

PFaffIAN STRUCTURE OF BASIN WALLS FOR COALESCING PARTICLES

PIOTR ŚNIADY

ABSTRACT. Coalescing particles on a line merge when they meet, creating boundaries—walls—between basins of attraction. Previous work established that for Brownian motions and specific discrete systems, the surviving particles form a *Pfaffian point process*, in which all correlations are determined by an antisymmetric kernel. These proofs relied on specific properties of each dynamics, limiting their scope. We give a combinatorial proof that works for any skip-free coalescing process (one where particles cannot pass through each other), with arbitrary deterministic initial conditions and possibly inhomogeneous transition rules. Our approach shows that the Pfaffian structure belongs naturally to the walls, not the particles: we prove an exact Pfaffian formula expressing the probability that intervals contain no wall in terms of probabilities that independent particles from the interval endpoints meet or cross. We derive a cumulant formula for the wall count as a signed sum over proper colorings of a path graph, with universal integer coefficients that depend only on the ordering pattern, not on the specific process. A structural property of this formula—indecomposability, meaning every nonzero term couples all wall positions—implies a central limit theorem. A duality between walls and particles recovers the known Pfaffian point process for surviving particles, extending it to totally asymmetric dynamics and arbitrary initial conditions.

1. INTRODUCTION

1.1. Coalescence and basin walls.

1.1.1. *Coalescence.* When identical particles on a line collide, they merge and continue as one (Figure 1). On a lattice, coalescence describes the voter model [HL75]; in continuous space, the $A + A \rightarrow A$ reaction-diffusion system [DA88]. For coalescing Brownian motions starting from the *maximal entrance law* (every point initially occupied), the system comes down from infinity: at any positive time, only finitely many survivors remain per bounded interval (Arratia [Arr79]).

Tribe and Zaboronski [TZ11] proved that for coalescing Brownian motions under the maximal entrance law, the surviving particle positions form a *Pfaffian point process*; Garrod, Poplavskyi, Tribe, and Zaboronski [GPTZ18] extended this to continuous-time random walks on \mathbb{Z} with spatially inhomogeneous rates and all deterministic initial conditions. In a Pfaffian point process, all correlation functions are Pfaffians of 2×2 matrix kernels. A Pfaffian is a signed sum over perfect

Date: March 21, 2026.

2020 Mathematics Subject Classification. Primary 60K35; Secondary 60F05, 60J65, 60G55, 15A15.

Key words and phrases. coalescing random walks, Pfaffian point process, basin walls, checkerboard duality, cumulant coloring formula, central limit theorem, voter model.




Figure 1. Coalescing random walks starting from every site. Paths merge on contact; line weight increases with each merger. Walls (triangles) mark the boundaries between basins of attraction at $t = 0$; survivors (large dots) are the particles that remain at $t > 0$, one per basin.

matchings of the indices of an antisymmetric matrix—the analogue of a determinant when the underlying combinatorics involves pairings rather than permutations.

1.1.2. *The wall perspective.* This paper takes seriously the “dual particles” of Garrod, Poplavskiy, Tribe, and Zaboronski [GPTZ18, Remark after Lemma 5], interprets their kernel as a crossing-or-meeting probability of independent particles, and develops the resulting wall perspective: the Pfaffian structure belongs naturally not to the surviving particles but to the *walls*—the boundaries between the particles’ basins of attraction (called partition points by Arratia [Arr79], see Figure 1). As particles coalesce, their basins merge and walls disappear; each surviving wall contributes a 2×2 block to the Pfaffian matrix. The surviving particles inherit the Pfaffian structure only because the *checkerboard duality* identifies them with walls of the dual process (Section 1.4).

1.1.3. *Scope.* The wall perspective yields more than the Pfaffian point process. From the empty-interval formula (Theorem 1.1 below) we derive an explicit formula for the cumulants of the wall count, and a structural property of this formula (indecomposability) that implies a central limit theorem. The results work for any *skip-free* process (transitions only to neighboring states, so that particles cannot change order without first meeting [KM59]) with any deterministic initial condition, including discrete-time systems with inhomogeneous transition probabilities and totally asymmetric dynamics. They treat pure coalescence only (not the mixed coalescence-annihilation models of [GPTZ18]); see Section 1.3 for a discussion of what this scope covers and why.

1.2. Results for walls.

1.2.1. *The empty-interval formula.* A wall is a boundary between two adjacent basins of attraction—a half-integer in $\mathbb{Z}' = \mathbb{Z} + \frac{1}{2}$ for processes on \mathbb{Z} , or a real number for processes on \mathbb{R} .

Theorem 1.1 (Pfaffian empty-interval formula). *Consider a coalescing skip-free process with every point initially occupied. Run the system for time $T > 0$, and let $a_1 < b_1 \leq a_2 < \dots \leq a_n < b_n$ be given. Then*

$$\mathbb{P}(\text{no wall in any } (a_i, b_i)) = \text{Pf}(A),$$

where A is the $2n \times 2n$ antisymmetric matrix indexed by the $2n$ endpoints $a_1, b_1, a_2, b_2, \dots, a_n, b_n$ (numbered $1, 2, \dots, 2n$ in this order). For $k < l$, the entry A_{kl} is the probability that two independent copies of the underlying process, started at the k -th and l -th endpoints and run for time T , cross or meet (their ordering reverses or they occupy the same position).

Specific instances include:

- coalescing Brownian motions on \mathbb{R} (Section 7);
- symmetric random walks (Arratia’s original setting; Section 9.2);
- totally asymmetric dynamics, where particles jump only in one direction (Section 9.3);
- continuous-time Poisson jumps with space-time varying rates (Section 9.4).

1.2.2. *Pfaffian point process on the lattice.* On the discrete lattice, the empty-interval formula directly implies that the wall configuration is a Pfaffian point process: occupancy at site $x \in \mathbb{Z}'$ indicates the presence of a wall in the unit interval $(x - \frac{1}{2}, x + \frac{1}{2})$, and inclusion-exclusion converts the empty-interval Pfaffian into a 2×2 matrix kernel whose entries are discrete differences of the crossing-or-meeting probabilities (Proposition 6.1).

1.2.3. *Brownian motion.* For coalescing Brownian motions, the crossing-or-meeting probabilities A_{kl} are complementary error functions of the endpoint separations (Section 7). Differentiating the Pfaffian empty-interval formula and taking the zero-length limit recovers the Tribe–Zaboronski Pfaffian point process (Section 7), the continuous analogue of the discrete kernel above.

1.2.4. *Cumulants and the central limit theorem.* The Pfaffian formula also yields explicit combinatorial expressions for the *cumulants* of the wall count—the connected part of the k -point moment (Theorem 3.1). These expressions cleanly separate combinatorics from probability: the cumulant is a weighted average of universal integer coefficients $\text{wt}(c)$ over random orderings of independent particles started at the interval endpoints; the coefficients depend only on the ordering pattern, not on the process.

A key structural property is *indecomposability*: every nonzero term in the k -th cumulant couples all k wall positions together, preventing the sum from splitting into independent pieces. This explains *why* the wall process has only local correlations, gives the sign of the covariance (repulsion) for free, and yields the cumulant bound $\kappa_k(N_L) = O(L)$ for $k \geq 2$ (Theorem 5.1).

1.3. Why Pfaffians?

1.3.1. *Three links.* Why should the wall process be Pfaffian? The proof of Theorem 1.1 has three links.

First, no wall in an interval (a, b) means that the particles starting at its endpoints a and b have met—a pairwise coalescence event.

Second, the *cancellative labeling* converts pairwise coalescence into annihilation (Section 2.2).

Third, annihilation naturally involves matchings: $2k$ particles must pair up to completely annihilate ($A + A \rightarrow \emptyset$), and the different pairings are exactly the perfect matchings of $\{1, \dots, 2k\}$. The particle trajectories interact (annihilation!), so the total annihilation probability does not decompose into pairwise terms in any obvious way. The companion paper [Śni26b] proves that the total annihilation probability nonetheless equals a Pfaffian of pairwise crossing probabilities, under the bare Karlin–McGregor assumptions (order preservation and the strong Markov property).

1.3.2. *Connection with Markov duality.* The above three-step reduction—from the interacting coalescing system to crossing probabilities of independent particles—has the same structure as a Markov duality. Garrod, Poplavskiy, Tribe, and Zaboronki [GPTZ18, Remark after Lemma 5] observed that their spin-pair duality fits the standard Markov duality framework [EK86, Chapter 4], with the dual process being instantly annihilating particles, and that the same dual process governs both coalescing and annihilating systems—suggesting a common origin for their Pfaffian structure. The cancellative labeling makes this precise: coalescence reduces to annihilation, so both are Pfaffian for the same reason. The annihilation formula [Śni26b] then decomposes the duality into two explicit combinatorial steps, replacing the generator computation by a sign-reversing involution.

1.3.3. *Wide scope.* The three-link reduction (no wall \rightarrow coalescence \rightarrow annihilation \rightarrow Pfaffian) uses only the Karlin–McGregor assumptions: order preservation, the strong Markov property, and meeting times that are stopping times. No generator, no PDE, and no spin-pair identity is needed. This replaces the PDE uniqueness argument of [TZ11; GPTZ18] and extends the Pfaffian structure beyond the time-homogeneous setting to discrete-time systems with arbitrary inhomogeneous transition probabilities, including totally asymmetric dynamics where every particle jumps in the same direction (Section 9.3). This is territory the spin-pair duality cannot reach, since it requires a Markov generator with a specific algebraic structure.

1.3.4. *The coloring formula.* The Pfaffian empty-interval formula is the starting point, not the end. The cumulant coloring formula (Theorem 3.1) reduces each cumulant to universal integer coefficients that are the same for every skip-free process—Brownian motion, random walks, Pólya walks [Urb25], and any other instance. This clean separation of combinatorics from probability has no analogue in the analytic approach: the spin-pair duality produces a 2×2 matrix kernel, but no combinatorial formula for cumulants. The kernel-based route to cumulants (trace formulas, Fredholm expansions) depends on operator-theoretic properties of the specific kernel at hand; the coloring formula bypasses the kernel entirely.

1.3.5. *Indecomposability and the CLT.* Among the coloring formula’s structural properties, the most consequential is *indecomposability* (previewed above): every nonzero coefficient $\text{wt}(c)$ couples all k wall positions together. The CLT follows from the formula’s internal structure—exact formula, combinatorial selection rule, cumulant bound $\kappa_k(N_L) = O(L)$ —rather than from verifying kernel-specific conditions such as finite-rank commutators [WX25], J -Hermiticity [LQW21], or spatial mixing [IL71].




Figure 2. Checkerboard duality on the (u, v) lattice. At each \mathbb{Z}^2 -vertex, a random binary choice (copy from West or South) determines two complementary forests. The **backward opinion forest** (double blue arrows, on \mathbb{Z}^2) traces ancestral lineages in the direction of decreasing $u + v$. The **forward boundary forest** (on \mathbb{Z}^2) has thick red arrows at boundary vertices (where neighboring opinions differ) and thin red arrows at non-boundary vertices. The thick diagonal line marks $u + v = 0$, where all opinions are initially distinct and every \mathbb{Z}^2 -vertex carries a boundary.

1.4. From walls to particles: the checkerboard duality. The results above describe walls. To translate them into statements about surviving particles—the objects most directly measured in applications—requires a duality on the planar lattice.

The duality builds on the graphical construction of Harris [Har78] and Arratia’s “percolation substructure” [Arr79] (Figure 2 and Section 8.1). The independent random choices that define the coalescing system simultaneously generate two complementary non-crossing forests on interleaved sublattices: a *backward opinion forest* carrying coalescent lineages, and a *forward boundary forest* carrying the walls. Different coordinate choices on the lattice yield different particle dynamics (Section 9).

The surviving boundaries of the forward process are precisely the walls of the backward process (Section 8.2). To study the surviving particles at time T —which are the walls of the dual backward process—one applies the wall-level results of Sections 2, 3 and 5 to the dual process. The wall-level results require the maximal entrance law, but a clamping argument extends the particle-level results to all deterministic initial conditions (Section 8.4).

1.5. Prior work.

1.5.1. *Analytic approaches.* All approaches to the Pfaffian structure—the spin-pair duality of [TZ11; GPTZ18] and the present checkerboard approach—require skip-free motion, but their scopes differ.

Tribe and Zaboronski [TZ11] established the Pfaffian point process for coalescing Brownian motions under the maximal entrance law. Garrod, Poplavskyi, Tribe, and Zaboronski [GPTZ18] extended this to continuous-time random walks on \mathbb{Z} with spatially inhomogeneous rates and all deterministic initial conditions via the spin-pair duality, yielding a 2×2 matrix kernel giving all correlation functions. Tribe and Zaboronski [TZ26] further extended it to every entrance law, again via the spin-pair duality. FitzGerald, Tribe, and Zaboronski [FTZ22] proved sharp asymptotics for the Fredholm Pfaffians arising from these kernels, rigorously confirming the gap exponents predicted by Derrida and Zeitak [DZ96].

The two approaches cover complementary territory (Section 1.3): the spin-pair duality treats mixed coalescence-annihilation and all entrance laws [GPTZ18; TZ26], while the checkerboard approach treats pure coalescence with inhomogeneous transition probabilities where no generator is available. Neither subsumes the other; they overlap for time-homogeneous random walks and Brownian motion. In this overlap, the two dualities produce the same mathematical object: the “dual particles” of the spin-pair duality are the backward particles of the checkerboard construction (Section 8.3).

For time-homogeneous coalescing systems, the Pfaffian point process structure itself is therefore not new—it was established in [TZ11; GPTZ18]—and a CLT was proved by Glinyanaya and Fomichov [GF17] via spatial mixing. The contributions of the present paper—wide scope, the coloring formula, and indecomposability—are described in Section 1.3.

1.5.2. *CLT for point processes.* Cumulants of determinantal point processes are well understood. Soshnikov [Sos02] proved $\kappa_k = (-1)^{k-1} \text{Tr}(K^k)$ for Hermitian kernels K , from which the CLT follows whenever the variance diverges. For *Pfaffian* point processes, the corresponding theory is more recent. Wang and Xu [WX25] proved the trace formula and a CLT under a finite-rank commutator property for the operators derived from the 2×2 matrix kernel, verified for the Sine_1 and Sine_4 processes. Lin, Qiu, and Wang [LQW21] proved a CLT for Pfaffian point processes whose 2×2 matrix kernel satisfies a symmetry condition (J -Hermiticity), without explicit cumulant formulas.

1.5.3. *CLT for coalescing Brownian motions.* Glinyanaya and Fomichov [GF17] proved a CLT (with Berry–Esseen bound) for the number of surviving coalescing Brownian motions in a bounded interval, using spatial mixing (the Ibragimov–Linnik CLT for weakly dependent sequences [IL71]). The coloring formula provides a different route, described in Section 1.3.

1.5.4. *Companion papers.* This paper is part of a series. The coalescence paper [ŠU26] proves exact determinantal formulas for coalescence probabilities on arbitrary planar weighted directed acyclic graphs. The annihilation paper [Šni26b] extends this to annihilation and proves the Pfaffian pairwise coalescence formula. A third companion [Šni26a] studies the *wall-particle system*—the joint system of survivors and their basin boundaries under the maximal entrance law. The present paper connects the combinatorial framework to Pfaffian point processes.

1.6. Organization. The paper has two parts. The first (Sections 2 to 7) develops the wall-level theory without any lattice structure: the Pfaffian empty-interval formula (Section 2), the cumulant coloring formula and its structural properties (Section 3), the coloring coefficients and proof (Section 4), the central limit theorem for the wall count (Section 5), correlation functions from empty-interval probabilities (Section 6), and the Brownian motion example recovering the Tribe–Zaboronski kernel [TZ11; GPTZ18] (Section 7). The second part (Sections 8 and 9) translates these results to surviving particles via the checkerboard duality (Section 8) and works out explicit formulas for several discrete particle systems (Section 9). Appendix A proves the underlying CLT for sums of dependent indicators via idempotence, and Appendix B tabulates the nonzero colorings for $k = 3$.

2. PFAFFIAN FORMULA FOR WALLS

This section proves the Pfaffian empty-interval formula (Theorem 1.1) for the wall process of any skip-free coalescing system. The proof uses only the definition of walls and the cancellative labeling; no duality or lattice structure is needed.

2.1. Walls and coalescence. Consider a coalescing *skip-free* process: particles move on a linearly ordered space, cannot change order without first meeting [KM59], and coalesce when they meet. On \mathbb{Z} this means nearest-neighbor transitions; on \mathbb{R} it means continuous paths (such as Brownian motion). The transition probabilities may be inhomogeneous in both space and time. Assume every point is initially occupied (the maximal entrance law). Fix a time $T > 0$.

The surviving particles partition the line into *basins of attraction*: the basin of a survivor y is the contiguous set of initial positions whose particle has coalesced with y by time T . A *wall* is a boundary between two adjacent basins (Section 1.1).

Proposition 2.1 (No wall equals coalescence). *For initial positions $a < b$, there is no wall in the interval (a, b) if and only if the particles starting at a and b have coalesced by time T .*

Proof. If the particles from a and b have coalesced, they belong to the same survivor. Since the process is skip-free, every particle starting between a and b must also have met them (it cannot escape without first meeting one of its neighbors). All points in $[a, b]$ share the same basin, so no wall exists in (a, b) .

Conversely, if no wall exists in (a, b) , then all points in $[a, b]$ belong to the same basin, so the particles from a and b have coalesced. \square

For n pairs $a_1, b_1, \dots, a_n, b_n$ with $a_1 < b_1 \leq a_2 < \dots \leq a_n < b_n$, the event {no wall in any (a_i, b_i) } equals *pairwise coalescence*: for each i , the particles from a_i and b_i have coalesced.

2.2. Cancellative labeling (coalescence to annihilation). Label each of the $2n$ particles (from positions $a_1, b_1, \dots, a_n, b_n$) with 1, and all other particles with 0. When two particles coalesce, let the merged particle carry the sum of their labels mod 2. A cluster of s coalesced particles therefore carries label $s \bmod 2$. Tracking only the clusters with label 1 gives an *annihilating* particle system: when two label-1 clusters meet, both labels become 0 ($1 \oplus 1 = 0$). *Total annihilation* means that no label-1 cluster survives; equivalently, every cluster contains an even number of labeled particles. This mod-2 rule is the *cancellative labeling* of Griffeath [Gri79] (see also [AS12; AB05]).

Proposition 2.2 (Pairwise coalescence equals total annihilation). *Pairwise coalescence of the $2n$ particles equals total annihilation under the cancellative labeling.*

Proof. Since the process is skip-free, particles that have met follow the same path thereafter, forming clusters of colocated particles. Clusters consist of consecutive original particles (a particle cannot join a non-adjacent cluster without first meeting its neighbor).

Under pairwise coalescence, each pair (a_i, b_i) is in the same cluster. Since pair members are adjacent in the ordering, each cluster is a union of consecutive pairs and therefore contains an even number of labeled particles. All labels are therefore 0: total annihilation.

Conversely, if some pair (a_i, b_i) is in different clusters, the clusters containing $a_1, b_1, \dots, a_{i-1}, b_{i-1}, a_i$ hold $2i - 1$ labeled particles in total (no later labeled particle can be in these clusters, since b_i is separated by a cluster boundary). Since labels add modulo 2, the total label of these clusters is 1, so at least one cluster carries label 1: not total annihilation. \square

2.3. Pfaffian empty-interval formula.

Theorem 2.3 (Pfaffian empty-interval formula). *Consider a coalescing skip-free process with every point initially occupied. Run the system for time $T > 0$, and let $a_1 < b_1 \leq a_2 < \dots \leq a_n < b_n$ be given. Then:*

$$\mathbb{P}(\text{no wall in any } (a_i, b_i)) = \text{Pf}(A),$$

where A is the $2n \times 2n$ antisymmetric matrix whose entry A_{kl} ($k < l$) is the probability that independent particles from the k -th and l -th endpoints cross or meet.

Proof. By Proposition 2.1, the event equals pairwise coalescence of the $2n$ particles. By Proposition 2.2, pairwise coalescence equals total annihilation. The Pfaffian annihilation formula [Śni26b] gives the probability as $\text{Pf}(A)$. \square

Remark 2.4 (Touching intervals). When $b_i = a_{i+1}$, two particles start at the same position. The Pfaffian annihilation formula [Śni26b] allows coincident starting positions, so the proof applies directly: the entry $A_{kl} = 1$ for each coincident pair (paths from the same position trivially cross or meet).

3. THE CUMULANT COLORING FORMULA

The Pfaffian empty-interval formula (Theorem 1.1) computes joint void probabilities for the wall process. This section extracts the *cumulants*—the connected part of the k -point moment—and expresses them in a form that cleanly separates the combinatorics (universal integer coefficients, the same for every skip-free process) from the probability (the random ordering of independent particles at time T).

Mixed cumulants select the connected part of joint moments via Möbius inversion on the partition lattice [Spe83]; see [PT11] for the diagrammatic perspective and [DJS22] for a survey of the cumulant method in normal approximation.

The recipe is as follows. Run $2k$ independent particles, two per interval: one from each endpoint. Call these the *flanking particles* of the wall. There is no coalescence, no interaction between different walls. Assign *color* w to both flanking particles of wall w . At time T , read their colors from left to right: this random sequence is the *coloring* c . The cumulant is a weighted average of a universal integer coefficient $\text{wt}(c)$ over all coloring outcomes (Theorem 3.1), but with one constraint: within each

wall’s pair, the two flanking particles must not have crossed during $[0, T]$. If they did, the integrand vanishes.

The coloring coefficient $\text{wt}(c)$ (defined in Section 4) vanishes for most colorings: of the $(2k)!/(2^k)$ possible colorings, only the *indecomposable* ones (Theorem 3.2) can contribute. Indecomposability controls the growth of cumulants with the interval length, which is the key input for the central limit theorem (Theorem 5.1).

Absorbing the non-crossing condition into a conditional measure gives a clean factorization: the cumulant equals a product of non-crossing probabilities (one per wall) times a conditional expectation of $\text{wt}(c)$ (Section 3.4).

The definition of $\text{wt}(c)$ and the proof of the coloring formula are given in Section 4.

3.1. The coloring formula.

3.1.1. *Setup and statement.* Fix k disjoint intervals $(a_1, b_1), \dots, (a_k, b_k)$ with $a_1 < b_1 \leq a_2 < \dots \leq a_k < b_k$. On \mathbb{Z} , the endpoints a_w, b_w are integers (so that flanking particles exist at both endpoints). Let η_w denote the indicator that at least one wall lies in (a_w, b_w) .

Start $2k$ independent copies of the underlying process, two per interval: one from a_w and one from b_w , each carrying color w . Run them to time T and sort the $2k$ particles by their final location; call the rank of each particle its *position* ($1 =$ leftmost, $2k =$ rightmost; in the discrete case, ties are broken uniformly at random). Reading the colors in position order gives the *coloring* $c = (c_1, \dots, c_{2k})$ —a sequence of length $2k$ over the alphabet $\{1, \dots, k\}$, each letter appearing exactly twice. In examples we write $A = 1, B = 2, C = 3$, etc. for the wall colors.

Theorem 3.1 (Coloring formula). *For $k \geq 2$,*

$$(3.1) \quad \kappa(\eta_1, \dots, \eta_k) = \mathbb{E} \left[\text{wt}(c) \cdot \prod_{w=1}^k \mathbf{1}_{\text{wall } w \text{ non-crossing}} \right],$$

where $\text{wt}(c)$ is a universal integer depending only on the coloring (Definition 4.3), “non-crossing” means that wall w ’s two flanking particles did not cross during $[0, T]$, and the expectation is over the random trajectories and, in the discrete case, the random tie-breaking.

The proof occupies Section 4.

3.1.2. *The assumption $k \geq 2$.* The derivation converts wall indicators $\eta_w = \mathbf{1}_{\text{wall}}$ to no-wall indicators $\xi_w = 1 - \eta_w = \mathbf{1}_{\text{no wall}}$ via $\kappa_k(\eta_1, \dots, \eta_k) = (-1)^k \kappa_k(\xi_1, \dots, \xi_k)$. This uses invariance of mixed cumulants under additive constants, which holds for $k \geq 2$ but fails for $k = 1$: $\kappa_1(\eta) = 1 - \kappa_1(\xi) \neq -\kappa_1(\xi)$. The first cumulant $\kappa_1(\eta) = P(\text{wall in } (a, b))$ is computed directly from the Pfaffian for a single interval.

3.2. **Indecomposable colorings.** A coloring $c = (c_1, \dots, c_{2k})$ is *indecomposable* if there is no $0 < g < k$ such that $\{c_1, \dots, c_{2g}\} = \{1, \dots, g\}$ (each appearing twice); otherwise it is *decomposable*.

Theorem 3.2 (Indecomposability). *If c is decomposable, then $\text{wt}(c) = 0$.*

The proof is given in Section 4.7.

Remark 3.3 (Gap straddling). If $\text{wt}(c) \neq 0$, then a crossing cycle exists. Since the crossing cycle is a Hamiltonian cycle on all $2k$ positions, for every gap between consecutive positions r and $r + 1$, the cycle must cross from one side to the other —

either via a wall step (the wall straddles the gap) or via a crossing step (two walls on opposite sides). In either case, flanking particles from at least two different walls traverse the segment between positions r and $r + 1$. This is the key property behind the cumulant bounds in the CLT proof (Section 5): it ensures that each gap variable in the coloring integral (4.8) is controlled by a Gaussian factor from a straddling wall’s determinant.

3.3. Covariance. For $k = 2$, the only nonzero coloring is $c = (B, A, B, A)$, with $\text{wt}(c) = -4$ (computed in Section 4.2). This coloring means the four particles interleave: a wall- B particle finishes leftmost, then a wall- A particle, then wall- B , then wall- A . Combined with the non-crossing condition for the pairs of particles starting at each wall, Theorem 3.1 gives

$$\text{Cov}(\eta_1, \eta_2) = -4 \cdot P \left(\begin{array}{l} Y_{a_2} < Y_{a_1} < Y_{b_2} < Y_{b_1} \\ \text{and both walls non-crossing} \end{array} \right),$$

where Y_s denotes the time- T position of an independent particle started at s , and “non-crossing” means the two flanking particles of each wall never crossed during $[0, T]$. The negative sign reflects ‘repulsion’ between walls.

3.4. Non-crossing conditioning. In the coloring formula (3.1), the k pairs of particles evolve independently, and the non-crossing indicators can be absorbed into a conditional measure: condition each pair on its two particles not crossing during $[0, T]$. This gives

$$\kappa(\eta_1, \dots, \eta_k) = \prod_{w=1}^k P_w(\text{no crossing}) \cdot \mathbb{E}_{\text{cond}}[\text{wt}(c)],$$

where $P_w(\text{no crossing})$ is the probability that wall w ’s flanking particles maintain their order, and the expectation is over k independently conditioned non-crossing pairs.

As the flanking distance $\varepsilon = b_w - a_w \rightarrow 0$, the prefactor $P_w(\text{no crossing}) = O(\varepsilon)$ captures all the vanishing, while the conditional density $p_{\text{cond},w}$ of wall w ’s pair given non-crossing converges to the density of a non-colliding Brownian pair—the Doob h -transform with $h(a, b) = b - a$, which conditions the pair to maintain strict order throughout $[0, T]$. The bridge density of this conditioned process involves $\varphi(y - x)$ and $(y - x)\varphi(y - x)$ (where φ is the standard normal density), which are the row functions of the 2×2 kernel (Section 7). The factorization also explains the CLT proof (Section 5): integrating \det_w over the wall position computes $P_w(\text{no crossing})$ marginalized over the wall, which has Gaussian-type decay in the gap $y_{\alpha_{w,2}} - y_{\alpha_{w,1}}$.

4. COLORING COEFFICIENTS AND THE PROOF

This section defines the coloring coefficient $\text{wt}(c)$ appearing in the coloring formula (Theorem 3.1) and proves the formula. The proof reorganizes the Pfaffian expansion: extracting connected matchings, expanding via the Karlin–McGregor segment swap, and reparametrizing in terms of crossing cycles. The section closes with the proof of indecomposability (Theorem 3.2).

4.1. The coloring coefficient. Represent the coloring as a bipartite diagram (Figure 3): positions $1, \dots, 2k$ along the top, walls $1, \dots, k$ along the bottom, with an edge from each position i to its wall c_i .




Figure 3. The coloring $c = (B, A, B, A)$. Positions 1, 2, 3, 4 (top) connect to walls $A < B$ (bottom). Two edges cross visually when the corresponding positions form an inversion: $i < j$ but $c_i > c_j$. Black arrows trace the unique crossing cycle: starting at the crossing of edges $1 \rightarrow B$ and $2 \rightarrow A$, it passes through wall B , then through the crossing of $3 \rightarrow B$ and $4 \rightarrow A$, then through wall A , and back.

4.1.1. *Inversions.* An *inversion* is a pair of positions $i < j$ with $c_i > c_j$: the left position belongs to a higher-numbered wall than the right one. In the diagram, the edges from i and j to their respective walls cross visually exactly when (i, j) is an inversion.

Example 4.1 ($k = 2$). In $c = (B, A, B, A)$ (Figure 3), wall A occupies positions $\{2, 4\}$ and wall B occupies $\{1, 3\}$. Three of the four inter-wall pairs are inversions: $(1, 2)$, $(1, 4)$, and $(3, 4)$. Only the pair $(2, 3)$ is not an inversion.

4.1.2. *Crossing cycles.* A *crossing cycle* \mathbf{p} is a Hamiltonian cycle on the $2k$ positions that alternates between:

- *wall steps*, connecting the two positions with the same color ($c_i = c_j$) —in the diagram, traveling down to a wall node and back up;
- *crossing steps*, connecting two positions that form an inversion—moving from one edge to a crossing edge.

Each wall is visited exactly once: the cycle enters a wall through one position and exits through the other. Equivalently, each edge of the bipartite diagram is traversed exactly once.

Each undirected Hamiltonian cycle has $4k$ representations as a sequence (p_1, \dots, p_{2k}) (two choices of direction times $2k$ choices of starting position). We choose a canonical representative by fixing $p_1 = 1$ and requiring the first step to be a wall step ($c_{p_1} = c_{p_2}$).

A *descent* at index $i \in \{1, \dots, 2k\}$ means $p_i > p_{i+1}$ (cyclically), and $\text{des}(\mathbf{p})$ denotes the total number of descents. Since des is invariant under cyclic rotation and $(-1)^{\text{des}}$ is invariant under reversal (reversal sends $\text{des} \rightarrow 2k - \text{des}$), the coloring coefficient (4.1) depends only on the undirected Hamiltonian cycle, not on the choice of representative.

Example 4.2 ($k = 2$, continued). In $c = (B, A, B, A)$ (Figure 3), there is one crossing cycle:

$$1 \xrightarrow{\text{wall } B} 3 \xrightarrow{\text{cross}} 4 \xrightarrow{\text{wall } A} 2 \xrightarrow{\text{cross}} 1,$$

i.e., $\mathbf{p} = (1, 3, 4, 2)$ with $\text{des}(\mathbf{p}) = 2$ (descents at $4 > 2$ and $2 > 1$).

4.1.3. The coloring coefficient.

Definition 4.3 (Coloring coefficient). For a coloring c of length $2k$ with $k \geq 2$, the *coloring coefficient* is

$$(4.1) \quad \text{wt}(c) = -2^k \sum_{\mathbf{p}} (-1)^{\text{des}(\mathbf{p})},$$

where the sum runs over all crossing cycles \mathbf{p} of c .

4.2. Small examples.

4.2.1. *Coloring weights for $k = 2$.* For $c = (B, A, B, A)$ (Example 4.2), the unique crossing cycle has $\text{des} = 2$, contributing $(-1)^2 = 1$, so $\text{wt}(c) = -2^2 \cdot 1 = -4$.

The five other colorings of length 4 all have $\text{wt}(c) = 0$. Four of them— (A, A, B, B) , (A, B, A, B) , (A, B, B, A) , and (B, A, A, B) —have a position that forms no inversion with any other position, so no crossing step can reach it and no Hamiltonian cycle exists. The remaining coloring (B, B, A, A) has two crossing cycles, $(1, 2, 3, 4)$ and $(1, 2, 4, 3)$, with $\text{des} = 1$ and $\text{des} = 2$ respectively; their contributions $(-1)^1 + (-1)^2 = 0$ cancel.

Thus $c = (B, A, B, A)$ is the only nonzero coloring for $k = 2$ (Figure 3).

4.2.2. *Other small examples.* For $k = 3$, there are 90 colorings (sequences of length 6 over $\{A, B, C\}$ with each letter twice), of which 14 have $\text{wt}(c) \neq 0$. All nonzero values satisfy $|\text{wt}(c)| = 8$; the complete list appears in Appendix B. Only one coloring has $\text{wt}(c) < 0$; the rest are positive. For $k = 4$, there are 386 nonzero colorings (out of 2520 total), with $|\text{wt}(c)| \in \{16, 32, 48\}$ (all multiples of 2^k).

Remark 4.4 (Sign of cumulants). For $k = 2$, the unique nonzero coloring has $\text{wt}(c) = -4 < 0$, so the covariance is manifestly non-positive for every skip-free process (and strictly negative whenever interleaving has positive probability). For $k \geq 3$, both positive and negative coefficients appear (Appendix B). Whether the total cumulant κ_k has a definite sign for all skip-free processes remains an open question.

4.3. **From Pfaffian to connected matchings.** We now derive Theorem 3.1; the argument spans Sections 4.3 to 4.5.

4.3.1. *Cumulants as connected Pfaffian terms.* Assume first that the transition law has a continuous density $p(s, y)$ (no atoms), so that final positions are almost surely distinct; the discrete case is treated in Section 4.3.4.

Since $\eta_w = 1 - \mathbf{1}_{\text{no wall in } (a_w, b_w)}$ and cumulants of order $k \geq 2$ are invariant under additive constants, the mixed cumulants of the wall indicators equal $(-1)^k$ times those of the void indicators. The void probabilities are Pfaffians (Theorem 1.1), and the cumulant is the connected part of the joint moment [Spe83]. For Pfaffians, Möbius inversion on the partition lattice selects precisely the *connected* matchings from the Pfaffian expansion (see also [WX25]). Here a perfect matching M on the $2k$ initial particles is *connected* if the graph formed by the matching edges of M together with the wall edges $\{2w-1, 2w\}$ ($w = 1, \dots, k$) is connected; equivalently, it forms a single Hamiltonian cycle on $2k$ vertices.

4.3.2. *Expanding matching edges via segment swap.* The Pfaffian expansion is an algebraic sum over matchings and permutations with signs; the rest of the derivation reorganizes it into the coloring formula, which separates the combinatorial signs from the probabilistic content.

By the Karlin–McGregor segment swap [KM59], the antisymmetric matrix entry A_{ij} in the Pfaffian (the probability that particles i and j cross or, in the discrete case, meet) satisfies

$$(4.2) \quad A_{ij} = 2 \int_{z_i > z_j} p(s_i, z_i) p(s_j, z_j) dz_i dz_j,$$

where s_i, s_j are the starting positions. The factor 2 arises because the segment swap bijects paths that crossed and ended unswapped with paths that crossed and ended swapped; both contribute equally. Expanding all k matching edges simultaneously, each connected matching M produces a factor 2^k (one factor 2 per edge) times a $2k$ -dimensional integral over z_1, \dots, z_{2k} with the constraint $z_i > z_j$ for every matching edge $\{i, j\} \in M$ with $i < j$.

4.3.3. *Lifts and the simplex integral.* Rewriting on the ordered simplex $y_1 < y_2 < \dots < y_{2k}$ via a permutation $\sigma \in S_{2k}$ (particle i lands at position $y_{\sigma(i)}$), the constraint from the integral in (4.2) becomes $\sigma(i) > \sigma(j)$. We call σ a *lift* of the matching M ; it is *compatible* with M if $\sigma(i) > \sigma(j)$ for every matching edge $\{i, j\} \in M$ with $i < j$. The contribution of each compatible lift is the simplex integral

$$(4.3) \quad J(\sigma) = \int_{y_1 < \dots < y_{2k}} \prod_{i=1}^{2k} p(s_i, y_{\sigma(i)}) dy,$$

where s_i is the starting position of particle i ($s_{2w-1} = a_w, s_{2w} = b_w$) and $p(s, y)$ the transition density. Combining the steps above, the cumulant becomes

$$(4.4) \quad \kappa(\eta_1, \dots, \eta_k) = (-1)^k \cdot 2^k \sum_{M \text{ connected}} \text{sgn}(M) \sum_{\substack{\sigma \text{ lift} \\ \text{compat. with } M}} J(\sigma).$$

4.3.4. *Discrete tie-breaking.* In the discrete setting, ties are broken by attaching a private copy of \mathbb{R} to each final vertex n : a particle at discrete position n receives an independent uniform continuous coordinate on a small interval confined to n 's copy (from which it cannot reach another vertex's interval). This does not alter the wall system: particles sharing a copy of \mathbb{R} have already coalesced in the discrete dynamics, so they belong to the same basin regardless of the continuous perturbation. The lexicographic order breaks all ties almost surely, so the strict simplex formulation applies unchanged. Since the transition probabilities depend only on the discrete coordinate, integrating out the continuous tie-breakers recovers the $1/\prod_j m_j!$ multiplicity weights on the weak simplex $y_1 \leq y_2 \leq \dots \leq y_{2k}$. When both flanking particles of a wall land at the same vertex, the \det_w factor vanishes (the determinant has identical columns), so only configurations with distinct within-wall positions contribute.

4.4. From simplex to coloring formula.

4.4.1. *Change of variables.* We reparametrize the double sum (4.4) over connected matchings M and compatible lifts σ . A lift σ sends each initial particle i to position $\sigma(i)$ (its final position in the ordered simplex); it therefore sends the matching M on initial particles to a perfect matching $\sigma(M)$ on positions: position $\sigma(i)$ is paired with position $\sigma(j)$ for each matching edge $\{i, j\} \in M$. The lifted matching $\sigma(M)$ pairs positions from different walls (different colors). Linking each pair of same-colored positions (the two positions of each wall) gives k additional edges; together with the k edges of $\sigma(M)$, they form an undirected Hamiltonian cycle \mathbf{p} on the $2k$ positions, alternating between same-color edges and lifted matching edges—precisely the structure of a crossing cycle (Section 4.1.2).

For each wall w , define the *wall orientation*

$$\epsilon_w := \begin{cases} +1 & \text{if } \sigma(2w-1) < \sigma(2w), \\ -1 & \text{if } \sigma(2w-1) > \sigma(2w). \end{cases}$$

The pair $(\mathbf{p}, (\epsilon_w)_{w=1}^k)$ determines σ uniquely (and hence M): the cycle \mathbf{p} specifies which two positions each wall occupies, and ϵ_w specifies which initial particle maps to which position. Conversely, every pair $(\mathbf{p}, (\epsilon_w))$ arises from a unique (M, σ) . The change of variables

$$\sum_{M \text{ connected}} \sum_{\substack{\sigma \text{ lift} \\ \text{compat. with } M}} \longleftrightarrow \sum_{\mathbf{p}} \sum_{(\epsilon_w) \in \{+1, -1\}^k}$$

reparametrizes the double sum in (4.4), where the sum over \mathbf{p} is restricted by the compatibility condition (see below).

4.4.2. *Compatibility and crossing cycles.* The compatibility condition on σ — that $\sigma(i) > \sigma(j)$ for each matching edge $\{i, j\}$ with $i < j$ — translates into a condition on \mathbf{p} alone, independent of (ϵ_w) . Each lifted matching edge connects two positions α and β with different colors $c_\alpha \neq c_\beta$. If $c_\alpha < c_\beta$, compatibility requires $\alpha > \beta$: the position with the smaller color has a larger index—which is exactly the inversion condition defining a crossing step. This depends only on \mathbf{p} , not on (ϵ_w) . Combined with the same-color edges, \mathbf{p} is a crossing cycle of the coloring c .

Conversely, every crossing cycle arises from a compatible lift. The compatibility condition therefore restricts the outer sum to crossing cycles:

$$(4.5) \quad \kappa(\eta_1, \dots, \eta_k) = (-1)^k \cdot 2^k \sum_{\mathbf{p} \text{ crossing}} \sum_{(\epsilon_w)} \text{sgn}(M) J(\sigma).$$

The sign $\text{sgn}(M)$ depends implicitly on \mathbf{p} and (ϵ_w) ; we now compute this dependence explicitly.

4.5. Computing the Pfaffian sign.

4.5.1. *The canonical matching.* Define the *canonical matching* M_0 as the matching induced by the all-positive orientation $\epsilon_w = +1$ for every w . Flipping ϵ_w at a single wall swaps which particle maps to which position, which is a transposition in M that changes $\text{sgn}(M)$ by a factor of -1 . Therefore

$$(4.6) \quad \text{sgn}(M) = \text{sgn}(M_0) \cdot \prod_{w=1}^k \epsilon_w.$$

The factor $\text{sgn}(M_0)$ depends only on \mathbf{p} ; the ϵ_w -dependence is entirely in the product.

4.5.2. *Computing $\text{sgn}(M_0)$.* The crossing cycle \mathbf{p} visits the k walls in a cyclic order (w_1, \dots, w_k) ; we normalize by $w_1 = 1$. Each matching edge of M_0 connects the *exit particle* of wall w_j to the *entry particle* of wall w_{j+1} (indices cyclic). In M_0 ($\epsilon_w = +1$ for all w), particle $2w-1$ maps to the smaller position of wall w and $2w$ to the larger. If the wall step at w_j is an ascent (the cycle enters at the smaller position), the entry particle is $2w_j-1$ and the exit particle is $2w_j$; if the wall step is a descent, these are reversed.

The Pfaffian sign $\text{sgn}(M_0)$ is the sign of the permutation π_{M_0} that lists matching edges oriented (smaller element first) and sorted by first element. We compute it by decomposing π_{M_0} into three factors.

4.5.3. *Edge-listing permutation.* List the matching edges in cycle order, each oriented exit-to-entry:

$$\lambda = (\text{exit}(1), \text{ent}(2), \text{exit}(2), \text{ent}(3), \dots, \text{exit}(k), \text{ent}(1)).$$

Each particle appears exactly once, so $\lambda \in S_{2k}$. To obtain π_{M_0} from λ :

- (1) *Orient* each edge: swap the two elements of pair j when the exit particle exceeds the entry particle. Since particles of wall w occupy $\{2w-1, 2w\}$, this happens exactly when $w_j > w_{j+1}$. Writing d_c for the number of such cyclic descents in the wall sequence, the orientation step has sign $(-1)^{d_c}$.
- (2) *Sort* pairs by their smaller element. This permutes k blocks of size 2; swapping two adjacent 2-blocks is a product of two disjoint transpositions (sign $+1$), so every block permutation has sign $+1$.

Therefore $\text{sgn}(M_0) = (-1)^{d_c} \cdot \text{sgn}(\lambda)$.

4.5.4. *Cyclic shift.* Define γ by listing *walls* in cycle order, each contributing its entry particle followed by its exit particle:

$$\gamma = (\text{ent}(1), \text{exit}(1), \text{ent}(2), \text{exit}(2), \dots, \text{ent}(k), \text{exit}(k)).$$

Then λ is a cyclic left-shift of γ by one position. A cyclic shift of $2k$ elements has sign $(-1)^{2k-1} = -1$, so $\text{sgn}(\lambda) = -\text{sgn}(\gamma)$.

4.5.5. *Wall decomposition.* The permutation γ reorders the k wall-pairs $\{2w-1, 2w\}$ into cycle-visit order (a block permutation, sign $+1$) and, within each wall whose wall step is a descent, swaps the two particles. Writing d_w for the number of wall step descents, $\text{sgn}(\gamma) = (-1)^{d_w}$.

4.5.6. *Combining the sign.*

$$\text{sgn}(M_0) = (-1)^{d_c} \cdot (-1) \cdot (-1)^{d_w} = (-1)^{d_c+d_w+1}.$$

The crossing condition (Section 4.4) provides the final link: a crossing step from wall w_j to wall w_{j+1} is a descent if and only if $w_j < w_{j+1}$ (ascending wall label), because compatibility forces the position with smaller color to have larger index. Therefore the number of crossing step descents equals $k - d_c$, giving $\text{des}(\mathbf{p}) = d_w + (k - d_c)$. Substituting $d_c = k + d_w - \text{des}(\mathbf{p})$:

$$\text{sgn}(M_0) = (-1)^{(k+d_w-\text{des}(\mathbf{p}))+d_w+1} = (-1)^{k+1+\text{des}(\mathbf{p})},$$

since $2d_w$ is even. Equivalently,

$$(4.7) \quad (-1)^k \text{sgn}(M_0) = -(-1)^{\text{des}(\mathbf{p})}.$$

4.6. **Completing the proof.**

4.6.1. *Factoring over walls.* Write $\alpha_{w,1} < \alpha_{w,2}$ for the two positions of wall w . Substituting (4.6) and (4.7) into (4.5),

$$\kappa(\eta_1, \dots, \eta_k) = -2^k \sum_{\mathbf{p} \text{ crossing}} (-1)^{\text{des}(\mathbf{p})} \sum_{(\epsilon_w)} \prod_{w=1}^k \epsilon_w \cdot J(\sigma).$$

The simplex integral $J(\sigma)$ from (4.3) factors over walls: wall w 's two particles contribute $p_w^{(\epsilon_w)}(y_{\alpha_{w,1}}, y_{\alpha_{w,2}})$, where $p_w^{(+1)}(u, v) = p(a_w, u)p(b_w, v)$ and $p_w^{(-1)}(u, v) = p(a_w, v)p(b_w, u)$ (the a_w -particle lands at the smaller or larger position depending on ϵ_w). Since each ϵ_w appears in exactly one factor of $J(\sigma)$ and one factor of the product $\prod \epsilon_w$, the sum over (ϵ_w) factorizes wall by wall:

$$\sum_{\epsilon_w = \pm 1} \epsilon_w p_w^{(\epsilon_w)}(u, v) = p(a_w, u)p(b_w, v) - p(a_w, v)p(b_w, u) =: \det_w(u, v).$$

By the Karlin–McGregor theorem [KM59], $\det_w(u, v) du dv$ is the probability that two independent particles from a_w and b_w end at (u, v) without crossing; in particular, $\det_w(u, v) \geq 0$. Thus

$$(4.8) \quad \kappa(\eta_1, \dots, \eta_k) = -2^k \sum_{\mathbf{p} \text{ crossing}} (-1)^{\text{des}(\mathbf{p})} \int_{y_1 < \dots < y_{2k}} \prod_{w=1}^k \det_w(y_{\alpha_{w,1}}, y_{\alpha_{w,2}}) dy.$$

4.6.2. *Conclusion.* Grouping crossing cycles by the coloring c they induce,

$$\kappa(\eta_1, \dots, \eta_k) = \sum_c \left(-2^k \sum_{\mathbf{p} \text{ crossing}} (-1)^{\text{des}(\mathbf{p})} \right) \int_{y_1 < \dots < y_{2k}} \prod_{w=1}^k \det_w(y_{\alpha_{w,1}}, y_{\alpha_{w,2}}) dy.$$

The parenthesized factor is exactly $\text{wt}(c)$ from (4.1), and the integral is the probability that the coloring equals c and all k walls are non-crossing. This completes the proof of Theorem 3.1. \square

4.7. Proof of indecomposability.

Proof of Theorem 3.2. The first group of walls $\{1, \dots, g\}$ occupies positions $\{1, \dots, 2g\}$, and the second group $\{g+1, \dots, k\}$ occupies $\{2g+1, \dots, 2k\}$. Every crossing cycle must contain a crossing step between the two groups. But for any position $p \leq 2g$ and $q \geq 2g+1$, we have $p < q$ and $c_p \leq g < g+1 \leq c_q$, so (p, q) is not an inversion. Since the cycle cannot cross between the two groups, no crossing cycle exists and $\text{wt}(c) = 0$. \square

5. VARIANCE AND THE CENTRAL LIMIT THEOREM

For discrete skip-free processes at a fixed time horizon, the wall indicators have finite-range dependence, so a central limit theorem follows from classical results for locally dependent variables. The coloring formula gives more than the CLT itself: it provides the cumulant bound $\kappa_k(N_L) = O(L)$ for each $k \geq 2$ (Section 5.1), and immediately shows that the covariance is negative (repulsion) for every skip-free process (Section 3.3).

5.1. Central limit theorem, discrete case. Write $N_L = N([0, L])$ for the number of walls in the interval $[0, L]$.

Theorem 5.1 (Cumulant bound: discrete). *Consider a coalescing process on \mathbb{Z} with crossing-or-meeting probabilities satisfying*

$$(5.1) \quad \sum_{d=1}^{\infty} d^N \sup_{a \in \mathbb{Z}} \mathbb{P}(X_a \geq X_{a+d}) < \infty \quad \text{for every } N \geq 0,$$

where X_a, X_{a+d} are independent particles started at a and $a + d$. Then for each fixed $k \geq 2$,

$$\kappa_k(N_L) = O(L) \quad \text{as } L \rightarrow \infty.$$

Morally, condition (5.1) says that particles far apart almost never cross, and the crossing probability decays faster than any polynomial in their separation.

Proof. Fix $k \geq 2$. The wall count

$$N_L = \sum_{x \in \mathbb{Z}' \cap [0, L]} \eta_x,$$

where $\eta_x = \mathbf{1}\{\text{wall at } x\}$, is a finite sum of 0-1 variables. By Proposition A.1, it suffices to show that for each fixed k ,

$$(5.2) \quad \sup_{x_0 \in \mathbb{Z}' \cap [0, L]} \sum_{\substack{A \ni x_0, |A|=k \\ A \subseteq \mathbb{Z}' \cap [0, L]}} |\kappa_A| = O(1).$$

Step 1: Reduction to indecomposable colorings. Fix any $x_0 \in \mathbb{Z}'$ and a set A of size k containing x_0 . The coloring formula (Theorem 3.1) gives

$$|\kappa_A| \leq \sum_c |\text{wt}(c)| \cdot \mathbb{P}_A(\text{coloring} = c, \text{ all non-crossing}),$$

where \mathbb{P}_A is over $2k$ independent particles (two flanking each wall). By Theorem 3.2, $\text{wt}(c) = 0$ unless c is indecomposable. Restricting to indecomposable c and exchanging the sums over A and c :

$$(5.3) \quad \text{LHS of (5.2)} \leq \sum_{c \text{ indecomp.}} |\text{wt}(c)| \sum_{\substack{A \ni x_0 \\ |A|=k}} \mathbb{P}_A(\text{coloring} = c, \text{ all non-crossing}).$$

Step 2: Crossing at the widest gap. We show that the inner sum is finite. Write $A = \{x_1 < \dots < x_k\}$ with $x_0 = x_j$ for some j , and let $d_{\max} = \max_g (x_{g+1} - x_g)$ be the widest gap between consecutive walls.

Consider the split at this widest gap: walls x_1, \dots, x_{g^*} on the left, x_{g^*+1}, \dots, x_k on the right. Since c is indecomposable, the $2k$ final positions (sorted left to right) do not decompose at g^* : some particle from the left group ends up to the right of some particle from the right group. These two particles have crossed or met. By the union bound over the at most $4k^2$ pairs of flanking particles,

$$\mathbb{P}_A(\dots) \leq 4k^2 \cdot \max_{\substack{a \leq x_{g^*} + \frac{1}{2} \\ b \geq x_{g^*+1} - \frac{1}{2}}} \mathbb{P}(X_a \geq X_b).$$

Step 3: Summation over A . Since $b - a \geq d_{\max} - 1$ for any such pair, and $d_{\max} \geq \max_i |x_i - x_0|/(k-1)$, summing over all sets A with $d_{\max} = d$ (at most $(2(k-1)d+1)^{k-1}$ choices) gives

$$\sum_{\substack{A \ni x_0 \\ |A|=k}} \mathbb{P}_A(\cdots) \leq 4k^2 \sum_{d=1}^{\infty} (2(k-1)d+1)^{k-1} \cdot \sup_a \mathbb{P}(X_a \geq X_{a+d-1}),$$

which converges by (5.1) with $N = k - 1$. Since the outer sum in (5.3) has finitely many indecomposable colorings, (5.2) holds. \square

Remark 5.2 (Local dependence and scaling limits). For a discrete-time skip-free process with T steps, the wall process has finite-range dependence (range $2T$), and the CLT is classical. The per-site bound depends on the time horizon: when T grows with the window L , the crude bound is useless. For example, if $L \sim \sqrt{T}$, the dependence range $2T \sim L^2$ grows faster than the window. This is where indecomposability (Theorem 3.2) becomes essential: it ensures that the per-site sum remains bounded as $T \rightarrow \infty$, by localizing all simplex variables through gap straddling (Remark 3.3).

5.2. Central limit theorem, continuous case. For continuous processes, the tail condition (5.1) cannot be applied directly to a refined grid: as the grid mesh $\epsilon \rightarrow 0$, each wall's flanking particles start at distance ϵ , so $\mathbb{P}(X_a \geq X_{a+\epsilon}) \rightarrow 1/2$ and the sum diverges. The non-colliding conditioning (Section 3.4) resolves this by factoring out the vanishing prefactor.

Theorem 5.3 (Cumulant bound: continuous). *Consider a coalescing process on \mathbb{R} with continuous paths under the maximal entrance law. For flanking particles at distance ϵ , condition on non-crossing during $[0, T]$ (Section 3.4). Assume:*

(i) Wall intensity. *The limit*

$$(5.4) \quad \varrho(x) = \lim_{\epsilon \rightarrow 0} \frac{\mathbb{P}(\text{no crossing at distance } \epsilon \text{ near } x)}{\epsilon}$$

exists (the wall intensity may depend on the spatial position x). In particular, $\mathbb{E}[N_L] = \int_0^L \varrho(x) dx$, so the $k = 1$ bound is immediate.

(ii) Limiting conditioned process. *As $\epsilon \rightarrow 0$, the conditioned non-colliding pair converges to a limiting process (the Doob h -transform with $h(a, b) = b - a$).*

(iii) Tail condition. *The crossing-or-meeting probabilities of particles from different limiting conditioned pairs satisfy*

$$(5.5) \quad \int_0^{\infty} \delta^N \sup_a \mathbb{P}_{\text{cond}}(Y_a^{(1)} \geq Y_b^{(2)}) d\delta < \infty \quad \text{for every } N \geq 0,$$

where $Y_a^{(1)}, Y_b^{(2)}$ are final positions of particles from two different limiting conditioned pairs starting at distance $\geq \delta$.

Then for each fixed $k \geq 2$,

$$\kappa_k(N_L) = O(L) \quad \text{as } L \rightarrow \infty.$$

Proof. Detect walls on a grid of mesh ϵ : partition $[0, L]$ into intervals of length ϵ and let $\eta_j^{(\epsilon)} = \mathbf{1}\{\text{wall in the } j\text{-th interval}\}$, $N_L^{(\epsilon)} = \sum_j \eta_j^{(\epsilon)}$. Each $\eta_j^{(\epsilon)}$ is a 0-1 indicator, and the coloring formula (Theorem 3.1) applies with flanking particles at the interval endpoints.

By the non-colliding conditioning (Section 3.4), each cumulant factors as

$$\kappa_A = \prod_w P_w(\text{no crossing}) \cdot \mathbb{E}_{\text{cond}}[\text{wt}(c)],$$

so

$$|\kappa_A| \leq \prod_w P_w(\text{no crossing}) \cdot \sum_c |\text{wt}(c)| \cdot \mathbb{P}_{\text{cond}}(\text{coloring} = c).$$

Apply the discrete-case argument (Theorem 5.1) to the conditional probability: indecomposability (Theorem 3.2) and the union bound control the per-site sum. By (5.4), $P_w(\text{no crossing}) = \varrho(x_w)\epsilon + o(\epsilon)$ for each wall at position x_w , contributing a factor $\prod_w \varrho(x_w)\epsilon$ to $|\kappa_A|$. The sums over the $k - 1$ free wall positions on the ϵ -grid are Riemann sums: each sum contributes a factor ϵ times an integral, and the integrals converge by the conditioned tail condition (iii). By (ii), the conditioned process converges as $\epsilon \rightarrow 0$, so the Riemann sums converge to the corresponding integrals. Altogether, $\kappa_k(N_L^{(\epsilon)}) = O(L)$ uniformly in ϵ .

As $\epsilon \rightarrow 0$, each interval contains at most one wall (the expected wall count $\int_0^L \varrho(x) dx$ is finite by (i), so at most finitely many walls exist almost surely), and $N_L^{(\epsilon)} \rightarrow N_L$. Since $0 \leq N_L^{(\epsilon)} \leq N_L$ and N_L has finite moments (it is bounded by the number of walls in $[0, L]$, which is almost surely finite by (i)), dominated convergence gives convergence of all moments, hence $\kappa_k(N_L) = O(L)$. \square

Remark 5.4. For coalescing Brownian motion, the conditioned pairs are non-colliding Brownian motions (the Doob h -transform with $h(a, b) = b - a$). The conditioned final positions have densities involving $\varphi(y - x)$ (where φ is the standard normal density) and $(y - x)\varphi(y - x)$ (Section 3.4), so the crossing-or-meeting probability between different conditioned pairs decays with Gaussian tails, and (5.5) holds.

For coalescing Brownian motion, the $k = 2$ case of Theorem 5.3 gives $\kappa_2(N_L) = O(L)$; the exact variance constant $(\sqrt{2} - 1)^2$ per unit length was computed by Glinyanaya and Fomichov [GF17], giving $\kappa_2(N_L) = \Theta(L)$. The standard CLT criterion (see, e.g., [Sos02]) then gives:

Corollary 5.5 (CLT for the wall count). *For coalescing Brownian motion under the maximal entrance law,*

$$\frac{N_L - \mathbb{E}[N_L]}{\sqrt{\text{Var}(N_L)}} \xrightarrow{d} \mathcal{N}(0, 1) \quad \text{as } L \rightarrow \infty.$$

5.3. Discussion.

5.3.1. *Structural insight from indecomposability.* The coloring formula adds structural insight beyond the classical CLT: indecomposability explains *why* the dependence is local (every gap between wall positions must be straddled by some particle pair, preventing long-range correlations), and the $k = 2$ analysis gives the sign of the covariance for free (Section 3.3).

5.3.2. *Comparison with kernel-based CLTs.* For time-homogeneous processes, a CLT can also be obtained from the explicit Pfaffian kernel [GF17; WX25; LQW21]. The coloring formula provides a different route that bypasses the kernel entirely, and therefore applies to inhomogeneous processes where no kernel is available. Moreover, the wall-level CLT (Theorem 5.1) is entirely self-contained: it uses only the coloring formula and the crossing-or-meeting probabilities of the original process, with no

duality or lattice structure needed. The checkerboard enters only when translating the wall CLT into a CLT for the surviving particles (Section 8.2.3).

6. FROM EMPTY INTERVALS TO POINT PROCESSES

The Pfaffian empty-interval formula (Theorem 2.3) gives the probability that specified intervals contain no wall as a Pfaffian whose entries are crossing-or-meeting probabilities of independent particles. This section shows how to recover correlation functions—the probability of finding walls at specified sites—from the empty-interval formula, yielding a Pfaffian point process for the wall configuration.

6.1. Continuous setting. In the continuous (Brownian motion) setting, differentiation suffices: the n -point correlation function is obtained by differentiating the n -interval empty probability with respect to the interval lengths and taking the zero-length limit. For coalescing Brownian motions under the maximal entrance law, this recovers the Tribe–Zaboronski Pfaffian point process [TZ11; GPTZ18]; Section 7 records the explicit kernel.

6.2. Discrete setting. In the discrete setting, Möbius inversion on the Boolean lattice converts empty-interval probabilities into correlation functions. Garrod, Poplavskiy, Tribe, and Zaboronski [GPTZ18, Theorem 1] showed that the alternating sum of Pfaffians collapses to a single Pfaffian with a 2×2 matrix kernel: the passage from empty to occupied indicators acts as a forward difference on the Pfaffian entries. Feeding our empty-interval formula (Theorem 2.3) into their reconstruction gives the following.

Proposition 6.1 (Discrete Pfaffian point process [GPTZ18, Theorem 1]). *Under the maximal entrance law for any discrete skip-free coalescing process with arbitrary inhomogeneous transition probabilities, the wall configuration at time T is a Pfaffian point process: the n -point correlation function at sites $y_1 < \dots < y_n$ on \mathbb{Z}' is*

$$\rho_n(y_1, \dots, y_n) = \text{Pf}[K(y_i, y_j)]_{1 \leq i, j \leq n},$$

where K is the 2×2 antisymmetric matrix kernel

$$(6.1) \quad K(y, z) = \begin{pmatrix} C(y, z) & -\Delta_z C(y, z) \\ -\Delta_y C(y, z) & \Delta_y \Delta_z C(y, z) \end{pmatrix}$$

with $C(y, z) := A(y - \frac{1}{2}, z - \frac{1}{2})$ the crossing-or-meeting probability of Theorem 2.3 evaluated at positions $y - \frac{1}{2}$ and $z - \frac{1}{2}$, and $\Delta_y f(y) := f(y+1) - f(y)$ the forward difference.

7. THE BROWNIAN MOTION CASE

The Pfaffian empty-interval formula (Theorem 2.3) applies to any skip-free coalescing system. This section works out the crossing-or-meeting probabilities for Brownian motion on \mathbb{R} —a setting where these probabilities have an explicit closed form. No lattice duality or checkerboard construction is needed; the formula operates directly on the process.

7.1. The crossing-or-meeting probability. Consider coalescing Brownian motions on \mathbb{R} under the maximal entrance law: every point of \mathbb{R} is initially occupied. The paths are continuous, so the process is skip-free. Fix a time $T > 0$. Walls are real-valued: each wall marks a boundary between two adjacent basins of attraction.

For two independent Brownian motions starting at $a < b$, the difference $D(s) = B_b(s) - B_a(s)$ is a Brownian motion starting at $b - a > 0$ with variance $2s$. The particles cross or meet when D hits zero. By the reflection principle,

$$(7.1) \quad \mathbb{P}(\text{crossing or meeting by time } T) = 2\Phi\left(-\frac{b-a}{\sqrt{2T}}\right) = \operatorname{erfc}\left(\frac{b-a}{2\sqrt{T}}\right),$$

where Φ is the standard normal distribution function.

7.2. The Pfaffian formula. Applying Theorem 2.3, the probability that no wall lies in any of the intervals $(a_1, b_1), \dots, (a_n, b_n)$ is $\operatorname{Pf}(A)$, where the $2n \times 2n$ antisymmetric matrix A has entries

$$A_{kl} = \operatorname{erfc}\left(\frac{|x_l - x_k|}{2\sqrt{T}}\right), \quad k < l,$$

and x_1, \dots, x_{2n} are the $2n$ endpoints in order.

For a single interval (a, b) , this gives

$$\mathbb{P}(\text{no wall in } (a, b)) = \operatorname{erfc}\left(\frac{b-a}{2\sqrt{T}}\right),$$

recovering the empty-interval probability of Doering and ben-Avraham [DA88]. For multiple intervals, the Pfaffian of erfc entries recovers the formulas of ben-Avraham [Avr98]. ben-Avraham and Brunet [AB05] used the same empty-interval method to derive consecutive-particle distributions and the thinning relation between coalescence and annihilation correlations.

7.3. The Tribe–Zaboronski kernel. In rescaled coordinates (measuring positions in units of \sqrt{T}), the crossing-or-meeting probability for positions x and y is $\mathcal{F}(|y-x|)$, where

$$\mathcal{F}(z) = \operatorname{erfc}(z/2) = \frac{1}{\sqrt{\pi}} \int_z^\infty e^{-u^2/4} du.$$

Differentiating the Pfaffian empty-interval formula with respect to the interval lengths and taking the zero-length limit (Section 6.1) yields the n -point correlation function as a Pfaffian with a 2×2 matrix kernel:

$$(7.2) \quad K(x, y) = \begin{pmatrix} -\mathcal{F}''(y-x) & -\mathcal{F}'(y-x) \\ \mathcal{F}'(y-x) & \operatorname{sgn}(y-x)\mathcal{F}(|y-x|) \end{pmatrix}.$$

This recovers the Tribe–Zaboronski Pfaffian point process [TZ11; GPTZ18].

Remark 7.1. For Brownian motion with position- and time-dependent diffusion coefficient and drift, the empty-interval probabilities are still given by a Pfaffian (the crossing-or-meeting probabilities depend on the variable coefficients). To extract a 2×2 matrix kernel requires differentiating with respect to the interval lengths, which imposes regularity conditions on the variable coefficients. This analytical problem is beyond the scope of this paper.

Remark 7.2. Coalescing Brownian motions satisfy the assumptions of Theorem 5.3; see Remark 5.4.

8. FROM WALLS TO PARTICLES: CHECKERBOARD DUALITY

This section constructs the checkerboard duality that translates the wall-level results of Sections 2, 3 and 5 into statements about surviving particles.

8.1. The dual forests.

8.1.1. *The lattice and the voter model.* Consider the integer lattice \mathbb{Z}^2 and the half-integer lattice $\mathbb{Z}'^2 = (\mathbb{Z} + \frac{1}{2})^2$, jointly embedded in the (u, v) plane. On each diagonal $u + v = n$ (where n is an integer), the two lattices alternate— \mathbb{Z}^2 -vertices at integer u -coordinates and \mathbb{Z}'^2 -vertices at half-integer u -coordinates—like black and white squares on a checkerboard (rotating the (u, v) plane by 45° makes the two sublattices visually checkerboard-like). In the voter model interpretation [HL75], \mathbb{Z}'^2 -vertices hold *opinions* and \mathbb{Z}^2 -vertices carry *boundaries* where neighboring opinions differ.

8.1.2. *The backward opinion forest.* Each \mathbb{Z}'^2 -vertex $(u + \frac{1}{2}, v + \frac{1}{2})$ copies its opinion from one of two \mathbb{Z}'^2 -neighbors on the previous diagonal:

- from the West neighbor $(u - \frac{1}{2}, v + \frac{1}{2})$ with probability $p_{u,v}$;
- from the South neighbor $(u + \frac{1}{2}, v - \frac{1}{2})$ with probability $1 - p_{u,v}$;

where the weights $p_{u,v} \in [0, 1]$ may vary from vertex to vertex. The backward edge connects each vertex to the neighbor whose opinion it copied. Following backward edges traces the chain of opinion inheritance to a common ancestor. The selected backward edges form the *backward opinion forest* on \mathbb{Z}'^2 , with edges going West or South in the direction of decreasing $u + v$.

8.1.3. *The forward boundary forest.* On diagonal $u + v = n$, a boundary exists at $(u, n - u) \in \mathbb{Z}^2$ when the opinions at the adjacent \mathbb{Z}'^2 -vertices $(u - \frac{1}{2}, n - u + \frac{1}{2})$ and $(u + \frac{1}{2}, n - u - \frac{1}{2})$ disagree. Each backward choice at a \mathbb{Z}'^2 -vertex determines a forward edge at the adjacent \mathbb{Z}^2 -vertex (u, v) : copying from the West produces a forward East edge (from (u, v) to $(u+1, v)$); copying from the South produces a forward North edge (from (u, v) to $(u, v+1)$). These forward edges form the complementary *forward boundary forest* on \mathbb{Z}^2 , with edges in the direction of increasing $u + v$. The two forests are non-crossing: the forward and backward edges at each vertex do not cross in the planar embedding (East pairs with West, North with South; the crossing pairs East–South and North–West are excluded).

Harris [Har78] introduced the graphical construction: a single collection of independent random choices that couples all initial configurations on the same probability space. Arratia [Arr79] observed that the same random choices produce two complementary non-crossing forests on interleaved sublattices (his “percolation substructure”); the forward forest carries boundaries and the backward forest carries opinions.

8.2. Boundaries as coalescing walks.

8.2.1. *Boundary propagation and coalescence.* Since all initial opinions are distinct (the maximal entrance law), every \mathbb{Z}^2 -vertex on diagonal $u + v = 0$ carries a boundary. Each boundary propagates forward along the forward forest: if a boundary exists at (u, v) and the forward edge goes East, then the same boundary appears at $(u+1, v)$ on the next diagonal (because the backward edge copies the opinion from the same side, preserving the disagreement). When two boundary paths meet at the same

vertex, they coalesce into one. The surviving boundaries on \mathbb{Z}^2 form coalescing walks—these are the coalescing random walks of Arratia [Arr79]. Symmetrically, starting with opinions on diagonal $u + v = T$ on \mathbb{Z}'^2 and following the backward edges yields coalescing opinion lineages on \mathbb{Z}'^2 .

At a fixed diagonal $u + v = T$, the coalescing boundaries are precisely the walls of the backward process (Section 2.1): each boundary separates two adjacent opinion basins, and the absence of a boundary between two vertices means their backward lineages have coalesced.

8.2.2. *Time as a coordinate choice.* Time is derived, not primitive: setting $t = u + v$ foliates the lattice into time slices, and different space functions $x = g(u, v)$ yield different particle dynamics (Section 9).

8.2.3. *From walls to particles.* To study the surviving particles at time T , apply the wall-level results of Sections 2, 3 and 5 to the dual backward process, whose walls are precisely the surviving boundaries of the forward process.

8.3. Connection with the spin-pair duality.

8.3.1. *The dual particles.* The spin-pair duality of [TZ11; GPTZ18] arrives at the same mathematical object by a different route. Their scalar kernel $K_t(y, z) = \mathbb{E}[\sigma_{y,z}(\eta_t)]$, where $\sigma_{y,z}(\eta) = (-\theta)^{\eta[y,z]}$ is the spin-pair function, satisfies the PDE $(L_y + L_z) K_t = \partial_t K_t$, where L is the generator of a single particle. Garrod, Poplavskyi, Tribe, and Zaboronski [GPTZ18] describe L as “the generator for a single dual particle.” These dual particles are our backward particles: the PDE says that two independent particles with generator L evolve at positions y and z , which is precisely what the Pfaffian formula computes—crossing-or-meeting probabilities of independent backward particles from the interval endpoints.

8.3.2. *Boundary condition and reconstruction.* The boundary condition $K_t(y, y) = 1$ is automatic in the checkerboard framework: two particles starting at the same position have crossing-or-meeting probability 1. The reconstruction step—passing from the scalar kernel K_t to the 2×2 matrix kernel via discrete differences (Proposition 6.1)—also coincides: in [GPTZ18], the identity $\eta(y) = (1 - \sigma_{y,y+1})/(1 + \theta)$ produces the same forward differences Δ_y applied to the same scalar kernel.

8.3.3. *Algebraic vs. geometric discovery.* The two approaches discover the dual particles by different means: the spin-pair duality finds them algebraically (the generator L emerges from computing the action on spin pairs), while the checkerboard construction makes them geometrically visible as paths on the lattice. For time-homogeneous dynamics with $\theta = 0$, the resulting Pfaffian point processes are identical.

8.3.4. *Non-maximal initial conditions.* The correspondence extends to non-maximal initial conditions. In [GPTZ18], the initial condition η enters through the PDE initial data $K_0(y, z) = \sigma_{y,z}(\eta)$; for $\theta = 0$, this is $K_0(y, z) = \mathbf{1}\{\eta[y, z] = 0\}$, the indicator that the interval $[y, z]$ is initially empty. In the checkerboard framework, the analogous role is played by the clamping step (Section 8.4). Both encode the same data—which intervals are initially empty—by different means.

8.4. General initial conditions.

8.4.1. *Two meanings of “no wall”.* The wall-level results (Sections 2, 3 and 5) assume the maximal entrance law: every site is initially occupied. In applications one often starts from a non-maximal initial condition—for instance, a finite configuration. The proof of the Pfaffian formula (Theorem 2.3) relies on the equivalence: “no wall in (a, b) ” equals “the flanking particles from a and b have coalesced” (Proposition 2.1). Under a general initial condition this equivalence breaks: “no wall in (a, b) ” can also hold when the flanking particles have not coalesced but no initial particle exists between a and b .

8.4.2. *The clamping step.* To restore the equivalence, adjoin one deterministic step to the dual backward process (from time 0 to time -1): clamp together all backward particles that have no initial forward particle between them. After clamping, “no wall in (a, b) ” once again implies that the (augmented) backward particles have coalesced—either during the T random steps (from time T to time 0) or by the clamp. The clamping step is automatically skip-free: a particle cannot jump over another when there is nothing between them to jump over.

The augmented process— T random steps and one deterministic—is an inhomogeneous skip-free process, and the Pfaffian formula [Śni26b] allows arbitrary inhomogeneous transition probabilities, including deterministic ones.

8.5. The augmented kernel.

Theorem 8.1 (Pfaffian formula for general initial conditions). *Consider a coalescing skip-free process with deterministic initial condition η (a subset of occupied sites). Run the system for time T , and let $a_1 < b_1 \leq a_2 < \dots \leq a_n < b_n$ be given. Then:*

$$\mathbb{P}(\text{no wall in any } (a_i, b_i)) = \text{Pf}(A),$$

where A is the $2n \times 2n$ antisymmetric matrix whose entry A_{kl} ($k < l$) is the probability that independent backward particles from x_k and x_l either (a) cross or meet during the T steps, or (b) do not cross or meet but η has no particle between their final positions.

Under the maximal entrance law the second alternative is vacuous, recovering Theorem 2.3.

Proof. The clamping step restores the equivalence between “no wall” and “pairwise coalescence” for the augmented process (Proposition 2.1). The augmented process— T random steps and one deterministic—is an inhomogeneous skip-free process, so the Pfaffian annihilation formula [Śni26b] applies. The crossing-or-meeting probability for the augmented process is exactly the entry A_{kl} described above. \square

Since the augmented process is skip-free and inhomogeneous, the cumulant coloring formula (Theorem 3.1) and the central limit theorem (Theorem 5.1) also hold for the augmented kernel, extending both results to general initial conditions.

9. CHECKERBOARD EXAMPLES

The checkerboard lattice carries no intrinsic notion of space and time. A choice of time function $t = f(u, v)$ and space function $x = g(u, v)$ converts the forests into particle trajectories: on each level set $t = \text{const}$, the vertices become particle positions and forward-forest paths become trajectories. This section works out explicit formulas for several particle systems arising from different coordinate choices.

9.1. Biased random walk. Fix a probability $p \in (0, 1)$ and set $q = 1 - p$. Set all edge weights equal: $p_{u,v} = p$ for every vertex. In Arratia’s coordinates $t = u + v$, $x = u - v$, the forward boundary forest on \mathbb{Z}^2 gives coalescing random walks on \mathbb{Z} : at each time step, a boundary particle jumps $+1$ (East in the (u, v) plane) with probability p , or -1 (North) with probability q .

9.1.1. Backward dynamics. The Pfaffian formula (Theorem 2.3) involves *backward* particles on \mathbb{Z}^2 . Each backward particle follows the edges of the backward forest—these are random edges, determined by the same random choices that generate the forward forest (Section 8.1). The backward particle does not make its own random choices; it traces the pre-existing backward arrows.

When the random choice at $(u, v) \in \mathbb{Z}^2$ selects East (probability p), the complementary backward edge goes West. In (t, x) -coordinates, West means $x \rightarrow x - 1$. When the choice selects North (probability q), the backward edge goes South, meaning $x \rightarrow x + 1$. Since the edge weights are constant, the backward steps are independent and identically distributed, giving a random walk with *swapped probabilities*: the backward particle jumps -1 with probability p and $+1$ with probability q .

Remark 9.1. When $p \neq q$, the backward walk differs from the forward walk: forward boundaries jump right with probability p , while backward particles jump right with probability q . The reversal arises because the forward and backward edges at each vertex point in opposite directions. For $p = q = \frac{1}{2}$ (Section 9.2), the reversal is trivial and both walks are simple symmetric random walks.

For inhomogeneous edge weights $p_{u,v}$, the backward particle follows a path through a random environment where the step probabilities vary from vertex to vertex, determined by the specific weights $p_{u,v}$ encountered along the backward path. The transition weight (9.1) must then be computed for the inhomogeneous probabilities.

9.1.2. The crossing-or-meeting probability. The backward transition weight—the probability that a backward particle at position x on diagonal $u + v = T$ reaches position y on diagonal $u + v = 0$ after T backward steps—is

$$(9.1) \quad w_T(x, y) = \binom{T}{\frac{T+y-x}{2}} q^{\frac{T+y-x}{2}} p^{\frac{T-y+x}{2}},$$

where $\frac{T+y-x}{2}$ counts the rightward jumps (each with probability q in the backward walk), and the expression is zero unless $\frac{T+y-x}{2}$ is a non-negative integer at most T .

For two backward particles at x -positions $x_I < x_J$ on diagonal $u + v = T$, the probability that they cross or meet (Section 2.2; see [Šni26b]) is

$$(9.2) \quad A_{IJ} = 2 \sum_{y_1 < y_2} w_T(x_I, y_2) w_T(x_J, y_1) + \sum_y w_T(x_I, y) w_T(x_J, y),$$

where y_1, y_2 range over all positions reachable in T backward steps. In each term, particle I (starting at the smaller position x_I) ends strictly right of particle J : the paths have crossed. The factor 2 arises because expanding $A_{IJ} = 1 - \det(\dots)$ yields the crossing probability twice (once from the total, once from the sign reversal in the determinant). The second sum counts pairs ending at the same position.

Equivalently, A_{IJ} equals the total probability minus the Karlin–McGregor non-crossing probability [KM59]:

$$A_{IJ} = 1 - \sum_{y_1 < y_2} \det \begin{pmatrix} w_T(x_I, y_1) & w_T(x_I, y_2) \\ w_T(x_J, y_1) & w_T(x_J, y_2) \end{pmatrix}.$$

We set $A_{JI} = -A_{IJ}$ to obtain an antisymmetric matrix.

9.1.3. *The Pfaffian empty-interval formula.* In (t, x) -coordinates, Theorem 2.3 reads as follows. For half-integer positions $a < b$, the 2×2 Pfaffian reduces to a single matrix entry:

$$\mathbb{P}(\text{no wall in } (a, b) \text{ at time } T) = A_{12} \quad \text{with } x_1 = a, x_2 = b.$$

For two disjoint intervals (a_1, b_1) and (a_2, b_2) with $a_1 < b_1 < a_2 < b_2$, the 4×4 Pfaffian expands as

$$\text{Pf}(A) = A_{12} A_{34} - A_{13} A_{24} + A_{14} A_{23},$$

where each A_{IJ} is computed from (9.2). Every entry is an explicit double sum of products of binomial coefficients.

9.2. **Symmetric random walk.** When $p = q = \frac{1}{2}$, the forward and backward dynamics coincide: both are simple symmetric ± 1 random walks. This is the original setting of Arratia [Arr79].

9.2.1. *The crossing-or-meeting probability.* The transition weight (probability of moving from x to y in T steps) is

$$(9.3) \quad w_T(x, y) = \binom{T}{\frac{T+y-x}{2}} \frac{1}{2^T},$$

where the expression is zero unless $\frac{T+y-x}{2}$ is a non-negative integer at most T .

For two independent particles starting at $a < b$, the crossing-or-meeting probability equals 1 minus the Karlin–McGregor [KM59] non-crossing probability:

$$(9.4) \quad A_{ab} = 1 - \sum_{y_1 < y_2} \det \begin{pmatrix} w_T(a, y_1) & w_T(a, y_2) \\ w_T(b, y_1) & w_T(b, y_2) \end{pmatrix}.$$

9.2.2. *The Pfaffian formula.* By Theorem 2.3, the probability that no wall lies in any of the intervals $(a_1, b_1), \dots, (a_n, b_n)$ is $\text{Pf}(A)$, where A is the $2n \times 2n$ antisymmetric matrix with entries (9.4). Every entry is an explicit finite sum of products of binomial coefficients.

Under diffusive scaling, the transition weight w_T converges to the Gaussian density and the crossing-or-meeting probability (9.4) converges to the complementary error function (7.1), recovering the Brownian motion case.

9.3. **Totally asymmetric dynamics.** In coordinates $t = u + v$, $x = u$, a forward boundary particle jumps $+1$ (East) with probability p or stays (North) with probability q . The backward particle jumps -1 or stays, giving the transition weight

$$w_T(x, y) = \binom{T}{x-y} p^{x-y} q^{T-x+y}$$

for $0 \leq x - y \leq T$. Substituting into the Karlin–McGregor complement gives A_{IJ} as a finite double sum of products of binomial coefficients, and Theorem 2.3 yields a Pfaffian empty-interval formula for this totally asymmetric system. All

existing Pfaffian point process results for coalescing particles require bidirectional dynamics [TZ11; GPTZ18].

9.3.1. *Poisson jumps.* With jump probability $\lambda\epsilon$ and step size $\epsilon \rightarrow 0$, the forward particles perform independent Poisson jumps of $+1$ at rate λ and coalesce upon meeting. The backward transition weight becomes

$$w_t(x, y) = e^{-\lambda t} \frac{(\lambda t)^{x-y}}{(x-y)!}$$

for $x \geq y$. For backward particles separated by $\Delta = x_J - x_I$, the Karlin–McGregor complement gives

$$A_{IJ} = 1 - e^{-2\lambda t} \left[I_0(2\lambda t) + 2 \sum_{k=1}^{\Delta-1} I_k(2\lambda t) + I_\Delta(2\lambda t) \right],$$

where I_k is the modified Bessel function of the first kind. To derive this, observe that the separation $D(t) = X_J(t) - X_I(t)$ of two independent backward particles is a symmetric continuous-time random walk on \mathbb{Z} starting at Δ , making ± 1 jumps at rate λ each. Its marginal distribution is [Ske46]

$$\mathbb{P}(D(t) = k) = e^{-2\lambda t} I_{|k-\Delta|}(2\lambda t).$$

Since the walk is skip-free and symmetric, the reflection principle gives

$$A_{IJ} = 2\mathbb{P}(D(t) < 0) + \mathbb{P}(D(t) = 0).$$

Indeed, every path from Δ to $k > 0$ that touches 0 bijects with a path from $-\Delta$ to k , which by symmetry has the same weight as a path from Δ to $-k$. Summing the tail and using $\sum_{k=-\infty}^{\infty} \mathbb{P}(D(t) = k) = 1$ gives the Bessel formula.

9.4. **Bidirectional Poisson coalescence.** In coordinates $t = u + v$, $x = u - \lfloor t/2 \rfloor$, a boundary particle jumps $+1$ on even time steps and -1 on odd time steps. With alternating edge weights $p_{u,v} = \lambda_+\epsilon$ on even diagonals and $p_{u,v} = 1 - \lambda_-\epsilon$ on odd diagonals, the $\epsilon \rightarrow 0$ limit gives coalescing particles with independent Poisson jumps of $+1$ at rate λ_+ and -1 at rate λ_- . The backward particle has swapped rates ($+1$ at rate λ_- , -1 at rate λ_+), giving the transition weight

$$w_t(x, y) = e^{-\lambda t} \left(\frac{\lambda_-}{\lambda_+} \right)^{(y-x)/2} I_{|y-x|} \left(2\sqrt{\lambda_+\lambda_-} t \right),$$

where $\lambda = \lambda_+ + \lambda_-$ is the total jump rate. The crossing-or-meeting probability depends only on λ , not on the individual rates: the difference of two independent particles is a symmetric process regardless of the drift. The Karlin–McGregor complement gives

$$A_{IJ} = 1 - e^{-2\lambda t} \left[I_0(2\lambda t) + 2 \sum_{k=1}^{\Delta-1} I_k(2\lambda t) + I_\Delta(2\lambda t) \right]$$

with $\Delta = x_J - x_I$. This coincides with the totally asymmetric formula (Section 9.3): in both cases the difference of two independent backward particles performs the same symmetric random walk.

More generally, the rates λ_+ and λ_- may vary with position and time, modeling the continuous-time voter model with inhomogeneous copying rates. The Pfaffian formula (Theorem 2.3) still applies; the transition weight w_t must then be computed for the inhomogeneous rates along each backward path.

APPENDIX A. CLT FOR SUMS OF INDICATORS VIA IDEMPOTENCE

This appendix proves a central limit theorem for sums of dependent 0-1 random variables, using only the joint cumulants of *distinct* indices. The key simplification is that indicator variables are idempotent ($X_i^2 = X_i$), which collapses repeated-index cumulants to distinct-index ones. The ingredients are individually well known (multilinearity of cumulants, the idempotent multinomial identity, and the method of cumulants). In the continuous setting, the analogous reduction is classical: factorial cumulant measures (see Daley and Vere-Jones [DVJ08, Chapter 5]) are defined on tuples of distinct points, and since the diagonal has measure zero in \mathbb{R}^k , ordinary and factorial cumulants of a count coincide. In the discrete setting, however, repeated indices contribute nontrivially, and the idempotence reduction below handles them; we include the short proof for completeness.

A.1. The CLT criterion.

Proposition A.1. *For each $n \geq 1$, let X_1, \dots, X_n be 0-1 random variables (possibly dependent), and let $S_n = X_1 + \dots + X_n$. Write $\kappa_A = \kappa(X_i : i \in A)$ for the joint cumulant indexed by a set $A \subseteq [n]$ of distinct indices. Assume:*

(i) *For each fixed $L \geq 1$,*

$$\sup_{\substack{i \in [n] \\ A \subseteq [n] \\ i \in A, |A|=L}} |\kappa_A| = O(1),$$

where the implied constant may depend on L but not on n .

(ii) $\text{Var}(S_n) \asymp n$.

Then for each fixed $m \geq 1$, $\kappa_m(S_n) = O(n)$, and

$$\frac{S_n - \mathbb{E}[S_n]}{\sqrt{\text{Var}(S_n)}} \xrightarrow{d} \mathcal{N}(0, 1) \quad \text{as } n \rightarrow \infty.$$

A.2. Proof via idempotence. By multilinearity, $\kappa_m(S_n)$ is a sum over tuples $(\alpha_1, \dots, \alpha_m) \in [n]^m$. Idempotence ($X_i^2 = X_i$) ensures that every repeated-index cumulant can be expressed as a polynomial in the distinct-index cumulants κ_A and the marginal probabilities $p_i = \mathbb{E}X_i$. The following example illustrates this reduction for $m = 2$ and 3; although the resulting formulas are explicit, they grow rapidly in complexity with m , and we do not attempt to find closed forms. Instead, Lemma A.3 below provides the bound $\kappa_m(S_n) = O(n)$ that suffices for the CLT, bypassing the need for exact expressions.

Example A.2 (Small cases). Write $p_i = \mathbb{E}X_i$.

For $m = 2$, multilinearity gives $\kappa_2(S_n) = \sum_{i,j} \text{Cov}(X_i, X_j)$. On the diagonal, idempotence ($\mathbb{E}X_i^2 = p_i$) yields $\text{Cov}(X_i, X_i) = p_i(1 - p_i)$, so

$$\kappa_2(S_n) = \sum_i p_i(1 - p_i) + \sum_{i \neq j} \text{Cov}(X_i, X_j).$$

For $m = 3$, the Leonov–Shiryaev formula and idempotence reduce every repeated-index cumulant to lower-order ones:

$$\begin{aligned} \kappa(X_i, X_i, X_k) &= (1 - 2p_i) \text{Cov}(X_i, X_k), \\ \kappa_3(X_i) &= p_i(1 - p_i)(1 - 2p_i). \end{aligned}$$

Thus

$$\begin{aligned} \kappa_3(S_n) &= \sum_{\substack{i,j,k \\ \text{distinct}}} \kappa(X_i, X_j, X_k) \\ &\quad + 3 \sum_{i \neq j} (1 - p_i - p_j) \text{Cov}(X_i, X_j) \\ &\quad + \sum_i p_i(1 - p_i)(1 - 2p_i). \end{aligned}$$

Every term is determined by the distinct-index cumulants κ_A and the marginal probabilities p_i .

Lemma A.3 (Cumulant bound for idempotent variables). *Under assumption (i) of Proposition A.1, for each fixed $m \geq 1$ define*

$$\kappa_m^{*,i}(n) := \sum_{\alpha_2, \dots, \alpha_m \in [n]} |\kappa(X_i, X_{\alpha_2}, \dots, X_{\alpha_m})|.$$

Then

$$\sup_{i \in [n]} \kappa_m^{*,i}(n) = O(1),$$

where the implied constant depends on m but not on n .

Proof of Lemma A.3. We proceed by induction on m . The base case $m = 1$ is immediate: $\kappa_1^{*,i}(n) = |\mathbb{E}X_i| \leq 1$ for every i .

For the inductive step, fix $i \in [n]$ and bound $\kappa_m^{*,i}(n)$. Cover the tuples $(\alpha_2, \dots, \alpha_m)$ by $\binom{m}{2} + 1$ classes (not necessarily disjoint). For each pair of positions $1 \leq p < q \leq m$, let

$$T_{pq} = \{(\alpha_2, \dots, \alpha_m) \in [n]^{m-1} : \text{positions } p \text{ and } q \text{ carry the same value}\},$$

where position 1 carries the fixed value i . Every tuple with at least one repeated index belongs to some T_{pq} , so

$$(A.1) \quad \kappa_m^{*,i}(n) \leq \underbrace{(m-1)! \sum_{\substack{A \subseteq [n], i \in A \\ |A|=m}} |\kappa_A|}_{\text{all distinct: } O(1)} + \sum_{1 \leq p < q \leq m} T_{pq}^{*,i},$$

where $T_{pq}^{*,i}$ is the sum of $|\kappa(X_i, X_{\alpha_2}, \dots, X_{\alpha_m})|$ over T_{pq} . The all-distinct term is $O(1)$ by assumption (i). It remains to show $T_{pq}^{*,i} = O(1)$ for each pair (p, q) .

Idempotence reduction. Fix a pair (p, q) and a tuple $(\alpha_2, \dots, \alpha_m) \in T_{pq}$. Write $r = \alpha_p = \alpha_q$ for the repeated value (with the convention $\alpha_1 = i$), and write Z_1, \dots, Z_{m-2} for the variables X_{α_s} , $s \in [m] \setminus \{p, q\}$, in increasing order of s . (When $p \geq 2$, the pinned variable X_i is one of the Z_j .) The m -argument cumulant reads

$$\kappa(X_{\alpha_1}, \dots, X_{\alpha_m}) = \kappa(X_r, X_r, Z_1, \dots, Z_{m-2}),$$

since X_r occupies both positions p and q .

The Leonov–Shiryaev product formula [LS59, Theorem 1] (see also Speed [Spe83, Proposition 4.3]) gives

$$(A.2) \quad \kappa(X_r \cdot X_r, Z_1, \dots, Z_{m-2}) = \kappa(X_r, X_r, Z_1, \dots, Z_{m-2}) + R,$$

where

$$(A.3) \quad R = \sum_{S \sqcup T = [m-2]} \kappa(X_r, Z_s : s \in S) \cdot \kappa(X_r, Z_t : t \in T).$$

The sum runs over all 2^{m-2} ordered partitions of $[m-2]$ into two blocks (either block may be empty, giving $\kappa_1 = \mathbb{E}$). Every term in R is a product of two cumulants whose orders sum to m and are each strictly less than m .

By idempotence ($X_r^2 = X_r$), the left-hand side of (A.2) equals $\kappa(X_r, Z_1, \dots, Z_{m-2})$. Rearranging:

$$(A.4) \quad \kappa(X_r, X_r, Z_1, \dots, Z_{m-2}) = \kappa(X_r, Z_1, \dots, Z_{m-2}) - R.$$

This is the key step: the m -argument cumulant (left) equals an $(m-1)$ -argument cumulant minus the remainder R . Taking absolute values and summing over all tuples $(\alpha_2, \dots, \alpha_m) \in T_{pq}$:

$$T_{pq}^{*,i} \leq \sum_{(\alpha_2, \dots, \alpha_m) \in T_{pq}} |\kappa(X_r, Z_1, \dots, Z_{m-2})| + \sum_{(\alpha_2, \dots, \alpha_m) \in T_{pq}} |R|.$$

Bounding the first sum. The $(m-1)$ -argument cumulant $\kappa(X_r, Z_1, \dots, Z_{m-2})$ has $m-2$ free indices (the constraint $\alpha_p = \alpha_q$ eliminates one).

If $p = 1$ (so $r = i$), the cumulant is $\kappa(X_i, Z_1, \dots, Z_{m-2})$; summing over the $m-2$ free indices gives at most $\kappa_{m-1}^{*,i}(n) = O(1)$ by induction.

If $p \geq 2$, the pinned variable X_i is one of the Z_j , so the cumulant still involves X_i ; summing over the $m-2$ free indices (including r) again gives at most $\kappa_{m-1}^{*,i}(n) = O(1)$.

Bounding the remainder sum. Each term in R (see (A.3)) is

$$|\kappa(X_r, Z_s : s \in S)| \cdot |\kappa(X_r, Z_t : t \in T)|$$

with $S \sqcup T = [m-2]$. Both factors contain X_r ; the pinned variable X_i sits in one factor (or equals X_r when $p = 1$).

Case $p = 1$ ($r = i$): both factors have X_i as their first argument, and the free indices in S and T are disjoint. The sum factors:

$$\sum_{(\alpha_2, \dots, \alpha_m) \in T_{1q}} |\kappa(X_i, Z_S)| \cdot |\kappa(X_i, Z_T)| = \kappa_{|S|+1}^{*,i} \cdot \kappa_{|T|+1}^{*,i} = O(1)$$

by the inductive hypothesis (both orders are $< m$).

Case $p \geq 2$: X_i is one of the Z_j ; say $Z_{j_0} = X_i$ belongs to the S -factor. Write $S_0 = S \setminus \{j_0\}$, so the first factor is $\kappa(X_r, X_i, Z_s : s \in S_0)$ and the second is $\kappa(X_r, Z_t : t \in T)$. Sum the second factor over its free indices (for each fixed r) to get at most $\sup_j \kappa_{|T|+1}^{*,j}(n) = O(1)$; then sum the first factor over its free indices (including r) to get at most $\kappa_{|S_0|+2}^{*,i}(n) = O(1)$. The product is $O(1)$.

Since R has 2^{m-2} terms (a constant depending only on m), $\sum_{T_{pq}} |R| = O(1)$.

Combining: $T_{pq}^{*,i} = O(1)$ for each pair (p, q) , uniformly in i , and (A.1) gives $\sup_i \kappa_m^{*,i}(n) = O(1)$. \square

Remark A.4 (Growth of the implied constant). Write $E_m = \sup_{i,n} \kappa_m^{*,i}(n)$ for the implied constant in Lemma A.3, and D_m for the distinct-index per-site bound from

assumption (i). The proof gives the recursion

$$(A.5) \quad E_m \leq (m-1)! D_m + \binom{m}{2} \left[E_{m-1} + 2^{m-2} \max_{\substack{j+l=m \\ j,l \geq 1}} E_j E_l \right].$$

Setting $F_m = \log E_m$, the dominant term is $\max_{j+l=m} (F_j + F_l) + O(m)$. Evaluating at $j = l = m/2$ gives $F_m \leq 2F_{m/2} + O(m)$, so $F_m = O(m \log m)$ by the master theorem. In particular, if D_m is bounded by $e^{O(m \log m)}$, then so is E_m .

Proof of Proposition A.1. By multilinearity of cumulants,

$$\kappa_m(S_n) = \sum_{\alpha_1, \dots, \alpha_m \in [n]} \kappa(X_{\alpha_1}, \dots, X_{\alpha_m}).$$

Taking absolute values and grouping by α_1 ,

$$|\kappa_m(S_n)| \leq \sum_{i \in [n]} \kappa_m^{*,i}(n) = O(n)$$

by Lemma A.3. For the standardized variable $\tilde{S}_n = (S_n - \mathbb{E}S_n)/\sqrt{\text{Var } S_n}$:

$$\kappa_m(\tilde{S}_n) = \frac{\kappa_m(S_n)}{(\text{Var } S_n)^{m/2}} = O(n^{1-m/2}) \rightarrow 0 \quad \text{for } m \geq 3,$$

while $\kappa_1(\tilde{S}_n) = 0$ and $\kappa_2(\tilde{S}_n) = 1$. Since all cumulants of order ≥ 3 vanish, the method of cumulants gives $\tilde{S}_n \xrightarrow{d} \mathcal{N}(0, 1)$. \square

A.3. Remarks.

Remark A.5 (Translation-invariant processes). For translation-invariant indicators (e.g., a stationary point process on $\mathbb{Z} \cap [1, n]$), assumption (i) follows from the weaker global bound $\sum_{|A|=L} |\kappa_A| = O(n)$: by stationarity, each i contributes equally to $\sum_i \sum_{A \ni i, |A|=L} |\kappa_A| = L \sum_{|A|=L} |\kappa_A|$, so the per-site sum is $O(1)$.

Remark A.6 (Berry–Esseen bound). The bound $|\kappa_m(S_n)| = O(n)$ from Lemma A.3 gives $|\kappa_m(\tilde{S}_n)| = O(n^{1-m/2})$ for the standardized sum. Since $|X_i| \leq 1$, this satisfies the Statulevičius condition $|\kappa_j(\tilde{S}_n)| \leq j^{1+\gamma}/\Delta^{j-2}$ with $\Delta \asymp \sqrt{n}$ and suitable γ (the boundedness of the indicators ensures $\gamma < \infty$). This yields a Berry–Esseen bound $\sup_x |F_{\tilde{S}_n}(x) - \Phi(x)| = O(n^{-1/2})$ via Corollary 2.1 of Saulis and Statulevičius [SS91], and moderate deviation estimates via their Lemma 2.3 (due to Rudzkis, Saulis, and Statulevičius [RSS78]); see also Döring, Jansen, and Schubert [DJS22], Sections 2.1–2.2.

Remark A.7 (Small occupation probability). The growth $E_m = e^{O(m \log m)}$ in Remark A.4 arises from the Leonov–Shiryaev remainder (A.3), which introduces the product $E_j E_l$ in the recursion (A.5). When $\rho := \sup_i \mathbb{P}(X_i = 1) \rightarrow 0$ (for example, after rescaling from a lattice to a continuous process), the correction from repeated indices is suppressed.

More precisely, the m -th cumulant of $S_n = \sum X_i$ admits the exact decomposition $\kappa_m(S_n) = \sum_{r=1}^m S(m, r) C_r$, where $S(m, r)$ are Stirling numbers of the second kind and C_r is the r -th factorial cumulant of S_n (see Saulis and Statulevičius [SS91], Chapter 1). For 0-1 variables, idempotence gives $\mathbb{E}[S_n^{(r)}] = \sum_{\text{distinct}} \mathbb{E}[\prod X_i]$, so the factorial cumulants C_r approximate the distinct-index sums with corrections of order $O(\rho)$. The exponential generating function $\sum_m B_m(x) z^m/m! = \exp(x(e^z - 1))$ is entire, so $B_m(x) = \sum_r S(m, r) x^r$ grows slower than $C^m \cdot m!$ for any $C > 0$.

In particular, if the distinct-index per-site bound D_m grows at most exponentially in m and $\rho \rightarrow 0$, then the Statulevičius condition holds with $\gamma = 0$, giving Cramér-type large deviation estimates.


Remark A.8 (Role of idempotence). The inductive step (A.4) uses idempotence ($X_r^2 = X_r$) to replace $X_r \cdot X_r$ on the left-hand side of (A.2) by a single X_r , reducing the m -argument cumulant to an $(m-1)$ -argument one. For general bounded random variables, $\kappa(X_r \cdot X_r, Z_1, \dots)$ would be a genuinely different object from any ordinary cumulant, and the induction would not close without stronger summability assumptions.

APPENDIX B. NONZERO COLORINGS FOR $k = 3$

For $k = 3$, there are 90 colorings (sequences of length 6 over $\{A, B, C\}$ with each color used exactly twice), of which 14 have $\text{wt}(c) \neq 0$ (Definition 4.3). Each diagram below shows the six simplex positions $y_1 < \dots < y_6$ (top) connected to the three walls A, B, C (bottom) according to the coloring c . Wall A : solid blue; wall B : dashed red; wall C : dotted green.



The *mirror symmetry* reverses the coloring sequence and relabels $A \leftrightarrow C$, $B \leftrightarrow B$. The coefficient $\text{wt}(c)$ is invariant under this symmetry. The 14 nonzero colorings split into 4 self-mirror colorings and 5 mirror pairs.


B.1. The unique negative coloring ($\text{wt} = -8$). The only coloring with $\text{wt}(c) < 0$ is self-mirror:







B.2. Positive colorings ($\text{wt} = +8$). The remaining 13 nonzero colorings all have $\text{wt}(c) = +8$.


B.2.1. Self-mirror.





B.2.2. *Mirror pairs.*





ACKNOWLEDGMENTS

We thank Theodoros Assiotis, Balázs Bárány, Maciej Dołęga, Sho Matsumoto, Bálint Tóth, Ákos Urbán, Oleg Zaboronski, and Karol Życzkowski for stimulating discussions and helpful literature suggestions.

We thank Richard Arratia for generously providing access to his PhD thesis [Arr79].

Claude Code (Anthropic) was used as an assistant during manuscript preparation.

REFERENCES

- [AB05] Daniel ben Avraham and Éric Brunet. “On the relation between one-species diffusion-limited coalescence and annihilation in one dimension”. In: *J. Phys. A: Math. Gen.* 38 (2005), pp. 3247–3252. DOI: [10.1088/0305-4470/38/15/001](https://doi.org/10.1088/0305-4470/38/15/001). arXiv: [cond-mat/0412745](https://arxiv.org/abs/cond-mat/0412745).
- [Arr79] Richard Arratia. “Coalescing Brownian motions on the line”. PhD thesis. University of Wisconsin–Madison, 1979.
- [AS12] Siva R. Athreya and Jan M. Swart. “Systems of branching, annihilating, and coalescing particles”. In: *Electron. J. Probab.* 17.80 (2012), pp. 1–32. DOI: [10.1214/EJP.v17-2003](https://doi.org/10.1214/EJP.v17-2003). arXiv: [1203.6477](https://arxiv.org/abs/1203.6477).
- [Avr98] Daniel ben Avraham. “Complete exact solution of diffusion-limited coalescence, $A + A \rightarrow A$ ”. In: *Phys. Rev. Lett.* 81 (1998), pp. 4756–4759. DOI: [10.1103/PhysRevLett.81.4756](https://doi.org/10.1103/PhysRevLett.81.4756). arXiv: [cond-mat/9803281](https://arxiv.org/abs/cond-mat/9803281).
- [DA88] Charles R. Doering and Daniel ben Avraham. “Interparticle distribution functions and rate equations for diffusion-limited reactions”. In: *Phys. Rev. A* 38 (1988), pp. 3035–3042. DOI: [10.1103/PhysRevA.38.3035](https://doi.org/10.1103/PhysRevA.38.3035).
- [DJS22] Hanna Döring, Sabine Jansen, and Kristina Schubert. “The method of cumulants for the normal approximation”. In: *Probab. Surv.* 19 (2022), pp. 185–270. DOI: [10.1214/22-PS7](https://doi.org/10.1214/22-PS7). arXiv: [2102.01459](https://arxiv.org/abs/2102.01459).
- [DVJ08] D. J. Daley and D. Vere-Jones. *An Introduction to the Theory of Point Processes. Vol. II: General Theory and Structure*. 2nd. New York: Springer, 2008.
- [DZ96] Bernard Derrida and Reuven Zeitak. “Distribution of domain sizes in the zero temperature Glauber dynamics of the one-dimensional Potts model”. In: *Phys. Rev. E* 54.3 (1996), pp. 2513–2525. DOI: [10.1103/PhysRevE.54.2513](https://doi.org/10.1103/PhysRevE.54.2513).
- [EK86] Stewart N. Ethier and Thomas G. Kurtz. *Markov Processes: Characterization and Convergence*. Wiley Series in Probability and Mathematical Statistics. Wiley, 1986. DOI: [10.1002/9780470316658](https://doi.org/10.1002/9780470316658).

- [FTZ22] Will FitzGerald, Roger Tribe, and Oleg Zaboronski. “Asymptotic expansions for a class of Fredholm Pfaffians and interacting particle systems”. In: *Ann. Probab.* 50.6 (2022), pp. 2409–2474. DOI: [10.1214/22-AOP1586](https://doi.org/10.1214/22-AOP1586). arXiv: [2107.14504](https://arxiv.org/abs/2107.14504).
- [GF17] E. V. Glinyanaya and V. V. Fomichov. “The central limit theorem for the number of clusters of the Arratia flow”. In: *Theory of Stochastic Processes* 22(38).2 (2017), pp. 1–7. arXiv: [1712.05098](https://arxiv.org/abs/1712.05098).
- [GPTZ18] Barnaby Garrod, Mihail Poplavskiy, Roger Tribe, and Oleg Zaboronski. “Examples of interacting particle systems on \mathbb{Z} as Pfaffian point processes: annihilating and coalescing random walks”. In: *Ann. Henri Poincaré* 19 (2018), pp. 3635–3662. DOI: [10.1007/s00023-018-0719-x](https://doi.org/10.1007/s00023-018-0719-x).
- [Gri79] David Griffeath. *Additive and Cancellative Interacting Particle Systems*. Vol. 724. Lecture Notes in Mathematics. Springer, 1979.
- [Har78] T. E. Harris. “Additive set-valued Markov processes and graphical methods”. In: *Ann. Probab.* 6.3 (1978), pp. 355–378. DOI: [10.1214/aop/1176995523](https://doi.org/10.1214/aop/1176995523).
- [HL75] Richard A. Holley and Thomas M. Liggett. “Ergodic theorems for weakly interacting infinite systems and the voter model”. In: *Ann. Probab.* 3.4 (1975), pp. 643–663. DOI: [10.1214/aop/1176996306](https://doi.org/10.1214/aop/1176996306).
- [IL71] I. A. Ibragimov and Yu. V. Linnik. *Independent and Stationary Sequences of Random Variables*. Groningen: Wolters-Noordhoff, 1971.
- [KM59] Samuel Karlin and James McGregor. “Coincidence probabilities”. In: *Pacific J. Math.* 9.4 (1959), pp. 1141–1164. DOI: [10.2140/pjm.1959.9.1141](https://doi.org/10.2140/pjm.1959.9.1141).
- [LQW21] Zhaofeng Lin, Yanqi Qiu, and Kai Wang. “Gaussian limit for determinantal point processes with J -Hermitian kernels”. In: (2021). arXiv: [2101.00384](https://arxiv.org/abs/2101.00384).
- [LS59] V. P. Leonov and A. N. Shiryaev. “On a method of calculation of semi-invariants”. In: *Theory Probab. Appl.* 4.3 (1959), pp. 319–329. DOI: [10.1137/1104031](https://doi.org/10.1137/1104031).
- [PT11] Giovanni Peccati and Murad S. Taqqu. *Wiener Chaos: Moments, Cumulants and Diagrams*. Vol. 1. Bocconi & Springer Series. Springer, 2011. DOI: [10.1007/978-88-470-1679-8](https://doi.org/10.1007/978-88-470-1679-8).
- [RSS78] R. Rudzikis, L. Saulis, and V. Statulevičius. “A general lemma on probabilities of large deviations”. In: *Lithuanian Math. J.* 18 (1978), pp. 226–238. DOI: [10.1007/BF00972235](https://doi.org/10.1007/BF00972235).
- [Ske46] J. G. Skellam. “The frequency distribution of the difference between two Poisson variates belonging to different populations”. In: *J. Roy. Statist. Soc. Ser. A* 109.3 (1946), p. 296. DOI: [10.2307/2981372](https://doi.org/10.2307/2981372).
- [Śni26a] Piotr Śniady. “Coalescing random walks via the coalescence determinant”. Preprint. 2026. arXiv: [2602.20043](https://arxiv.org/abs/2602.20043) [[math.PR](https://arxiv.org/abs/2602.20043)].
- [Śni26b] Piotr Śniady. “Determinant and Pfaffian formulas for particle annihilation”. Preprint. 2026. arXiv: [2602.13183](https://arxiv.org/abs/2602.13183) [[math.PR](https://arxiv.org/abs/2602.13183)].
- [Sos02] Alexander Soshnikov. “Gaussian limit for determinantal random point fields”. In: *Ann. Probab.* 30.1 (2002), pp. 171–187. DOI: [10.1214/aop/1020107764](https://doi.org/10.1214/aop/1020107764).

- [Spe83] T. P. Speed. “Cumulants and partition lattices”. In: *Austral. J. Statist.* 25.2 (1983), pp. 378–388. DOI: [10.1111/j.1467-842X.1983.tb00391.x](https://doi.org/10.1111/j.1467-842X.1983.tb00391.x).
- [SS91] Leonas Saulis and Vytautas A. Statulevičius. *Limit Theorems for Large Deviations*. Kluwer, 1991. DOI: [10.1007/978-94-011-3530-6](https://doi.org/10.1007/978-94-011-3530-6).
- [ŚU26] Piotr Śniady and Ákos Urbán. “Exact determinant formulas for coalescing particle systems”. Preprint. 2026. arXiv: [2602.10782](https://arxiv.org/abs/2602.10782) [[math.PR](#)].
- [TZ11] Roger Tribe and Oleg Zaboronski. “Pfaffian formulae for one-dimensional coalescing and annihilating systems”. In: *Electron. J. Probab.* 16 (2011), pp. 2080–2103. DOI: [10.1214/EJP.v16-942](https://doi.org/10.1214/EJP.v16-942).
- [TZ26] Roger Tribe and Oleg Zaboronski. “Entrance laws for coalescing and annihilating Brownian motions”. Preprint. 2026. arXiv: [2602.16509](https://arxiv.org/abs/2602.16509) [[math.PR](#)].
- [Urb25] Ákos Urbán. “The Pólya Web”. MA thesis. Budapest University of Technology and Economics, 2025. arXiv: [2601.12172](https://arxiv.org/abs/2601.12172) [[math.PR](#)].
- [WX25] Kai Wang and Mei Xu. “Gaussian limit for Pfaffian point processes”. In: (2025). arXiv: [2504.14487](https://arxiv.org/abs/2504.14487).

INSTITUTE OF MATHEMATICS, POLISH ACADEMY OF SCIENCES, UL. ŚNIADECKICH 8,
00-656 WARSZAWA, POLAND
Email address: psniady@impan.pl