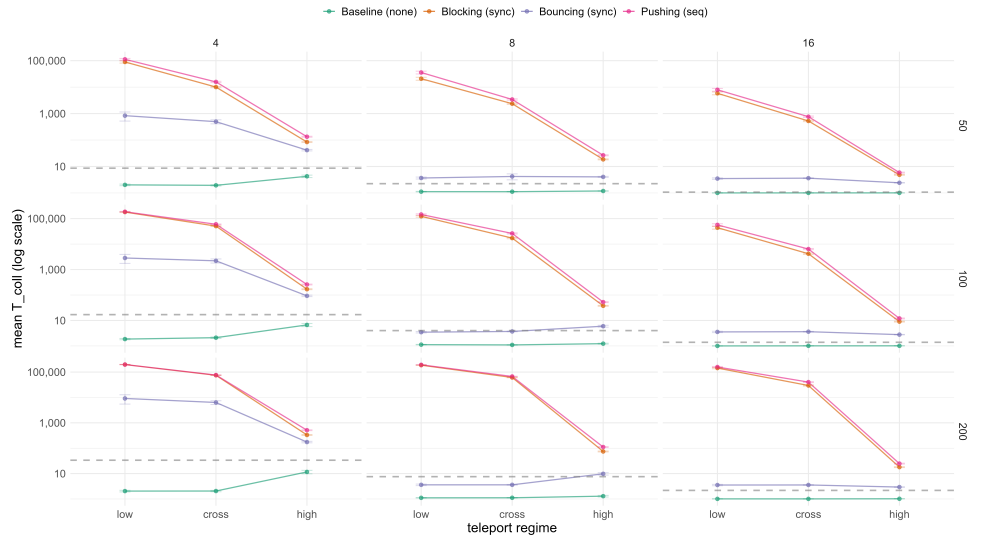


Fig. 8 Mean collision time $\mathbb{E}[T_{\text{coll}}]$ (log scale) across teleport regimes and (d, L) for baseline (none), blocking (sync), bouncing (sync), and pushing (seq). The dashed horizontal line denotes the well-mixed occupancy benchmark $1/\pi(\mathcal{C})$.

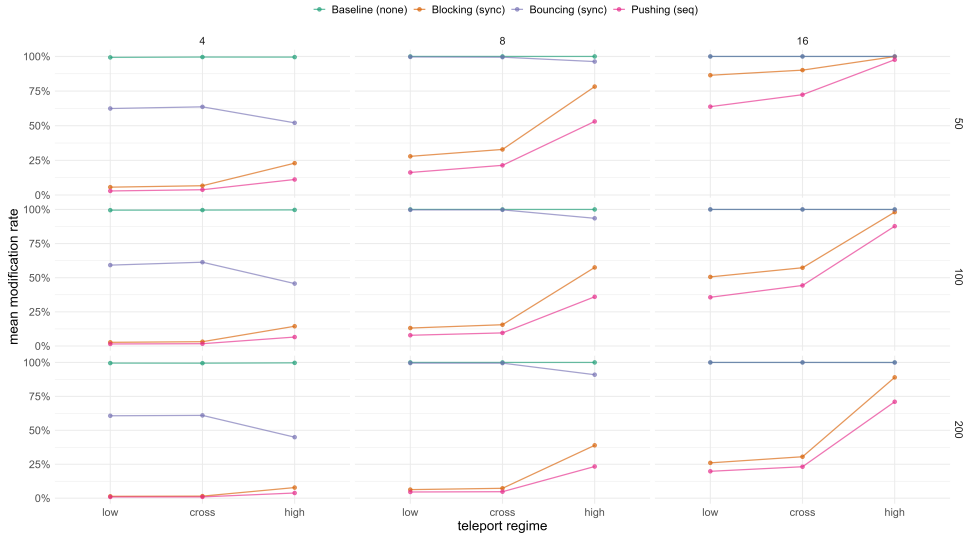


Fig. 9 Mean proposal-modification rate by teleport regime (`low`/`cross`/`high`) across (d, L). Values are normalized to the baseline (none) reference on the same grid, so 100% corresponds to the baseline rate.

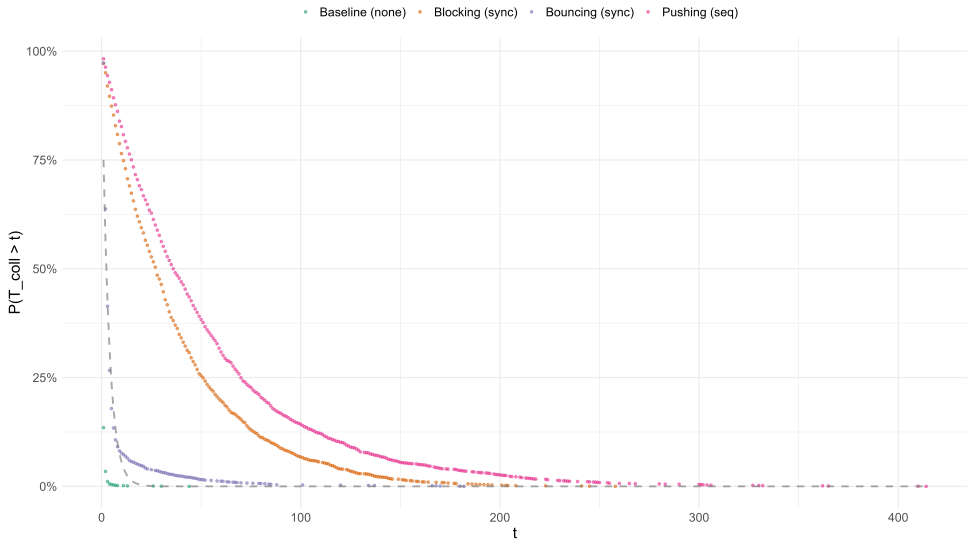


Fig. 10 Empirical survival function of the first collision time $\mathbb{P}(T_{\text{coll}} > t)$ for baseline (none), blocking (sync), bouncing (sync), and pushing (seq).