

Balanced two-type annihilation: mean-field asymptotics

John Haslegrave*

Peter Keevash[†]

May 6, 2026

Abstract

We consider an interacting particle system where equal-sized populations of two types of particles move by random walk steps on a graph, the two types may have different speeds, and meetings of opposite-type particles result in annihilation. The key quantity of interest is the expected extinction time. Even for the mean-field setting of complete graphs, the correct order of magnitude was not previously known. Under essentially optimal assumptions on the starting configuration, we determine not only the order of magnitude but also the asymptotics: the expected extinction time on K_{2n} is $(2 + o(1))n \log n$, independently of the relative speeds of the two types.

Keywords: interacting particle system; reaction model; annihilation; random walk; complete graph
MSC2020 subject classifications: 05C81; 82C22; 60K35.

1 Introduction

Interacting particle systems have been intensively studied since the 1970s to model a variety of phenomena in Statistical Physics and Mathematical Biology, such as spin systems, chemical reactions, population dynamics and the spread of infections. An important example is the two-type annihilation model for chemical reactions, studied classically on integer lattices \mathbb{Z}^d , where a series of papers by Bramson and Lebowitz [3, 4, 5] gives order of magnitude estimates for the site occupancy probability. In this model, there are two types of particles and interactions only occur between two particles of different type, which makes the analysis particularly difficult. A single type model introduced by Erdős and Ney [9] is easier to analyse, although even here the basic question of whether the origin would be occupied infinitely often was not solved in the original paper, but was subsequently resolved in one dimension by Lootgieter [17] and in higher dimensions by Arratia [1, 2]. Besides these diffusive models, ballistic annihilation models have also been studied in the physics literature since the 1990s [8, 15], and more recently analysed rigorously in the mathematics literature [14].

Considering the real-world motivations for interacting particle systems, it is natural to consider systems on finite graphs and ask for the asymptotics of key parameters when the number n of vertices is large. While a wide range of interacting particle systems have been well studied in a lattice setting (see e.g. [16] for an overview), the literature covering interacting particle systems on general, or random, finite graphs is less developed, and has mainly covered opinion-forming or information-spreading processes (see [6] for an overview), with one of the main models considered in this setting being the voter model, considered on general graphs by Donnelly and Welsh in the 1980s [12]. Consequently there has been significant interest in the dual process

*Department of Mathematics and Statistics, Lancaster University, UK. Supported by ERC Advanced Grant 883810.

[†]Mathematical Institute, University of Oxford, UK. Supported by ERC Advanced Grant 883810.

of coalescing random walks on finite graphs (e.g. [10]). Coalescence is simpler to analyse than annihilation due to methods based on duality and monotonicity. A more general framework, with a variety of models including annihilation, was analysed by Cooper, Frieze and Radzik [11], assuming that (a) the initial populations are polynomially small and randomly distributed, (b) the underlying geometry is a random regular graph, and (c) the types move simultaneously with equal speeds. By contrast, we will consider an annihilation model with individual rather than simultaneous movements starting from an arbitrary spanning configuration, so we cannot avoid considering the so-called ‘big bang regime’ in which there is rapid action at the beginning of the process that is harder to analyse.

We refer to our companion paper [13] for a more detailed study of particle systems on expanders, with more history and new results for a variety of models, including the balanced two-type annihilation to be considered here; our main result here will be a more precise asymptotic result in the mean-field setting. Our precise formulation of the model is taken from a recent paper by Cristali, Jiang, Junge, Kassem, Sivakoff and York [7]. Fix a graph with $2n$ vertices, and initialise with one particle at each vertex, with n particles being blue and n red. The particles represent molecules of two reactants; when a red and a blue particle occupy the same vertex, they react and are removed from the system, i.e. they mutually annihilate. At each time step, with probability p a red particle is chosen, uniformly at random from the remaining red particles, or otherwise a uniformly random blue particle is chosen. Without loss of generality, we assume $p \leq 1/2$. The chosen particle performs a simple random walk step. If it reaches a vertex with one or more particles of the opposite colour, it mutually annihilates with one such particle. This process almost surely eventually terminates with no particles remaining.

The key quantity of interest is the time taken for this to happen, which we call the *extinction time*. In analysing this we use standard asymptotic notation, with all limits applying as $n \rightarrow \infty$ unless otherwise stated: $f(n) = o(g(n))$ means $f(n)/g(n) \rightarrow 0$; $f(n) = \omega(g(n))$ means $f(n)/g(n) \rightarrow +\infty$; $f(n) = O(g(n))$ means $\limsup |f(n)/g(n)| < \infty$; and $f(n) = \Theta(g(n))$ means $0 < \liminf |f(n)/g(n)| \leq \limsup |f(n)/g(n)| < \infty$.

The above model was considered by Cristali et al. [7] for the complete graph K_{2n} (the mean-field case) and the star. For these graphs, the one-type annihilation model is easy to analyse, but the two-type process is much harder, owing to the possibility of multiple particles of the same type occupying a single site. While they were able to give precise results on the extinction time for stars, for the mean-field case, also assuming that p is bounded away from zero uniformly in n , they gave a lower bound of $2n \log n$ and an upper bound of $20n \log^2 n / \log \log n$, thus leaving the open problem of showing that it is $\Theta(n \log n)$. In [13] we solved this problem for any p on all regular graphs with sufficiently strong spectral expansion. Here we use different methods to give the following more precise result in the mean-field setting.

Theorem 1. *Let $p = p(n) \in [0, 1/2]$. Consider an initial possibly random configuration on K_{2n} having n particles of each type, with each site having at most one type of particle, where $n - A$ sites are initially occupied by red particles. Let T be the extinction time. If $\mathbb{E}(A) = o(pn \log n)$ then $\mathbb{E}(T) = (2 + o(1))n \log n$ as $n \rightarrow \infty$.*

Note that if $p = \omega(1/\log n)$, i.e. the reds are not that slow, then Theorem 1 applies to any starting configuration. The assumption on A cannot be significantly relaxed, as the result is not true for all starting configurations when p is of order $1/\log n$. To see this, notice that if all reds start on the same vertex v then extinction requires at least n occasions where either a red moves or a blue moves onto v . The number of these events in time $3n \log n$ is with high probability at most $3pn \log n + O(\log n)$, so if $p = 1/(4 \log n)$ then extinction will not occur.

The idea of the proof is to control the distribution of the particles by martingale arguments, categorising it as good, intermediate or bad for each colour separately. If both distributions

remain good or intermediate until very few particles remain, this will imply sufficiently good bounds on annihilation probabilities to give the desired result. We show that it is unlikely that random walk steps alone can cause a colour to cross the intermediate regime until late in the process. For blue particles, this also applies when we take annihilations into account, since they occur less frequently than blue walk steps, but for rarely-moving red particles it is possible that annihilations have a dominant effect. We thus bound both the time taken for the red distribution to recover from the effects of each annihilation, and the total number of annihilations that occur while the red distribution is bad. This part of the proof could be avoided if we assumed bounded speed ratios as in [7], but this would be an unnecessary and moreover a strange assumption, as the limiting case when one colour is stationary seems intuitively easier! The potential difficulty that must be avoided when p is small but nonzero is that the process may be slow to recover from a situation where due to blue annihilations the red particles have somehow become clustered on a few vertices.

2 Proof

Consider an initial possibly random configuration on K_{2n} having n particles of each type, with each site having at most one type of particle, where $n - A$ sites are initially occupied by red particles. For simplicity we will add a loop at every vertex, so that each random walk step is equally likely to go to any vertex, including the current vertex. This produces a modified process that is just a time change by a factor $1 + O(1/n)$, so this makes no difference to our result. We also assume throughout that n is sufficiently large.

At time t , we write M_t for the number of particles of each colour, R_t for the number of sites with red particles, B_t for the number of sites with blue particles, and \mathcal{F}_t for the σ -algebra generated by all previous movements of particles. When tracking R_t , we consider the effects of red and blue movements separately. To that end, we introduce the notation R_t^* , for $t \geq 1$, defined as follows: if a red particle moved at time t , then $R_t^* = R_t$; otherwise $R_t^* = R_{t-1}$. We can think of this as first choosing whether a red particle moves, if so moving it and resolving any annihilation that occurs, then updating the number of red-occupied sites to R_t^* , then moving a blue particle if appropriate, and finally updating the number of red-occupied sites to R_t . Since a blue movement can only affect the red particles via a single annihilation, we necessarily have $R_t \in \{R_t^* - 1, R_t^*\}$. Note that if $R_t < R_t^*$ then a blue particle moved onto a lone red particle at time t .

We will show that with high probability R_t and B_t remain close to M_t until almost all particles have been destroyed, except for short excursions. To quantify the closeness, we introduce a function $f(m) = \max(6m/n, 1/w(n))$, where $w(n) = \omega(1)$, $w(n)^2 = o(\log n)$ and $\mathbb{E}(A)w(n)^2 = o(pn \log n)$, noting that such a function w exists by our assumption on $\mathbb{E}(A)$. We say that blue is *good* if $B_t(1 + f(M_t)) \geq M_t$, *bad* if $B_t(1 + 2f(M_t)) < M_t$, or *intermediate* otherwise, and likewise for red with respect to R_t .

We will show that blue quickly becomes good, and thereafter is unlikely to become bad within the timeframe of interest. On the other hand, the slower red particles are harder to control due to fluctuations in their distribution caused by annihilations with the faster blue particles. To analyse the red distribution, we introduce the red ‘level’, ℓ_t , which we define below. We also maintain a stack of ℓ_t ‘thresholds’, denoted S_1, \dots, S_{ℓ_t} if $\ell_t > 0$. The stack can only change in one of three ways: a new threshold is added when the level increases; the most recently added threshold is removed when the level decreases; or the entire stack is emptied when the level is reset to 0. The initialisation and progression of the level and stack is as follows.

- (a) If the initial configuration of red is good then $\ell_0 = 0$. Otherwise, $\ell_0 = n - R_0$ and $S_i = n - i + 1$ for $i = 1, \dots, \ell_0$.

- (b) If red becomes good at time t (i.e. it is good at time t but not time $t - 1$), then we set $\ell_t = 0$ (and empty the stack).
- (c) If $R_t < R_t^*$ and red is not good, then we set $\ell_t = \ell_{t-1} + 1$ and add R_t^* to the stack.
- (d) If $R_t = S_{\ell_{t-1}}$ then we set $\ell_t = \ell_{t-1} - 1$ and remove $S_{\ell_{t-1}}$ from the stack.
- (e) If $t \geq 1$ but none of the previous three cases apply, then $\ell_t = \ell_{t-1}$.

We observe the following simple properties.

Lemma 2. *For each time t , provided $\ell_t > 0$ we have $S_1 > \dots > S_{\ell_t} > R_t$. Consequently, the next time t' with $\ell_{t'} < \ell_t$ is also the next time such that either $R_{t'} = S_{\ell_t}$ or red is good at time t' . Furthermore, if red was good at time $t - 1$, then red must be good at time t' .*

Proof. We prove the first statement by induction; by (a), it holds for $t = 0$. Suppose it holds at time $t - 1$. If (b) or (d) apply, it clearly holds at time t . If (c) applies, then $R_t < R_t^* = S_{\ell_t}$ and blue must have moved at time t , meaning $R_t^* = R_{t-1}$. Thus, provided $\ell_t > 1$, we also have $S_{\ell_t} = R_{t-1} < S_{\ell_{t-1}}$. Finally, if (e) applies then $R_t \leq R_{t-1} + 1 \leq S_{\ell_{t-1}}$, and since (d) does not apply at least one inequality must be strict.

Note that the next time t' with $\ell_{t'} < \ell_t$ coincides with the time that S_{ℓ_t} is removed from the stack. This can only happen if either red is good at time t' or $R_{t'} = S_{\ell_t}$. If the former holds then $\ell_{t'} = 0$ by (b). If the latter holds then S_{ℓ_t} can no longer be on the stack, since this would contradict the first statement. Thus the second statement follows.

To prove the third statement, note that red being good at time $t - 1$ implies $\ell_{t-1} = 0$, and so we must have $\ell_t = 1$, $R_t^* = R_{t-1}$ and $R_t = R_t^* - 1$, giving $S_{\ell_t} = R_{t-1}$. It is easy to check that the function $x \mapsto \frac{x}{1+f(x)}$ is increasing, and so $R_{t'} = S_{\ell_t}$ would imply red is good at t' . \square

We define a stopping time τ as the first time that one of the following occurs:

- (i) $M_t = \lfloor \log^2 n \rfloor$;
- (ii) blue is bad but has previously been good;
- (iii) red is bad at level 0; or
- (iv) $t = 3n^2$.

Our estimate for the extinction time T will be $T \leq T_{\text{blue}} + T_{\text{red}} + T_{\text{late}} + T_{\text{ok}}$, where

- T_{blue} is the number of steps before blue first becomes good,
- T_{red} is the number of steps when red is above level 0,
- T_{late} is the number of steps after τ with red at level 0, and
- T_{ok} is the number of steps before τ where neither colour is bad.

To see that this estimate is valid, we first write $T_{\text{red}} = T_{\text{red}}^{(1)} + T_{\text{red}}^{(2)}$, where the term $T_{\text{red}}^{(1)}$ counts the number of steps before τ when red is above level 0, and $T_{\text{red}}^{(2)}$ counts the number of steps after τ when red is above level 0. Now consider any time step before τ . If blue has not yet become good, this time step is counted by T_{blue} . If red is above level 0, this time step is counted by $T_{\text{red}}^{(1)}$. If neither of these apply then blue has previously become good and red is at level 0. By definition of τ , (ii) and (iii) have not yet occurred, meaning that blue is still good and red is not bad, so this time step is counted by T_{ok} . Thus

$$\tau \leq T_{\text{blue}} + T_{\text{red}}^{(1)} + T_{\text{ok}}. \quad (1)$$

Finally, consider any time step after τ . This time step is counted by either $T_{\text{red}}^{(2)}$ or T_{late} , depending on the level. Therefore $T - \tau = T_{\text{red}}^{(2)} + T_{\text{late}}$, which together with (1) gives the required bound.

We start our analysis with the following simple tail bound on the runtime.

Lemma 3. *From any starting position with $m \leq n$ particles of each type, $\mathbb{E}(T) \leq 2mn$ and $\mathbb{P}(T > 3n^2) \leq e^{-\Theta(n)}$.*

Proof. At each step, conditional on the previous evolution, an annihilation occurs with probability at least $\frac{1}{2n}$, so the expected time between consecutive annihilations is at most $2n$. Summing gives the bound for $\mathbb{E}(T)$. The probability of needing more than $3n^2$ steps for $m \leq n$ annihilations is at most $\mathbb{P}(\text{Bin}(3n^2, \frac{1}{2n}) < n)$, so the probability bound holds by Chernoff. \square

Next we consider the main term in the estimate for T .

Lemma 4. *We have $\mathbb{E}(T_{\text{ok}}) \leq 2n \log n + 100w(n)^{-1}n \log n$.*

Proof. Each step when neither colour is bad results in an annihilation with probability at least

$$\frac{pB_t + (1-p)R_t}{2n} \geq \frac{M_t}{2n(1+2f(M_t))}.$$

Thus we bound the expected number of steps with neither colour bad before the next annihilation by $26n/M_t$ while $M_t \geq n/6w(n)$ (so $f(M_t) \leq 6$) or by $(1+2/w(n))2n/M_t$ while $M_t \leq n/6w(n)$ (so $f(M_t) = 1/w(n)$). Summing and using $\sum_{m=k}^n 1/m \leq 1 + \log n - \log k$ for all $n > k \geq 1$,

$$\begin{aligned} \mathbb{E}(T_{\text{ok}}) &\leq (1+2/w(n)) \sum_{m=1}^n 2n/m + \sum_{m=\lceil n/6w(n) \rceil}^n 24n/m \\ &\leq 2n \log n + 100w(n)^{-1}n \log n, \end{aligned}$$

since $w(n) \leq \sqrt{\log n}$ and n is sufficiently large. \square

The evolution of B_t may be bounded as follows, writing $q := 1 - p$. We have

$$\mathbb{P}(B_{t+1} = B_t + 1 \mid \mathcal{F}_t) \geq q \frac{M_t - B_t}{M_t} \cdot \frac{2n - B_t - R_t}{2n}, \quad (2)$$

since this occurs when a blue particle at a site occupied by at least one other blue particle, of which there are at least $M_t - B_t$, moves to an empty site. Similarly, we have

$$\mathbb{P}(B_{t+1} = B_t - 1 \mid \mathcal{F}_t) \leq q \frac{B_t}{M_t} \cdot \frac{B_t + R_t}{2n} + p \frac{B_t}{2n}, \quad (3)$$

since for this to happen either the only blue particle at some site must move to a non-empty site, or a red particle must move to a site occupied by exactly one blue particle. In all other cases, $B_{t+1} = B_t$.

Recall that we distinguish changes in R_t owing to red movements from those owing to blue movements, which may be dominant if blue is much faster than red. Not including the effect of blue movements, the corresponding transitions are

$$\mathbb{P}(R_{t+1}^* = R_t + 1 \mid \mathcal{F}_t) \geq p \frac{M_t - R_t}{M_t} \cdot \frac{2n - R_t - B_t}{2n} \quad (4)$$

and

$$\mathbb{P}(R_{t+1}^* = R_t - 1 \mid \mathcal{F}_t) \leq p \frac{R_t}{M_t} \cdot \frac{R_t + B_t}{2n}. \quad (5)$$

We will see that when a colour is not good then increments are significantly more likely, giving a strong self-correcting effect. To describe this, we need the following simple facts about biased random walks.

Lemma 5. *Let X_t be an integer-valued discrete-time stochastic process adapted to a filtration $(\mathcal{F}_t)_{t \geq 0}$ and stopped at some stopping time S . Suppose, for all $t \leq S$, we have*

- $X_{t+1} \in \{X_t - 1, X_t, X_t + 1\}$; and
- $\mathbb{P}(X_{t+1} = X_t - 1 \mid F_t) \leq \mathbb{P}(X_{t+1} = X_t + 1 \mid F_t)/2$.

Then, given $X_0 = k - 1$, and conditioned on F_0 , the probability of reaching 0 before k is at most 2^{1-k} . If additionally there is some $0 < \alpha \leq 2/3$ such that $\mathbb{P}(X_{t+1} = X_t + 1 \mid F_t) \geq \alpha$ for all t , then the expected time to reach k (or stop without having hit k) is at most $2\alpha^{-1}$.

Proof. For the first statement, it is sufficient to prove for every finite time t' that the probability of reaching 0 before reaching k and before time t' is at most 2^{1-k} . Note that 2^{k-X_t} is a supermartingale, and the first time T for which $X_T \in \{0, k\}$ or $T = S \wedge t'$ is a stopping time. Since $2^{k-X_{T \wedge t'}}$ is bounded, the Optional Stopping Theorem gives $\mathbb{E}(2^{k-X_T} \mid F_0) \leq 2^{k-X_0} = 2$. Since $2^{k-X_T} \geq 1$ we obtain $\mathbb{P}(X_T > 0 \mid F_0) + 2^k \mathbb{P}(X_T = 0 \mid F_0) \leq 2$, and therefore $\mathbb{P}(X_T = 0 \mid F_0) \leq 2^{1-k} - 2^{-k} \mathbb{P}(X_T > 0 \mid F_0) < 2^{1-k}$.

For the second statement, let Z_t^+ be the number of $s < t$ such that $X_{s+1} = X_s + 1$, and similarly let Z_t^- be the number of $s < t$ such that $X_{s+1} = X_s - 1$. Let T' be the first time that $X_{T'} = k$ or $T' = S$. Note that T' is a stopping time of finite expectation, and that $Z_t^+ - 2Z_t^-$ is a submartingale. Thus the Optional Stopping Theorem gives $\mathbb{E}(Z_{T'}^+ \mid F_0) \geq 2\mathbb{E}(Z_{T'}^- \mid F_0)$, and by definition of T' we have $Z_{T'}^+ - Z_{T'}^- \leq 1$. Thus $\mathbb{E}(Z_{T'}^+ \mid F_0) - \mathbb{E}(Z_{T'}^- \mid F_0)/2 \leq 1$, so $\mathbb{E}(Z_{T'}^+ \mid F_0) \leq 2$. Finally, $Z_t^+ - \alpha t$ is a submartingale, so $\mathbb{E}(T' \mid F_0) \leq \alpha^{-1} \mathbb{E}(Z_{T'}^+ \mid F_0) \leq 2\alpha^{-1}$. \square

We next establish the necessary bias in our walks.

Lemma 6. *Provided that the blue distribution is not good, we have*

$$\frac{\mathbb{P}(B_{t+1} = B_t + 1 \mid \mathcal{F}_t)}{\mathbb{P}(B_{t+1} = B_t - 1 \mid \mathcal{F}_t)} \geq 2, \quad (6)$$

and similarly, provided that the red distribution is not good, we have

$$\frac{\mathbb{P}(R_{t+1}^* = R_t + 1 \mid \mathcal{F}_t)}{\mathbb{P}(R_{t+1}^* = R_t - 1 \mid \mathcal{F}_t)} \geq 2. \quad (7)$$

Proof. We prove (6). Since the bound in (4) matches that of (2), after exchanging R_t and B_t , and p and q , but that of (5) is stronger than (3) (after the same exchanges), the proof of (7) only differs in having more room to spare. Since $R_t \leq M_t$ and $p \leq q$, we have, using (2),

$$\mathbb{P}(B_{t+1} = B_t + 1 \mid \mathcal{F}_t) \geq q \frac{M_t - B_t}{M_t} \cdot \frac{2n - B_t - M_t}{2n}$$

and, using (3),

$$\mathbb{P}(B_{t+1} = B_t - 1 \mid \mathcal{F}_t) \leq q \frac{B_t}{M_t} \cdot \frac{B_t + M_t}{2n} + p \frac{B_t}{2n} \leq q \frac{B_t}{M_t} \cdot \frac{B_t + 2M_t}{2n},$$

and so it is sufficient to verify that

$$(M_t - B_t)(2n - B_t - M_t) \geq 2B_t(B_t + 2M_t),$$

or equivalently that

$$2n(M_t - B_t) \geq M_t^2 + 4M_t B_t + B_t^2. \quad (8)$$

First suppose $M_t \geq 7B_t$. Then, since $n \geq M_t$, the LHS of (8) is at least $\frac{12}{7}M_t^2$ and the RHS is at most $\frac{78}{49}M_t^2$, so (8) holds. Now suppose $7B_t \geq M_t$. As blue is not good, we have $M_t \geq (1 + 6M_t/n)B_t$. Then the LHS of (8) is at least $12M_t B_t$ and the RHS is at most $11M_t B_t$, so again (8) holds. \square

We deduce a suitable bound on the number of steps before blue first becomes good.

Lemma 7. *We have $\mathbb{E}(T_{\text{blue}}) \leq 2n$.*

Proof. We show that with probability $1 - \exp(-\Theta(n))$, blue becomes good before $5n/9$ annihilations have occurred and within $\lceil 6nq^{-1}/7 \rceil < 2n$ time steps. The result then follows since we can bound the expected additional number of steps required when this fails by $2n^2 \exp(-\Theta(n)) = o(1)$, using Lemma 3.

Note that blue is good whenever $B_t \geq n/7$. Provided $B_t < n/7$ and $M_t > 4n/9$, we can calculate using (2) and $R_t \leq M_t$ that

$$\mathbb{P}(B_{t+1} = B_t + 1 \mid \mathcal{F}_t) \geq q \min_{x \in [4/9, 1]} \left(\frac{x - 1/7}{x} \cdot \frac{13/7 - x}{2} \right) = 18q/49.$$

Similarly, since $p \leq q$, we have from (3) that

$$\mathbb{P}(B_{t+1} = B_t - 1 \mid \mathcal{F}_t) \leq q \max_{x \in [4/9, 1]} \left(\frac{1/7 + x}{14x} + \frac{1}{14} \right) = 65q/392.$$

Thirdly, the probability of an annihilation under the same assumptions is at most

$$q \frac{M_t}{2n} + p \frac{B_t}{2n} \leq 4q/7.$$

Now run the process until one of the following occurs: $M_t \leq 4n/9$, $B_t \geq n/7$ or $\lceil 6nq^{-1}/7 \rceil$ time steps have occurred. Note that the number of annihilations up to this point is dominated by $\text{Bin}(\lceil 6nq^{-1}/7 \rceil, 4q/7)$, which by a standard Chernoff bound has exponentially small probability of exceeding $5n/9$. Similarly, the number of times B_t decreases up to this point has exponentially small probability of exceeding $n/7$. Since $\text{Bin}(\lceil 6nq^{-1}/7 \rceil, 18q/49) < 2n/7$ with exponentially small probability, B_t reaches $n/7$ before time $\lceil 6nq^{-1}/7 \rceil$ with suitably high probability. \square

Recall that the stopping time τ is the first time when one of four possible conditions holds. The ‘success’ event is that (i) occurs at time τ , i.e. $M_\tau = \lfloor \log^2 n \rfloor$, and the ‘failure’ event, denoted F , is that one of (ii), (iii) or (iv) occurs before (i). We now bound the probability of failure.

Lemma 8. *We have $\mathbb{P}(F) \leq 3n^{-3}$.*

Proof. Note that $\mathbb{P}(F) \leq \mathbb{P}(F_{\text{blue}}) + \mathbb{P}(F_{\text{red}}) + \mathbb{P}(F_{\text{time}})$, where F_{blue} is the event that (ii) occurs (that is, blue is bad but has previously been good) at time τ , and likewise F_{red} that (iii) does (that is, red is bad at level 0) and F_{time} that (iv) does (that is, $\tau = 3n^2$).

Suppose that F_{blue} or F_{red} occurs. We say that the *genesis* of such an event is the last time τ^* prior to τ that the relevant colour was good. For F_{blue} , the existence of such a time is immediate from (ii); for F_{red} note that (a) ensures that red will be good when the level first reaches 0.

We first bound, for each fixed time $t^* < 3n^2$, the probability that blue is good at t^* , not good at $t^* + 1$, and thereafter becomes bad before becoming good. Denote this event by $E_{t^*, \text{blue}}$, and let $E'_{t^*, \text{blue}}$ be the event that blue is good at t^* but not good at $t^* + 1$. Note that, if F_{blue} occurs with genesis τ^* , then $E_{\tau^*, \text{blue}}$ occurs (and $\tau^* < \tau \leq 3n^2$), and so taking a union bound across all $t^* < 3n^2$ will bound $\mathbb{P}(F_{\text{blue}})$.

Clearly $\mathbb{P}(E_{t^*, \text{blue}}) \leq \mathbb{P}(E_{t^*, \text{blue}} \mid E'_{t^*, \text{blue}})$, so we bound the latter; note that $E'_{t^*, \text{blue}}$ is \mathcal{F}_{t^*+1} -measurable. Let $X_i = B_{t^*+1+i}$ for $i = 0, \dots, S$, where S is the smallest value such that blue is either good or bad at $t^* + 1 + S$. We clearly have $X_i \in \{X_i - 1, X_i, X_i + 1\}$ for all $0 \leq i < S$.

Furthermore, since blue is not good for all $t^* + 1 \leq t < S$, (6) of Lemma 6 applies. Thus both conditions of Lemma 5 are met with $F_t = \mathcal{F}_{t+t^*+1}$.

Note that $B_{t^*} \geq \frac{M_{t^*}}{1+f(M_{t^*})}$ but $X_0 = B_{t^*+1} < \frac{M_{t^*+1}}{1+f(M_{t^*+1})}$, which together with $B_{t^*+1} \geq B_{t^*} - 1$ and the fact that $\frac{M_t}{1+f(M_t)}$ is non-increasing implies $X_0 = \lceil \frac{M_{t^*}}{1+f(M_{t^*})} \rceil - 1$. In order for $E_{t^*,\text{blue}}$ to occur, X_i must reach $\lfloor \frac{M_{t^*}}{1+2f(M_{t^*})} \rfloor$ before $\lceil \frac{M_{t^*}}{1+f(M_{t^*})} \rceil$. Lemma 5 gives that this occurs with probability at most 2^{1-z} , where $z = \lceil \frac{M_{t^*}}{1+f(M_{t^*})} \rceil - \lfloor \frac{M_{t^*}}{1+2f(M_{t^*})} \rfloor$.

Next we consider F_{red} . For each fixed time $t^* < 3n^2$, denote by $E_{t^*,\text{red}}$ the event that red is good at t^* , red is not good at time $t^* + 1$, and that red becomes bad at level 0 before it becomes good, and let $E'_{t^*,\text{red}}$ be the event that red is good at t^* but not at time $t^* + 1$, and $\ell_{t^*+1} = 0$. Note that if $\ell_{t^*+1} > 0$ then, by Lemma 2, red is good at the next time t' that $\ell_{t'} = 0$, and so $E_{t^*,\text{red}}$ does not occur. Again, if F_{red} occurs with genesis τ^* , then $\tau^* < 3n^2$ and $E_{\tau^*,\text{red}}$ and $E'_{\tau^*,\text{red}}$ both occur, so $\mathbb{P}(F_{\text{red}}) \leq \sum_{t^* < 3n^2} \mathbb{P}(E_{t^*,\text{red}} \mid E'_{t^*,\text{red}})$.

In order to bound $\mathbb{P}(E_{t^*,\text{red}} \mid E'_{t^*,\text{red}})$, condition on the latter event and let $t^* + 1 = c_0 < c_1 < \dots$ be the times after t^* such that $\ell_{c_i} = 0$, let S be the smallest value such that red is either good or bad at c_S and let $X_i = R_{c_i}$ for $i = 0, 1, \dots, S$. In other words, we remove any excursions above level 0, each of which returns R_t to the value before the excursion. This ensures $X_{j+1} = R_{c_{j+1}}^*$ for each $0 \leq j < S$. In particular, $X_{j+1} \in \{R_{c_j} - 1, R_{c_j}, R_{c_j} + 1\}$, and, since red is not good, (7) of Lemma 6 applies. Thus again both conditions of Lemma 5 are met with $F_t = \mathcal{F}_{c_t}$. By the same argument as before, we have $X_0 = \lceil \frac{M_{t^*}}{1+f(M_{t^*})} \rceil - 1$, and in order for $E_{t^*,\text{red}}$ to occur, X_i must reach $\lfloor \frac{M_{t^*}}{1+2f(M_{t^*})} \rfloor$ before $\lceil \frac{M_{t^*}}{1+f(M_{t^*})} \rceil$, which by Lemma 5 has probability at most 2^{1-z} .

Note that $\frac{x}{(1+x)(1+2x)}$ is increasing for $0 \leq x \leq 2^{-1/2}$ and decreasing for $x \geq 2^{-1/2}$. Also, we have $1/w(n) \leq f(M_t) \leq 6$ and $M_{t^*} \geq \log^2 n$. Consequently, we have

$$\begin{aligned} \left\lceil \frac{M_{t^*}}{1+f(M_{t^*})} \right\rceil - \left\lfloor \frac{M_{t^*}}{1+2f(M_{t^*})} \right\rfloor &\geq \frac{M_{t^*} f(M_{t^*})}{(1+f(M_{t^*}))(1+2f(M_{t^*}))} \\ &\geq \min(M_{t^*}/(4w(n)), 6M_{t^*}/91) \geq 10 \log n. \end{aligned}$$

Thus $2^{1-z} \leq 2^{1-10 \log n} \leq n^{-6}$. By a union bound over all $t^* < 3n^2$ and choices of blue or red, we have $\mathbb{P}(F_{\text{blue}}) + \mathbb{P}(F_{\text{red}}) \leq 2n^{-3}$. Finally, the probability of failure due to taking too many steps is bounded by $\mathbb{P}(T \geq 3n^2) \leq n^{-3}$ by Lemma 3. \square

Next we control the time spent with red not at level zero.

Lemma 9. *We have $\mathbb{E}(T_{\text{red}}) \leq 12nw(n)^2 + 6p^{-1}\mathbb{E}(A)w(n)^2$.*

Proof. We break the calculation into the following two claims, where the first bounds the contribution from each increase in level, and the second the expected number of increases in level. Here we consider that the level increases to i at time t^* if either $t^* > 0$, $\ell_{t^*-1} = i - 1$ and $\ell_{t^*} = i$, or $t^* = 0$ and $\ell_0 \geq i \geq 1$. Thus we consider that the process starts with ℓ_0 simultaneous increases.

Claim 10. *Suppose the level increases to i at time t^* . Let t' be the next time with $\ell_{t'} < i$. Let $C = \{t : t^* \leq t \leq t' \text{ and } \ell_t = i\}$. Then $\mathbb{E}(|C| \mid \mathcal{F}_{t^*}) \leq 5p^{-1}(1 + w(n))$.*

Proof. We show the slightly stronger bound that $\mathbb{E}(|C| \mid \mathcal{F}_{t^*}, C \neq \emptyset) \leq 5p^{-1}(1 + w(n))$, which only differs if $t^* = 0$, since otherwise $t^* \in C$. Write $C = \{c_1, \dots, c_s\}$, where $c_1 < c_2 < \dots < c_s < t'$. As in the proof of Lemma 8, we consider the sequence R_{c_1}, \dots, R_{c_s} , i.e. we remove any excursions above level i , each of which returns R_t to the value before the excursion. This ensures $R_{c_{j+1}}^* = R_{c_{j+1}}$ for each $j < s$.

Note that, using Lemma 2, the level will fall below i if R_t exceeds R_{c_1} : if $t^* > 0$ then $c_1 = t^*$ and this follows by (c) and (d), whereas if $t^* = 0$ then it follows from (a) and (d). As red is not good at c_j for any $j \leq s$, we have

$$\begin{aligned} \mathbb{P}(R_{c_j+1}^* = R_{c_j} + 1 \mid \mathcal{F}_{c_j}) &\geq p \frac{M_{c_j} - R_{c_j}}{M_{c_j}} \cdot \frac{2n - R_{c_j} - B_{c_j}}{2n} \\ &\geq p \frac{f(M_{c_j})}{1 + f(M_{c_j})} \cdot \frac{3}{7} \geq \frac{3p}{7} \cdot \frac{1}{1 + w(n)}. \end{aligned}$$

By Lemma 6, the conditions of Lemma 5 are satisfied with $\alpha = \frac{3p}{7} \cdot \frac{1}{1+w(n)}$ and thus the expected time for this sequence to exceed its initial value is at most $14p^{-1}(1 + w(n))/3$. \square

Claim 11. *The expected number of increases in level is at most $(2pn + \mathbb{E}(A))w(n)$.*

Proof. Recall that $A = n - R_0$ and $q = 1 - p$. Let L be the number of remaining red particles when the level first increases. It suffices to show $\mathbb{E}(L \mid A = a) \leq (2pn + a)w(n)$ for each a such that $\mathbb{P}(A = a) > 0$. This is easy to see if red is not initially good, as then $a \geq 6n/7$ and $L = n < aw(n)$.

Now suppose red is initially good. We say that a ‘bad move’ by red is a movement of a red particle onto another site occupied by red. We write Ξ_t for the number of bad moves by time t . As $M_t - R_t$ can only increase via bad moves, we have $M_t - R_t \leq a + \Xi_t \leq a + \Xi_T$. When the level increases red is not good, so $M_t - R_t \geq R_t f(R_t) \geq R_t/w(n)$. We deduce $L \leq (a + \Xi_T)w(n)$.

It remains to bound Ξ_T . We note that a bad move at time t has probability at most $p \frac{R_t}{2n}$, whereas the probability of an annihilation is at least $q \frac{R_t}{2n}$, since this is the probability of a blue particle moving onto a site occupied by red particles. Thus $pC_t - q\Xi_t$ is a submartingale, where C_t is the number of annihilations by time t . At the stopping time T , since $C_T = n$, we have $\mathbb{E}(\Xi_T) \leq np/q \leq 2pn$, as required. \square

Now we deduce the lemma from the following calculation combining these two claims:

$$\begin{aligned} \mathbb{E}(T_{\text{red}}) &\leq (2pn + \mathbb{E}(A))w(n) \cdot 5p^{-1}(1 + w(n)) \\ &= (10n + 5\mathbb{E}(A)p^{-1})(w(n) + w(n)^2) \\ &\leq 12nw(n)^2 + 6p^{-1}\mathbb{E}(A)w(n)^2. \end{aligned} \quad \square$$

It remains to bound the time for completion once only $\log^2 n$ particles remain.

Lemma 12. *We have $\mathbb{E}(T_{\text{late}}) \leq 26n \log \log n$.*

Proof. Recalling that F is the failure event, we write $\mathbb{E}(T_{\text{late}}) = \mathbb{P}(F)\mathbb{E}(T_{\text{late}} \mid F) + \mathbb{E}(T_{\text{late}}1_{\neg F})$. By Lemma 8, we have $\mathbb{P}(F) \leq 3n^{-3}$. By Lemma 3, we have $\mathbb{E}(T_{\text{late}} \mid \mathcal{F}_\tau) \leq 2n^2$. Since F is \mathcal{F}_τ -measurable, also $\mathbb{E}(T_{\text{late}} \mid F) \leq 2n^2$ and so the first term is at most $6/n$. For the second, we assume F does not occur, and note that T_{late} only counts steps t after τ with red at level 0, so red must be good at such t . In particular, $R_t \geq M_t/2$, so the probability of an annihilation is at least $p \frac{M_t}{4n} \geq \frac{M_t}{8n}$. Thus this contribution has expected size at most $\sum_{m=1}^{\lfloor \log^2 n \rfloor} 8n/m \leq 16n \log \log n + 8n$. \square

Combining the bounds from Lemmas 4, 7, 9 and 12, using $\mathbb{E}(A)w(n)^2 = o(pn \log n)$ gives

$$\begin{aligned} \mathbb{E}(T) &\leq 2n \log n + 100w(n)^{-1}n \log n + 2n + 12nw(n)^2 + 6p^{-1}\mathbb{E}(A)w(n)^2 + 26n \log \log n \\ &= (2 + o(1))n \log n, \quad \text{as required to prove Theorem 1.} \end{aligned}$$

As a variant form of Theorem 1, one can obtain an explicit estimate for the second order term under the stronger assumption $\mathbb{E}(A) \leq pn$. Indeed, then applying the argument with $w(n) = \log^{1/3} n$ we obtain $\mathbb{E}(T) \leq 2n \log n + O(n \log^{2/3} n)$. This form of the second order term seems to be an artefact of the proof; we do not think it is optimal.

3 Concluding remarks

In a companion paper [13] we consider a more general class of interacting particle systems on expanders, for which we obtain similar $\Theta(n \log n)$ bounds on the equilibrium time. Our methods there are entirely different, and while they apply in much greater generality, they only reveal the order of magnitude, not the asymptotics. It may be interesting to extend the methods of this paper to obtain asymptotic results for other graphs with simple structure, such as the complete bipartite graph $K_{n,n}$. This question conceals another, more basic question: does $K_{n,n}$ have a well-defined asymptotic extinction time from any starting distribution, or do some initial distributions reach extinction much faster than others?

For $K_{n,n}$ we believe that the starting distribution does not make much difference. We give some evidence for this by sketching a proof in the case of stationary reds. In this case the extinction time does not depend on the order in which blue particles are chosen to move (the ‘abelian sandpile’ property, see [13]). Thus we may consider blues in turn, moving each one until it hits a red. Each time taken is a random variable of the following form: count the number of independent trials required for success, where odd trials have some success probability p_1 and even trials have some success probability p_2 . While p_1 and p_2 depend on the arrangement of remaining reds, $p_1 + p_2 = m/n$ does not, where m is the number of reds remaining. Now this is $2X - Y$ where X is a geometric random variable with success probability $p_1 + p_2 - p_1 p_2$, i.e. the probability that one of any given pair of trials succeeds, and Y is an independent Bernoulli random variable with probability $p_1 / (p_1 + p_2 - p_1 p_2)$. In particular, $\mathbb{E}(2X - Y)$ is minimised when $p_2 = 0$, at $2n/m - 1$. Assuming $m < n / \log n$ we have $p_1 p_2 < (p_1 + p_2)^2 / 4 < (p_1 + p_2) / (4 \log n)$, so $\mathbb{E}(2X - Y) \leq (2 + 1 / \log n) m / n$, and the sum of all terms with $m < n / \log n$ is at most $2n \log n - 2n \log \log n + O(n)$. The sum of all other terms is increased by at most a constant factor, and so has sum $O(n \log \log n)$. Thus for stationary reds on $K_{n,n}$ we have $\mathbb{E}(T) = (2 + o(1)) n \log n$ for any starting distribution.

Acknowledgements

Both authors were supported by ERC Advanced Grant 883810. We are grateful to the anonymous referees for their help in improving the exposition.

References

- [1] R. Arratia, Limiting point processes for rescalings of coalescing and annihilating random walks on Z^d . *Ann. Probab.* **9**:6 (1981), 909–936.
- [2] R. Arratia, Site recurrence for annihilating random walks on Z^d . *Ann. Probab.* **11**:3 (1983), 706–713.
- [3] M. Bramson and J. L. Lebowitz, Asymptotic behavior of densities in diffusion dominated two-particle reactions. *Phys. A* **168**:1 (1990), 88–94.
- [4] M. Bramson and J. L. Lebowitz, Spatial structure in diffusion-limited two-particle reactions. *J. Statist. Phys.* **65**:5–6 (1991), 941–951.

- [5] M. Bramson and J. L. Lebowitz, Spatial structure in low dimensions for diffusion limited two-particle reactions. *Ann. Appl. Probab.* **11:1** (2001), 121–181.
- [6] F. Capannoli and F. den Hollander, Interacting particle systems on random graphs. *Ensaaios Mat.* **40** (2024), 117–176.
- [7] I. Cristali, Y. Jiang, M. Junge, R. Kassem, D. Sivakoff, and G. York, Two-type annihilating systems on the complete and star graph. *Stochastic Process. Appl.* **139** (2021), 321–342.
- [8] Y. Elskens and H. L. Frisch, Annihilation kinetics in the one-dimensional ideal gas. *Phys. Rev. A*, **31** (1985), 3812–3816.
- [9] P. Erdős and P. Ney, Some problems on random intervals and annihilating particles, *Ann. Probab.* **2:5** (1974), 828–839.
- [10] C. Cooper, R. Elsässer, H. Ono, and T. Radzik, Coalescing random walks and voting on connected graphs. *SIAM J. Discrete Math.* **27:4** (2013), 1748–1758.
- [11] C. Cooper, A. Frieze and T. Radzik, Multiple random walks in random regular graphs. *SIAM J. Discrete Math.* **23** (2009), 1738–1761.
- [12] P. Donnelly and D. Welsh, (1983). Finite particle systems and infection models. *Math. Proc. Camb. Phil. Soc.* **94** (1983), 167–182.
- [13] J. Haslegrave and P. Keevash, Dissipative particle systems on expanders. *Probab. Theory Relat. Fields* (2025).
- [14] J. Haslegrave, V. Sidoravicius and L. Tournier, Three-speed ballistic annihilation: phase transition and universality. *Selecta Math. (N.S.)* **27:5** (2021), paper 84, 38pp.
- [15] P. L. Krapivsky, S. Redner and F. Leyvraz, Ballistic annihilation kinetics: The case of discrete velocity distributions. *Phys. Rev. E*, **51** (1995), 3977–3987.
- [16] T. M. Liggett, *Interacting particle systems*, Springer-Verlag, New York, 1985.
- [17] J.-C. Lootgieter, Problèmes de récurrence concernant des mouvements aléatoires de particules sur Z avec destruction. *Ann. Inst. H. Poincaré Sect. B (N.S.)* **13:2** (1977), 127–139.