

Topology-dependent density optima for efficient simultaneous network exploration

Daniel B. Wilson,¹ Ruth E. Baker,¹ and Francis G. Woodhouse²

¹*Wolfson Centre for Mathematical Biology, Mathematical Institute, University of Oxford, Radcliffe Observatory Quarter, Oxford OX2 6GG, United Kingdom.*

²*Department of Applied Mathematics and Theoretical Physics, Centre for Mathematical Sciences, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, United Kingdom.*

(Dated: April 27, 2018)

A random search process in a networked environment is governed by the time it takes to visit every node, termed the cover time. Often, a networked process does not proceed in isolation but competes with many instances of itself within the same environment. A key unanswered question is how to optimise this process: how many concurrent searchers can a topology support before the benefits of parallelism are outweighed by competition for space? Here, we introduce the searcher-averaged parallel cover time (APCT) to quantify these economies of scale. We show that the APCT of the networked symmetric exclusion process is optimised at a searcher density that is well predicted by the spectral gap. Furthermore, we find that non-equilibrium processes, realised through the addition of bias, can support significantly increased density optima. Our results suggest novel hybrid strategies of serial and parallel search for efficient information gathering in social interaction and biological transport networks.

I. INTRODUCTION

From animals foraging to T-cells hunting pathogens to proteins examining DNA, nature relies on carefully optimised random searches at many scales [1–10]. This concept is not limited to biology: robotic self-assembly [11], traffic flow management [12–14] and computer resource allocation [15] hinge on optimising decentralised exploration. In these distributed processes, the searcher, be it a protein, an animal or a network token, must often visit not just one site but many locations connected in a network [16]. The critical measure of efficiency is then the time to visit every node of the network, called the cover time [17]. Single-searcher cover times are known on simple networks such as linear chains, rings and other regular lattices [18–22], and significant progress has been made on establishing transport statistics for more general networks [16, 23–28], providing a means to design topologies that can be searched efficiently. However, when multiple parallel searchers compete for space or resources, as often occurs in the examples above, how to design optimal search strategies remains a significant open question. Our central result is that the optimal density of searchers depends heavily, yet predictably, on network topology, implying that search strategies can be made efficient by careful optimisation of topology and searcher quantity.

The exclusion process represents the most fundamental model of competition for space [29–31]. It has been key to understanding such diverse phenomena as cell migration [32], molecular traffic [33, 34], surface roughening [35] and queueing [36]. Capitalising on this breadth, we use the exclusion process to model parallel searching of a network and introduce the *searcher-averaged parallel cover time* (APCT), the average per-searcher time for all searchers to visit all nodes within a network. Strategy optimisation then demands an understanding of how both network topology and searcher density impact the

APCT. Consider, for example, a scenario with as many searchers as nodes. Placing all searchers on the network simultaneously results in an infinite APCT, while a simple ‘serial’ strategy where a single searcher is placed on the network, removed once it has visited every node and replaced with a new searcher, one at a time, is almost always inefficient. It is therefore critical to determine the optimal, or most efficient, density of parallel searchers minimising the APCT (Fig. 1(a)). Through analytic and numerical results, we find that this optimal density is heavily dependent on network topology. We demonstrate that the spectral gap, which quantifies the convergence rate of a single-searcher random walk, is a strong predictor of a network’s density optimum and outperforms simpler degree-based network statistics, as measured by mutual information. We provide strategies for optimal deployment of hybrid series–parallel searches allowing for construction of efficiently-explored networks. We broaden to non-equilibrium processes by generalising to flux-conserving asymmetric exclusion processes, finding a remarkable non-monotonic relationship between density optima and the spectral gap. Our work provides an accessible route into the design of optimal search strategies in complex environments involving equilibrium and non-equilibrium processes.

II. AVERAGE PARALLEL COVER TIME AND OPTIMAL DENSITY

We employ a symmetric exclusion process of M parallel searchers on a network with N vertices (where $M < N$). First, an initial configuration of searchers is generated uniformly at random with each searcher occupying its own node. The searchers then perform mutually-excluding continuous-time random walks (CTRWs) with i.i.d. exponential waiting times of mean $\tau = 1$. When

a searcher attempts to move, an adjacent node is picked uniformly at random; if the node is vacant the searcher moves there, and if not the move is aborted. Let T_r^i be the time for searcher i to first visit r distinct nodes. (In practice we evaluate T_r^i by counting the number of attempted jumps made by all parallel searchers until walker i visits the r -th distinct site and rescaling by the average waiting time $1/M$. This is equivalent to describing the CTRW by a discrete-time random walk with time step $1/M$ which is sufficient when seeking only first moments as we do here.) We define the parallel cover time $\mathcal{C}_M = \max_{1 \leq i \leq M} \{T_N^i\}$ as the time for all M searchers to each visit all N nodes, and the expectation $\langle \mathcal{C}_M \rangle$ to be taken simultaneously over the space of all initial configurations and random walk instances [37]. The APCT, $\mathcal{S}_M = \langle \mathcal{C}_M \rangle / M$, then quantifies the economy of scale in parallel searching: if $\langle \mathcal{C}_M \rangle < M\mathcal{S}_1$ for some $M > 1$ there is an efficiency gain in parallel searching beyond the simple serial search strategy. For a given network, let M^* be the optimal number of searchers, that is, the M for which \mathcal{S}_M is minimised, and let $\rho^* = M^*/N$ be the equivalent optimal density. For numerical methods, see Appendix A.

A. Efficient search for ring-like topologies

Parallel search of a ring lattice is optimal at a strikingly low density ($\rho^* = 0.05$ when $N = 100$; Fig. 1(a)), and it becomes increasingly inefficient as M increases. This is due to the strong confinement effects of single-file diffusion [38–43]. That said, parallel searching is still more efficient than simple serial searching for a range of ρ (Fig. 1(a), dashed line). Remarkably, however, a parallel search on a ring lattice can be made significantly more efficient by introducing only a small number of additional edges. To demonstrate, we consider a form of the Newman–Watts ensemble [44] that interpolates between the two extremal topologies of the ring lattice and the complete network. Starting from a ring lattice, we add a fixed number of random additional edges, or shortcuts (Fig. 1(b), inset). The optimal density rapidly approaches that of the complete network, $\rho^* = 0.47$, even with only 3–4% of the possible shortcuts added on $N = 100$ nodes (Fig. 1(b)).

Parallel search efficiency on the complete network can be evaluated exactly. In Appendix B, we derive the APCT

$$\mathcal{S}_M = \frac{(N-1)^2}{M(N-M)} h(M(N-1)), \quad (1)$$

where $h(i)$ is the i -th harmonic number. The optimal density of parallel searchers, ρ^* , for a given N can then be calculated from $M^* = \operatorname{argmin}_M \mathcal{S}_M$, which gives $\rho^* \rightarrow 0.5$ in the limit $N \rightarrow \infty$ (Appendix B). Equation (1) can be written as $\mathcal{S}_M = \phi \tilde{\mathcal{S}}_M$ where $\tilde{\mathcal{S}}_M$ is the non-interacting APCT (Appendix B) and $\phi = (1 - N^{-1})/(1 - \rho)$ is a mean-field correction. This accounts




FIG. 1. (a) APCT on the ring lattice for interacting (\times) and non-interacting ($+$) searchers as a function of searcher density $\rho = M/N$, where $N = 100$. The optimal density ρ^* (marked) can be seen as the minimum of the APCT; dashed horizontal line denotes the area below which parallel search is faster than serial search. Each APCT was calculated from 10^5 random walk instances. (b) Mean optimal density of Newman–Watts networks with $N = 100$ as a function of the number of shortcuts (example inset). Coloured band indicates ± 1 s.d. Optimal densities were calculated (Appendix A) for 1000 Newman–Watts realisations for a range of added edges.

for the average slow-down due to aborted moves by assuming two-site occupancy probabilities to be products of single-site densities (Appendix C). While it is exact for the complete network by spatial homogeneity, this need not hold in general.

For general networks, naïve mean-field approximations (MFAs) [29–31] can lead to drastically inaccurate estimation of the APCT and optimal density. Taking the near-Gumbel-distributed single searcher cover times on random networks [16] and attempting to incorporate exclusion through rescaling by ϕ fails for many networks: under this MFA the predicted optimal density as N becomes large approaches $\rho^* = 0.5$ (Appendix C), but for networks of low average degree, and for the ring lattice (Fig. 1(a)) in particular, this is far from the true ρ^* evaluated numerically (Fig. 2). Insight into this inaccuracy can be drawn from asymptotic estimates of the APCT on the ring lattice in the high density regime. Consider $M = N - k$ searchers with $k \ll N$ vacancies. Inspired by particle-hole duality [38] we can follow the vacancies instead of the searchers. The parallel cover time is then the time taken for the net displacement of the vacancies to first reach $(N-1)(N-k)$ (Appendix B). For large N the vacancies are approximately non-interacting, meaning their net displacement is approximately that of a single vacancy moving k times faster, that is, $\tau_{\text{vac}} = 1/k$. Standard results [20] then imply

$$\mathcal{S}_{N-k} \sim \frac{(N-1)[(N-1)(N-k)+1]}{2k}. \quad (2)$$

This estimate does extremely well in predicting the APCT (Appendix B, Fig. 5). It is exact for $k = 1$, and only noticeably deviates when $\rho \lesssim 0.85$ for $N = 100$. Equation (2) reveals that the APCT at high density is $\mathcal{O}(N^3)$, while the MFA suggests an APCT of $\phi \tilde{\mathcal{S}}_N \sim \mathcal{O}(N^2)$, an order of magnitude difference, highlighting the failure of the MFA to capture spatial correlations.




FIG. 2. Spectral gap predicts optimal parallel search density. Optimal densities were calculated for 1500 random networks on $N = 100$ nodes with average degree between 2 and 5 (Appendix D). Each point represents one network with average degree indicated by colour, with three topologically contrasting networks highlighted. Solid grey curve denotes the mean MFA-predicted optimal density as a function of \mathcal{G} from numerically-determined single-searcher cover times. Dashed red curve shows best-fit expansion $\rho^*(\mathcal{G}; 0.7556, 0.2933)$. Two real-world transport networks are indicated by stars: the London Underground (lower) and American airports (upper).

B. General topologies and the spectral gap

Can we identify a topological heuristic to predict density optima for general networks? Given the adjacency matrix \mathbf{A} and the diagonal matrix \mathbf{D} whose entries are the node degrees, define the random walk transition matrix $\mathbf{P} = \mathbf{D}^{-1}\mathbf{A}$. The eigenvalues $\{\lambda_i\}$ of \mathbf{P} determine how fast the probability distribution of a single-searcher random walk converges to its equilibrium (that is, mixes) [45]. The largest eigenvalue is necessarily $\lambda_1 = 1$, and we define the spectral gap, \mathcal{G} , as the difference in magnitude between the first and second largest eigenvalues $\mathcal{G} = 1 - \max_{2 \leq i \leq N} \{|\lambda_i|\}$. The closer \mathcal{G} is to zero, the slower a random walk converges [45]. The spectral gap is then a natural candidate for predicting density optima as the features of networks that slow down convergence, such as bottlenecks (identified through nodes of high betweenness centrality [46]), also significantly increase the parallel cover time. Furthermore, single-searcher cover times can be related to the eigenvalues of the combinatorial Laplacian $\mathbf{L} = \mathbf{D} - \mathbf{A}$ [23], which in turn relate to those of $\mathbf{P} = \mathbf{I} - \mathbf{D}^{-1}\mathbf{L}$ via the normalised Laplacian [47].

There is a tight relationship between the spectral gap and the optimal density of searchers for general networks (Fig. 2), quantified through high mutual information (Appendix E). Sampling over random networks of minimum degree two and average degree between two and five (Appendix D) reveals that topology has a huge impact on optimal parallel search strategies: for a small spectral gap

(Fig. 2, left-most inset) networks have low average degree and a high concentration of nodes of degree two, resulting in linear chains along which the single-file diffusion of searchers significantly increases the cover time [38–43]. Indeed, for a linear chain the parallel cover time is finite only for a single searcher; for two or more searchers the left- (right-) most searcher will never reach the right- (left-) most site due to the effect of single-file diffusion. The naïve MFA fails to capture this relationship (Fig. 2, grey curve), demonstrating the importance of long-lived occupancy correlations. The MFA only becomes valid for networks of average degree nearing five, where the density optima approach that of the complete network ($\rho^* = 0.47$ for $N = 100$). Typically, these networks have a small fraction of degree two nodes and are more highly connected (Fig. 2, right-most inset). A power series $\rho^*(\mathcal{G}; a, b) = \mathcal{G}^b [1/2 - (1/2 - a)(1 - \mathcal{G})]$, constructed such that $\rho^*(1) = 1/2$ to match the $N \rightarrow \infty$ limit of the complete graph, matches the data well (Fig. 2, red curve) and provides an easily-computed predictive approximation.

Beyond random networks, we find that two real-world transport networks further validate this relationship in practice (Fig. 2). Firstly we consider the London Underground network (from Ref. [48], which includes the Docklands Light Railway, with one spurious disconnected node deleted from the dataset). The average degree of this network is 2.31 where the number of nodes is 306, the spectral gap of the symmetric random walk of a single searcher is 0.003 and the estimated optimal density is 0.0915 (Fig. 2). Our second real world network is the top 500 busiest US airports, where the nodes represent each airport and an edge connects two nodes if there is at least one direct flight between the two airports [49]. The average degree of the network is 11.9 where the number of nodes is 500. The spectral gap of the symmetric random walk of a single searcher is 0.0237 and the numerically estimated optimal density is 0.26. Despite the network’s high average degree the optimal density is still prohibitively low, which is excellently predicted via its spectral gap (Fig. 2). Simpler statistics based on the degree distribution fail in such cases and have lower mutual information over our graph ensemble (Appendix E).

III. OPTIMISING PARALLEL SEARCH STRATEGIES

Our results enable the design of optimal search strategies for arbitrary numbers of searchers. Suppose M_{tot} searchers, potentially more than N , need to each cover a given network. A hybrid series-parallel strategy comprises sequentially introducing batches of searchers, with each batch covering the network in parallel to completion before being removed and replaced by the next batch. Let k_i be the number of times i parallel searchers are introduced sequentially; for example, $k_1 = 3$ and $k_2 = 4$ denotes performing three successive searches




FIG. 3. (a) Optimal search strategy as M_{tot} is increased for a random network of $N = 20$ nodes (inset) with $\rho^* = 0.35$. Values at fixed M_{tot} give the non-zero components of the strategy vector \mathbf{k}_{opt} . (b) Mean number of searchers per instance ($\langle \mathbf{k}_{\text{opt}} \rangle$) for the optimal search strategies of the complete network, the random network in (a) and the ring lattice on $N = 20$ nodes, showing the approach of $\langle \mathbf{k}_{\text{opt}} \rangle$ towards the optimal APCT densities $\rho^* = 0.45, 0.35$ and 0.15 respectively for large M_{tot} . APCTs are averages over 10^5 random walk instances.

with one searcher and four successive searches with two searchers. A hybrid strategy is then defined by the vector $\mathbf{k} = (k_1, \dots, k_{N-1})$. The optimal strategy, \mathbf{k}_{opt} , minimises the overall cover time $\mathcal{T} = \sum_{i=1}^{N-1} k_i \langle \mathcal{C}_i \rangle$ subject to the constraint $\sum_{i=1}^{N-1} i k_i = M_{\text{tot}}$. As M_{tot} increases the solution of this integer programming problem typically yields a strategy comprising a mixture of searcher numbers near the optimal density (Fig. 3).

To investigate the optimal strategy in the regime where $M_{\text{tot}} \gg N \gg 1$ we consider the corresponding continuous optimisation problem. Let $\bar{\mathcal{C}}(\rho)$ be the mean parallel cover time of the search process with searcher density $\rho \sim i/N$ for $N \gg 1$. We define $k(\rho)$ as our real-valued strategy function (as $M_{\text{tot}} \gg N$) that gives the frequency of serial searches that have searcher density ρ . Thus our optimisation process is to select $k(\rho)$ such that we minimise the total cover time functional

$$\Gamma[k] = \int_0^1 k(\rho) \bar{\mathcal{C}}(\rho) d\rho, \quad (3)$$

subject to the constraint

$$N \int_0^1 \rho k(\rho) d\rho = M_{\text{tot}}. \quad (4)$$

We introduce a rescaled strategy function $f(\rho) = N\rho k(\rho)/M_{\text{tot}}$ and total cover time $\tilde{\Gamma}[f] = \Gamma[k]/M_{\text{tot}}$. We note that the rescaled cover time is

$$\tilde{\Gamma}[f] = \int_0^1 \frac{f(\rho) \bar{\mathcal{C}}(\rho)}{N\rho} d\rho = \int_0^1 f(\rho) \mathcal{A}(\rho) d\rho, \quad (5)$$

where $\mathcal{A}(\rho) = \bar{\mathcal{C}}(\rho)/(N\rho)$ is the APCT. Our final version of the optimisation problem is to minimise

$$\tilde{\Gamma}[f] = \int_0^1 f(\rho) \mathcal{A}(\rho) d\rho, \quad (6)$$

subject to the constraint

$$\int_0^1 f(\rho) d\rho = 1. \quad (7)$$

The APCT, $\mathcal{A}(\rho)$, has a minimum value at the optimal density $\rho^* \in (0, 1)$, thus the function that will solve our optimisation problem is $f(\rho) = \delta(\rho - \rho^*)$ or equivalently the optimal strategy is $k(\rho) = M_{\text{tot}} \delta(\rho - \rho^*) / (N\rho)$, i.e. to have all serial searches operating at the optimal density ρ^* . Thus for large networks whose optimal density becomes difficult to evaluate, our results present a simple strategy to optimise mutually excluding search: approximate ρ^* via the easily-computed spectral gap (Fig. 2), and then perform searches in batches at this density.

IV. ASYMMETRIC SEARCH PROCESSES

Until now, we have only considered the symmetric exclusion process, where searchers uniformly sample a target node from their neighbours. However, many biological and physical systems are not in equilibrium, possessing non-zero probability currents in stationary state that markedly change the possible physical behaviours [34]. To model this, we now explore directed networks with bias in the choice of target node. To avoid searchers accumulating at nodes we restrict to flux-conserving (balanced) networks, where each vertex has an equal number of inwards- and outwards-biased edges (and is therefore necessarily of even degree; Fig. 4(a)). For each edge biased $u \rightarrow v$, we define the transition probabilities as $p_{u \rightarrow v} = (1 + 2\varepsilon)/d_u$ and $p_{v \rightarrow u} = (1 - 2\varepsilon)/d_v$, where d_u, d_v are the vertex degrees and $\varepsilon \in [0, 1/2]$ controls the bias. Thus the probability a searcher exits a node through an outwardly biased edge is $1/2 + \varepsilon$.

Asymmetric processes facilitate a significantly greater optimal density than the equivalent unbiased process on the ring lattice. We find numerically that the optimal density for $N = 100$ and $\varepsilon = 0.25$ is $\rho^* = 0.51$ (with high confidence; Appendix A) compared to $\rho^* = 0.05$ in the symmetric case $\varepsilon = 0$. This is a marked increase in ρ^* that even exceeds the symmetrically-biased complete network for which $\rho^* < 0.5$. Here, conservation of flux forces there to be only two viable orientations for the edges, all either clockwise or anticlockwise, effectively adding a constant drift to the symmetric process. As for symmetric exclusion, intuition can be drawn from asymptotic estimates of the APCT in the high density regime for the ring lattice. Let p and $q = 1 - p$ be the probabilities of moving clockwise and anticlockwise, respectively. In Appendix B we show that, for large N and small vacancy count k , the APCT for $M = N - k$ searchers with $p > q$ is approximately $\mathcal{S}_{N-k} \sim (N - 1)/(p - q)k$. This is $\mathcal{O}(N)$ for a non-negligible bias, two orders of magnitude smaller than Eq. (2). In addition, the APCTs in the low and high density regimes are both $\mathcal{O}(N)$ (Appendix B) in contrast to the unbiased process, whose low and high




FIG. 4. (a) A random network for the asymmetric exclusion process, where $N = 100$. The arrows depict the directional bias of the random searchers between nodes. (b) Frequency of optimal densities for a variety of biases. (c) Mean optimal density as a function of spectral gap for symmetric ($\varepsilon = 0$) and asymmetric ($\varepsilon = 0.1$) exclusion processes. Coloured bands are ± 1 s.d. (d) APCT at optimal searcher density for a variety of biases. Results in (b) and (d) are from the same ensemble of 2250 random networks of average degree between 2 and 5 (Appendix G). In (c) an additional 2000 networks are included with average degree between 2 and 3. APCTs in (d) are averages over 10^5 random walk instances.

density APCTs we recall as $\mathcal{O}(N^2)$ and $\mathcal{O}(N^3)$, respectively, indicating bias induces a qualitative change in the APCT.

Beyond the ring lattice, general flux-conserving networks of low degree also see a significantly enhanced optimal density for sufficient bias, with a far narrower distribution of optimal densities (Fig. 4(b)). Dramatic increases in optimal density particularly occur for networks with $\mathcal{G} \approx 0$ (Fig. 4(c)): these have a high concentration of degree two nodes implying many linear chains, the edges of which must have the same directional bias (Fig. 4(a)). More surprisingly, for networks with extremely small spectral gaps ($\mathcal{G} \approx 0$) we see a non-monotonic relationship between \mathcal{G} and ρ^* on average. This phenomenon is also captured by the APCT at optimal density (Fig. 4(d)) where networks with $\mathcal{G} \approx 0$ typically have APCTs below those of networks with greater \mathcal{G} . This shows that the addition of flux-conserving bias can have a counterintuitive impact on search efficiency and optimality (which can be orientation dependent; Appendix F). For example, consider $M_{\text{tot}} = 100$ searchers on the random network in Fig. 4(a) using the optimal serial-parallel strategy versus the same search on the ring lattice. Without bias ($\varepsilon = 0$), the random network has an average search time of 6.6×10^4 and the ring lattice has time 4.8×10^5 . In contrast, with $\varepsilon = 0.25$, these times become 8.5×10^3 and 8.3×10^2 , respectively. Thus, random flux-conserving

bias inverts the networks' relative efficiency, implying care must be taken when attempting to improve the efficiency of network search processes by naïve biasing.

V. SUMMARY

To conclude, we have introduced the APCT as a fundamental measure of how efficiently a network can support mutually excluding random search processes. We found that the optimal density of concurrent searchers can vary from 1–2% to nearly 50% depending on topology, and showed that this topological dependence can be efficiently captured by the spectral gap for both artificial random graphs and real-world transport networks. Finally, we generalised to a biased, non-equilibrium process and uncovered a qualitative change in the topological dependence of optimal densities. This evidences how the transport process itself must be taken into account when optimising random search, and leads us to ask whether there is an efficient statistical characterisation of the topology-transport interplay for general active networked processes [33, 50–52]. More broadly, our work paves the way towards new strategies for topology optimisation in process allocation and flow transport problems across physical and biological networked systems.

ACKNOWLEDGMENTS

This work was supported by the EPSRC Systems Biology DTC EP/G03706X/1 (D.B.W.), a Royal Society Wolfson Research Merit Award (R.E.B.), a Leverhulme Research Fellowship (R.E.B.), the BBSRC UK Multi-Scale Biology Network BB/M025888/1 (R.E.B. and F.G.W.), and Trinity College, Cambridge (F.G.W.).

Appendix A: Numerical evaluation of cover time statistics and optimal densities

All cover time statistics presented in this work were calculated as follows. We do not explicitly sample from the exponential waiting time distribution; rather, we sample the next attempted jump and update the time statistic of interest, t , by the average waiting time, $t = t + \tau/M$. As we are only interested in the first moments of the cover times, linearity of expectation implies that this suffices.

We now detail our numerical procedure for evaluating optimal densities. Suppose we have a network with N nodes and we wish to calculate the optimal number of parallel searchers, M^* , that minimises the APCT. For efficiency we begin using a small number of random walk instances (initially 10^3) that we increase as we get closer to M^* . Starting at $M = 1$ we calculate the average cover time \mathcal{S}_1 , then increase M by one and run another sample of random walk instances, calculating \mathcal{S}_2 and so on. After generating a new APCT for $j + 1$ parallel searchers

we ask whether $\mathcal{S}_{j+1} > \mathcal{S}_j$ to see whether j is a candidate for M^* . If not we increase M by one, however if so then we increase the number of random walk instances to 10^4 and return to the previous number of parallel searchers j , where we recalculate \mathcal{S}_j and \mathcal{S}_{j+1} , and check again to see if $\mathcal{S}_{j+1} > \mathcal{S}_j$. Once we have moved to 10^4 random walk instances we still check to see if the new APCT is greater than the old one, as before. Once we find a density of parallel searchers that appears to be the minimum we increase the number of random walk instances used for the final time to 10^5 . Then the number of searchers that are found to minimise the APCT using 10^5 random walk instances is taken to be M^* , the optimal number of searchers. This was typically enough realisations such that the standard error when calculating the APCTs resulted in mutually excluding 95% confidence intervals either side of the optimal density.

Appendix B: Average parallel cover times for ring-like topologies

1. Complete network

We consider the complete network with N nodes. First we calculate the APCT without interactions. Let X_r^i be the time taken for the i^{th} searcher to first visit $r+1$ distinct sites, given it has just visited r distinct sites. Note that X_r^i is geometrically distributed with success probability $(N-r)/(N-1)$. Then the mean cover time for the single searcher is $\langle \mathcal{C}_1 \rangle = \sum_{r=0}^{N-1} \langle X_r^i \rangle = (N-1)h(N-1)$, where $h(i)$ is the i -th harmonic number. For M non-interacting searchers we start with M distinct sites covered (if every searcher initially occupies its own node) and the parallel cover time is the time taken to visit MN distinct sites (N distinct sites for each of the M searchers). We define Y_r to be the time taken for the overall number of distinct sites visited by the entire population to hit $r+1$ given r distinct sites have just been visited. Y_r is now geometrically distributed with success probability $(MN-r)/(M(N-1))$. Therefore, $\langle \mathcal{C}_M \rangle = 1/M \sum_{r=M}^{MN-1} \langle Y_r \rangle = (N-1)h(M(N-1))$, where the pre-factor of $1/M$ appears as the expected time for the next jump for M random searchers. Thus the APCT for the complete network for non-interacting searchers is

$$\tilde{\mathcal{S}}_M = \frac{N-1}{M} h(M(N-1)). \quad (\text{B1})$$

In order to include the effect of interactions, we need to calculate the probability that an attempted jump is not blocked by another searcher. Due to the unique topology of the complete network this probability is simply $(N-M)/(N-1)$. We can redefine the geometrically distributed variables Y_r to have success probability $(MN-r)/(M(N-1)) \times (N-M)/(N-1)$. The APCT for the complete network with M interacting searchers is

therefore

$$\mathcal{S}_M = \frac{(N-1)^2}{M(N-M)} h(M(N-1)), \quad (\text{B2})$$

where we note that $\mathcal{S}_M = (N-1)/(N-M) \tilde{\mathcal{S}}_M$ is exactly the mean-field approximation (MFA), discussed in Appendix C. Re-writing in terms of the density, $\rho = M/N$, and taking the limit $N \rightarrow \infty$, gives

$$\frac{(N-1)^2}{N\rho(N-N\rho)} h(N\rho(N-1)) \sim \frac{1}{\rho(1-\rho)} \log(\rho N^2). \quad (\text{B3})$$

To find the optimal density, we then differentiate the right-hand side of Eq. (B3) and equate the derivative to zero, this yields the relation

$$\frac{1-\rho^*}{1-2\rho^*} = \log(\rho^*) + \log(N^2). \quad (\text{B4})$$

For the equality (B4) to hold as $N \rightarrow \infty$ we need that $\rho^* \rightarrow 1/2$ (as $\rho^* \rightarrow 0$ yields a trivial solution). Thus the optimal density for parallel search on the complete network is $\rho^* = 1/2$. Substituting in $\rho^* = 1/2 + \xi$ and linearising in ξ we arrive at

$$\xi \approx \frac{1}{2} [1 - 2\log(2N^2)]^{-1}. \quad (\text{B5})$$

For large N , $\xi \sim -(8\log N)^{-1}$ and so $\xi < 0$ and $\xi \rightarrow 0$, thus $\rho^* \rightarrow 1/2$ from below. More over, for $N = 100$ we have $\xi \approx -0.03$ and so $\rho^* \approx 0.47$.

2. Symmetric search on the ring lattice in the high density regime

a. One vacancy

Consider a ring lattice of N nodes with $M = N - 1$ searchers each occupying their own node. Rather than modelling the searchers moving we can instead model the vacancy as a random walker. If the vacancy jumps $N - 1$ times clockwise (anti-clockwise) then every single searcher has jumped one site anti-clockwise (clockwise). Thus we can see that the mean cover time for all searchers to visit all nodes is the mean first passage time of the vacancy to visit $(N-1)^2 + 1$ distinct nodes. On an infinite lattice in one dimension, it is a well known result [20] that the mean first passage time for the span of the random walk to visit r distinct nodes is $r(r-1)/2$. Therefore the mean cover time for $N - 1$ interacting searchers on a ring is

$$\langle \mathcal{C}_{N-1} \rangle = \binom{(N-1)^2 + 1}{2}. \quad (\text{B6})$$

The APCT for the ring lattice with $N - 1$ parallel searchers is therefore

$$\mathcal{S}_{N-1} = \frac{1}{N-1} \binom{(N-1)^2 + 1}{2}. \quad (\text{B7})$$

For large N , this is $\mathcal{O}(N^3)$.




FIG. 5. APCT for the ring lattice in the high density regime where $N = 100$. The +’s represent data generated from simulations of the search process, the solid curve is an interpolation of the APCT calculated using the asymptotic estimates in Eq. (2). All APCTs were calculated from 10^5 random walk instances.

b. Several vacancies

We again have a ring lattice with N nodes. Let $M = N - k$ be the number of searchers and suppose that $k \ll N$ so we are still in the high density regime. We now consider, as before, the random walk of the k vacancies instead of the searchers directly. In the limit of large N the vacancies are non-interacting and therefore we can treat the vacancies as being approximately independent. The net displacement of the k vacancies can then be approximated as the sum of k independent normally distributed random variates with mean zero. Equivalently we consider a single vacancy that moves on average k times faster, that is, with mean waiting time $\tau_{\text{vac}} = 1/k$. This random walker needs to visit $(N - 1)(N - k) + 1$ distinct sites in order for the all the searchers to have visited every site, thus the mean cover time is

$$\langle \mathcal{C}_{N-k} \rangle \sim \frac{1}{k} \left(\frac{(N-1)(N-k) + 1}{2} \right). \quad (\text{B8})$$

Note that the factor of $1/k$ appears because the collective vacancy moves k times faster than a single vacancy. The APCT for the ring lattice in the high density regime is therefore

$$\mathcal{S}_{N-k} \sim \frac{1}{k(N-k)} \left(\frac{(N-1)(N-k) + 1}{2} \right). \quad (\text{B9})$$

We see that the APCT in the high density regime is still of magnitude $\mathcal{O}(N^3)$. The asymptotic estimate (B9) of the APCT matches well to simulations in the high density regime (Fig. 5).

3. Asymmetric search on the ring lattice in the high density regimes

a. Single searcher

We define a random walk on the integers, \mathbb{Z} , where the searcher starts at the origin and at each timestep moves to the right with probability p or left with probability $q = 1 - p$. Chong et. al [53] showed that the expected number of steps to reach m distinct sites (not including the origin) is

$$\langle T_m \rangle = \frac{g_{m-1} h_{m-1}}{s_{m-1} s_m}, \quad (\text{B10})$$

where $g_k = \sum_{i=0}^k (i+1) p^{k-i} q^i$, $h_k = \sum_{i=0}^k (i+1) p^i q^{k-i}$, and $s_k = \sum_{i=0}^k p^i q^{k-i}$. For a lattice of length N we can write down directly the expected cover time,

$$\langle \mathcal{C}_1 \rangle = \frac{g_{N-2} h_{N-2}}{s_{N-2} s_{N-1}}. \quad (\text{B11})$$

For $p > q$ and large N it is known that the cover time is asymptotically

$$\langle \mathcal{C}_1 \rangle \sim \frac{1}{p-q} \left(N - 1 - \frac{q}{p-q} + \mathcal{O}(N^2 (q/p)^N) \right). \quad (\text{B12})$$

b. Single vacancy

We now consider $N - 1$ interacting searchers on a ring lattice of length N with asymmetric transition probabilities p and q . In order to calculate the parallel cover time we consider the vacancy as a random walker on the integers \mathbb{Z} . The vacancy has to visit $(N - 1)^2$ distinct sites (not including the origin) in order for every searcher to have visited every site. Therefore, if we let $m = (N - 1)^2$, the average number of jumps the vacancy has to make is given in Eq. (B10). In the limit of large N we have that the APCT for $N - 1$ searchers is

$$\mathcal{S}_{N-1} \sim \frac{1}{(p-q)(N-1)} \left((N-1)^2 - \frac{q}{p-q} \right). \quad (\text{B13})$$

Thus, for moderate bias where $p - q$ is not small, the APCT for $N - 1$ searchers is two orders of magnitude smaller than the APCT in the case of the symmetric exclusion process.

c. Multiple vacancies

We consider the case where we have k vacancies, where $k \ll N$. As for the symmetric exclusion process we consider instead a single asymmetric random walker that moves on average k times faster than a single vacancy, and we are interested in the time taken for the vacancy




FIG. 6. The APCTs of an asymmetric exclusion process on the ring lattice, where $p = 0.75$ and $q = 0.25$, from simulation (asterisks) and from the analytical approximation in the high density regime (solid line; (B14)). All APCTs were calculated using 10^5 random walk instances.

to visit $m = (N-1)(N-k)$ new sites (not including the origin). Therefore we can write down the APCT for the high density regime as

$$\begin{aligned} \mathcal{S}_{N-k} &\sim \frac{1}{(p-q)k(N-k)} \left((N-1)(N-k) - \frac{q}{p-q} \right) \\ &\sim \frac{N-1}{(p-q)k}. \end{aligned} \quad (\text{B14})$$

In Fig. 6 we provide a plot of the APCT for an asymmetric random walk on a ring lattice with transition probabilities $p = 0.75$ and $q = 0.25$ for all numbers of searchers, as well as the estimates in the high density regime, where we see excellent agreement for $k \leq 5$.

Appendix C: The mean-field approximation

The MFA rescales non-excluding cover times by approximating how many additional jumps are needed when exclusion is considered, thus accounting for aborted moves. In order to approximate this scaling we consider the occupancy probabilities of nodes adjacent to a searcher attempting a move. Let \mathcal{O}_ν denote the occupancy of node ν , where a value 1 corresponds to the node being occupied and a value 0 to the node being vacant. For a searcher at node ω , let ν be the neighbouring site which the searcher attempts a move to. The probability the neighbouring node is vacant (i.e. the move is successful) is $\mathbb{P}(\mathcal{O}_\nu = 0 | \mathcal{O}_\omega = 1)$. Bayes' rule allows this probability to be written in terms of a two-site occupancy probability, $\mathbb{P}(\mathcal{O}_\nu = 0, \mathcal{O}_\omega = 1) / \mathbb{P}(\mathcal{O}_\omega = 1)$. The MFA comes from approximating the two-site occupancy probability to be the product of single-site occupancy probabilities:

$\mathbb{P}(\mathcal{O}_\nu = 0, \mathcal{O}_\omega = 1) \approx \mathbb{P}(\mathcal{O}_\nu = 0)\mathbb{P}(\mathcal{O}_\omega = 1)$. Therefore the probability of the attempted move being successful is approximately $\mathbb{P}(\mathcal{O}_\nu = 0)$ and thus is independent of the searcher location ω . This probability is the fraction of the number of vacant sites over the number of total sites that are not ω , namely, $(N-M)/(N-1)$. Denoting $\langle \mathcal{C} \rangle$ and $\langle \tilde{\mathcal{C}} \rangle$ to be the excluding and non-excluding cover times respectively, we approximate $\langle \mathcal{C} \rangle = \langle \tilde{\mathcal{C}} \rangle \times (N-1)/(N-M)$ or equivalently $\langle \mathcal{C} \rangle = \phi \langle \tilde{\mathcal{C}} \rangle$, where $\phi = (1 - N^{-1})/(1 - \rho)$ and $\rho = M/N$.

We consider the cover time of a single searcher on a general network. Following Ref. [16], suppose that the cover time is Gumbel distributed with parameters $\mu \in \mathbb{R}$ and $\beta > 0$. The probability density function $p(t)$ of the cover time is

$$p(t) = \frac{1}{\beta} \exp \left[-\frac{t-\mu}{\beta} - \exp \left(-\frac{t-\mu}{\beta} \right) \right], \quad (\text{C1})$$

where γ is the Euler–Mascheroni constant and the mean cover time is $\gamma\beta + \mu > 0$. Introducing the general scaling $x = (t - \mu)/\beta$ gives the standard Gumbel distribution.

Now let the maximum cover time of M independent random walkers be $\sigma = \max\{t_1, \dots, t_M\}$, or in the rescaled variables $y = \max\{x_1, \dots, x_M\}$. Due to independence of the individual cover times y is also Gumbel distributed with $\langle y \rangle = \gamma + \log(M)$. We transform the coordinates back to give $\langle \sigma \rangle = \beta\{\gamma + \log(M) + \mu/\beta\}$. The APCT for M non-interacting searchers is then

$$\tilde{\mathcal{S}}_M = \frac{\beta}{M} \left\{ \gamma + \log(M) + \frac{\mu}{\beta} \right\}. \quad (\text{C2})$$

(Note that for the APCT with interactions we cannot make analytic progress with $\sigma = \max\{t_1, \dots, t_M\}$, as before, because the individual cover times, t_i , are no longer independent and so we cannot easily apply extremal value theory.) The MFA $\mathcal{S}_M = \phi \tilde{\mathcal{S}}_M$, where $\phi = (1 - N^{-1})/(1 - \rho)$, then gives an approximate expression for the APCT with interactions

$$\mathcal{S}_M = \frac{\beta(1 - N^{-1})}{N\rho(1 - \rho)} \left\{ \gamma + \log(N\rho) + \frac{\mu}{\beta} \right\}. \quad (\text{C3})$$

We would like to find the density that minimises the APCT for a network that has a single particle cover time that is Gumbel distributed with parameters (μ, β) . Setting $d\mathcal{S}_M/d\rho = 0$ to find the optimal $\rho = \rho^*$ yields

$$\frac{1 - \rho^*}{1 - 2\rho^*} = \gamma + \frac{\mu}{\beta} + \log(N) + \log(\rho^*), \quad (\text{C4})$$

similar to Eq. (B4). Note that $\gamma + \mu/\beta > 0$ as the mean cover time must be positive and β is always positive. Then for Eq. (C4) to hold as $N \rightarrow \infty$ we require that $\rho^* \rightarrow 1/2$ (as $\rho^* \rightarrow 0$ yields a trivial solution). Substituting $\rho^* = 1/2 + \xi$ into Eq. (C4) and linearising in ξ we find that

$$\xi \approx \frac{1}{2} \left[1 - 2 \left(\gamma + \frac{\mu}{\beta} + \log \frac{N}{2} \right) \right]^{-1}. \quad (\text{C5})$$




FIG. 7. The rescaled distributions for the cover times of all 1500 networks used to generate Fig. 2 (red) against the standard Gumbel distribution (black). The cover time distributions for each network are estimated using 10^6 random walk instances.

Thus, as $N \rightarrow \infty$, we have $\xi < 0$ and $\xi \rightarrow 0$, so we conclude that $\rho^* \rightarrow 1/2$ from below.

In order to justify the assumption that the cover times are Gumbel distributed, for all 1500 networks used to generate Fig. 2 we simulated $n = 10^6$ cover times, t_i , for single searchers and used the data to calculate maximum likelihood estimators for the Gumbel distribution which are

$$\hat{\beta} = \frac{1}{n} \sum_{i=1}^n t_i - \frac{\sum_{i=1}^n t_i \exp(-t_i/\hat{\beta})}{\sum_{i=1}^n \exp(-t_i/\hat{\beta})}, \quad (\text{C6})$$

and

$$\hat{\mu} = -\hat{\beta} \log \left\{ \frac{1}{n} \sum_{i=1}^n \exp(-t_i/\hat{\beta}) \right\}. \quad (\text{C7})$$

Using the maximum likelihood estimators we rescaled the cover times using the transformation $x = (t - \hat{\mu})/\hat{\beta}$ and present the distribution of cover times for each of the networks against the standard Gumbel distribution. We see in Fig. 7 that across the entire range of the networks seen in Fig. 2 the Gumbel approximation is valid.

Appendix D: Generating random symmetric networks

In this section we explain the procedure used to sample the random networks seen in Fig. 2.

- 1: Select m and N , the required number of edges and the number of nodes of the network respectively, such that $m \geq N$.
- 2: Sample uniformly the target degree distribution, \vec{d} , such that every node has degree at least two, i.e. $\vec{d} = 2(1, \dots, 1) + \text{Multinomial}(2(m - N), N)$.




FIG. 8. The same networks as in Fig. 2, with optimal density now plotted against average network degree.

- 3: Allocate d_i stubs to the i -th node for $i \in \{1, \dots, N\}$.
- 4: **while** there are stubs remaining **do**
- 5: **if** the only remaining connections to choose from result in self-loops or multiple-edges **then** start over and return to step 2.
- 6: **else** select two distinct nodes (avoiding self-loops) with stubs remaining uniformly at random and connect them with an edge only if they have not already been connected (avoiding multiple-edges).
- 7: **end if**
- 8: **end while**
- 9: **if** the resulting network is not connected **then** return to step 2.
- 10: **else** take the resulting network to be a single realisation.
- 11: **end if**

Appendix E: Quantifying the predictive capacity of the spectral gap versus the average degree

Here, we evidence the lesser predictive capacity of the average degree compared to the spectral radius. We present estimates for the uncertainty coefficient between the optimal density, ρ^* , and two network properties: the spectral gap, \mathcal{G} , and the average degree, \bar{d} . For two random variables X and Y , the uncertainty coefficient is defined as $U(X|Y) = I(X, Y)/H(X)$, where $I(X, Y)$ is the mutual information between X and Y and $H(X)$ is the Shannon entropy of X . This measures what proportion of the total possible information about X we can learn by observing Y , with $0 \leq U(X|Y) \leq 1$. A coefficient of zero means that X and Y are independent, while a coefficient of one means that Y completely predicts X .

As we do not have access to the distributions for the optimal density or the spectral gap, we must rely on estimations of the mutual information. We use the algorithm




FIG. 9. Two sampled networks from Fig. 2 that have identical degree distribution but markedly different optimal densities: (a) $\rho^* = 0.01$, (b) $\rho^* = 0.11$.

proposed by Ross [54]. For the data presented in Fig. 2 the uncertainty coefficient between the optimal density and the spectral gap is estimated to be $U(\rho^*|\mathcal{G}) = 0.4835$ while for the same networks the average degree gives an uncertainty coefficient of $U(\rho^*|\bar{d}) = 0.3769$, demonstrating greater mutual information between the optimal density and the spectral gap compared to the average degree.

To see this, in Fig. 8 we replot Fig. 2 with average degree on the horizontal axis. While there is still a good correlation, visual inspection suggests that even our restricted ensemble networks of low average degree are not as well correlated as with spectral radius, with a greater number of outliers. Our two real-world networks are notable examples: the London Underground network, marked in Fig. 8, clearly deviates from the bulk, while the US airports network is an even worse match to the trend at the off-axis point $\bar{d} = 11.9, \rho^* = 0.26$. Indeed, consider two of the sampled networks from Fig. 2, shown in Fig. 9, which have identical degree distributions $\{3, 3, 3, 3, 2, \dots, 2\}$: we find that the network in Fig. 9(a) has an optimal density $\rho^* = 0.01$ while that in Fig. 9(b) has $\rho^* = 0.11$, differing by a factor of 11. Thus average degree (or any degree distribution statistic) can be a poor predictor of optimal density. However, the spectral gaps for these networks are, respectively, $\mathcal{G} = 0.0005$ and $\mathcal{G} = 0.0012$, with the more-than-doubled spectral gap of (b) compared to (a) accounting for its higher optimal density, particularly since small changes in \mathcal{G} near zero have large effects on ρ^* .

Appendix F: Optimal density and asymmetric orientations

There are typically multiple topologically distinct ways to bias the edges of the same undirected network, even when fluxes are balanced. These different biases can induce different cover time properties. For example, consider the five distinct orientations of the same network in Fig. 10(a-e). For bias strength $\varepsilon = 0.1$ these have optimal densities $\rho^* = 0.40, 0.41, 0.41, 0.40$ and 0.40 for (a) to (e) respectively (with high statistical confidence). While these differences are slight, the corresponding APCTs differ more markedly (Fig. 10(f)), and such differences will




FIG. 10. (a-e) Five distinct flux-conserving (balanced) bias orientations for a given fixed undirected network. (f) The APCTs for each of the graphs (a-e) as a function of searcher density near the optima. Each APCT is calculated from 10^6 random walk instances.

be particularly important if a large number of searchers are being employed in a hybrid series-parallel search strategy. In larger networks, we would expect this variation between different directional biases to be amplified.

Appendix G: Generating random balanced directed networks

In this section we describe the sampling procedure for the asymmetric random networks in Fig. 4.

- 1: Select m and N , the required number of edges and the number of nodes of the network respectively, such that $m \geq N$.
- 2: Sample uniformly the target degree distribution for

- outwardly biased edges, \vec{d}_{out} , such that every node has outward degree at least one, i.e. $\vec{d}_{\text{out}} = (1, \dots, 1) + \text{Multinomial}(m - N, N)$.
- 3: Set the degree distribution for the inwardly biased edges to be $\vec{d}_{\text{in}} = \vec{d}_{\text{out}}$.
 - 4: Allocate $\vec{d}_{\text{out}}(i) = \vec{d}_{\text{in}}(i)$ outward stubs and inward stubs to the i -th node for $i \in \{1, \dots, N\}$.
 - 5: **while** there are stubs remaining **do**
 - 6: **if** the only remaining connections to choose from result in self-loops or multiple-edges **then** start over and return to step 2.
 - 7: **else** select uniformly at random a node with an outward stub and a distinct node with an inward stub. Connect these two nodes with an edge with preferential direction from the outward node to the inward node as long as these nodes have not been connected before.
 - 8: **end if**
 - 9: **end while**
 - 10: **if** the resulting network is not connected **then** return to step 2.
 - 11: **else** take the resulting network to be a single realisation.
 - 12: **end if**

-
- [1] O. Bénichou, C. Loverdo, M. Moreau, and R. Voituriez, *Rev. Mod. Phys.* **83**, 81 (2011).
 - [2] J.-F. Rupprecht, O. Bénichou, and R. Voituriez, *Phys. Rev. E* **94**, 012117 (2016).
 - [3] M. Coppey, O. Bénichou, R. Voituriez, and M. Moreau, *Biophys. J.* **87**, 1640 (2004).
 - [4] F. Bartmeus and J. Catalan, *J. Phys. A-Math. Theor.* **42**, 434002 (2009).
 - [5] M. L. Heuzé, P. Vargas, M. Chabaud, M. L. Berre, Y. J. Liu, O. Collin, P. Solanes, R. Voituriez, M. Piel, and A. M. Lennon-Duménil, *Immunol. Rev.* **256**, 240 (2013).
 - [6] G. M. Viswanathan, E. P. Raposo, and M. G. E. da Luz, *Phys. Life Rev.* **5**, 133 (2008).
 - [7] G. M. Viswanathan, S. V. Buldyrev, S. Havlin, M. G. E. da Luz, E. P. Raposo, and H. E. Stanley, *Nature* **401**, 911 (1999).
 - [8] E. A. Codling, M. J. Plank, and S. Benhamou, *J. R. Soc. Interface* **5**, 813 (2008).
 - [9] G.-W. Li, O. G. Berg, and J. Elf, *Nat. Phys.* **5**, 294 (2009).
 - [10] K. Schwarz, Y. Schröder, B. Qu, M. Hoth, and H. Rieger, *Phys. Rev. Lett.* **117**, 068101 (2016).
 - [11] M. Rubenstein, A. Cornejo, and R. Nagpal, *Science* **345**, 795 (2014).
 - [12] M. Herty and A. Klar, *SIAM J. Sci. Comput.* **25**, 1066 (2008).
 - [13] N. Bellomo and C. Dogbe, *SIAM Rev.* **53**, 409 (2011).
 - [14] A. Fügenschuh, M. Herty, A. Klar, and A. Martin, *SIAM J. Optim.* **16**, 1155 (2006).
 - [15] L. A. Adamic, R. M. Lukose, A. R. Puniyani, and B. A. Huberman, *Phys. Rev. E* **64**, 046135 (2001).
 - [16] M. Chupeau, O. Bénichou, and R. Voituriez, *Nat. Phys.* **11**, 844 (2015).
 - [17] J. D. Kahn, N. Linial, N. Nisan, and M. E. Saks, *J. Theor. Prob.* **2**, 121 (1989).
 - [18] A. Kundu, S. N. Majumdar, and G. Schehr, *Phys. Rev. Lett.* **110**, 220602 (2013).
 - [19] S. N. Majumdar, S. Sabhapandit, and G. Schehr, *Phys. Rev. E* **94**, 062131 (2016).
 - [20] C. S. O. Yokoi, A. Hernández-Machado, and L. Ramírez-Piscina, *Phys. Lett. A* **145**, 82 (1990).
 - [21] J. Ding, *Electron. J. Probab.* **17**, 1 (2012).
 - [22] M. J. A. M. Brummelhuis and H. J. Hilhorst, *Physica A* **176**, 387 (1991).
 - [23] B. F. Maier and D. Brockmann, *Phys. Rev. E* **96**, 042307 (2017).
 - [24] V. Tejedor, O. Bénichou, and R. Voituriez, *Phys. Rev. E* **80**, 065194 (2009).
 - [25] B. Meyer, E. Agliari, O. Bénichou, and R. Voituriez, *Phys. Rev. E* **85**, 026113 (2012).
 - [26] J. Palacios, *Braz. J. Probab. Stat.* **17**, 127 (2003).
 - [27] J. Jonasson, *Comb. Probab. Comput.* **7**, 265 (1998).
 - [28] T. Weng, J. Zhang, M. Small, and P. Hui, *Phys. Rev. E* **95**, 052103 (2017).
 - [29] T. Liggett, *Stochastic Interacting Systems: Contact, Voter and Exclusion Processes* (Springer, 1999).
 - [30] M. J. Simpson, B. D. Hughes, and K. A. Landman, *Australas. J. Eng. Ed.* **15**, 59 (2009).
 - [31] M. J. Plank and M. J. Simpson, *J. R. Soc. Interface* **9**, 2983 (2012).
 - [32] D. Chowdhury, A. Schadschneider, and K. Nishinari, *Phys. Life Rev.* **2**, 318 (2005).
 - [33] I. R. Graf and E. Frey, *Phys. Rev. Lett.* **118**, 128101 (2017).
 - [34] I. Neri, N. Kern, and A. Parmeggiani, *Phys. Rev. Lett.* **107**, 068702 (2011).
 - [35] M. Plischke, Z. Rácz, and D. Liu, *Phys. Rev. B* **35**, 3485 (1987).
 - [36] C. Kipnis, *Ann. Probab.* **14**, 397 (1986).
 - [37] This is contrary to the usual cover time, where a maximum over initial position is taken [45]. Here, this is combinatorially intractable on most networks when $M \gg 1$, and our definition allows us to consider topological effects removed from any influence of initial conditions.
 - [38] H. van Beijeren, K. W. Kehr, and R. Kutner, *Phys. Rev. B* **28**, 5811 (1983).
 - [39] L. Lizana and T. Ambjörnsson, *Phys. Rev. Lett.* **100**, 200601 (2008).
 - [40] C. Lutz, M. Kollmann, and C. Bechinger, *Phys. Rev. Lett.* **93**, 026001 (2004).
 - [41] Q. Wei, C. Bechinger, and P. Leiderer, *Science* **287**, 625 (2000).
 - [42] T. Imamura, K. Mallick, and T. Sasamoto, *Phys. Rev. Lett.* **118**, 160601 (2017).
 - [43] P. L. Krapivsky, K. Mallick, and T. Sadhu, *Phys. Rev. Lett.* **113**, 078101 (2014).
 - [44] M. E. J. Newman and D. J. Watts, *Phys. Lett. A* **263**, 341 (1999).
 - [45] L. Lovász, in *Combinatorics, Paul Erdős is Eighty*, Vol. 2 (János Bolyai Mathematical Society, Budapest, 1993) pp. 353–398.
 - [46] M. E. J. Newman, *Networks: An Introduction* (Oxford

- University Press, 2010).
- [47] H. Chen and F. Zhang, *Discrete Appl. Math.* **155**, 654 (2007).
- [48] <https://wikis.bris.ac.uk/display/ipshe/London+Tube>.
- [49] V. Colizza, R. Pastor-Satorras, and A. Vespignani, *Nat. Phys.* **3**, 276 (2007).
- [50] A. Caspi, R. Granek, and M. Elbaum, *Phys. Rev. Lett.* **85**, 5655 (2000).
- [51] T. Becker, K. Nelissen, B. Cleuren, B. Partoens, and C. Van den Broeck, *Phys. Rev. Lett.* **111**, 110601 (2013).
- [52] F. G. Woodhouse, A. Forrow, J. B. Fawcett, and J. Dunkel, *Proc. Natl. Acad. Sci. U.S.A.* **113**, 8200 (2016).
- [53] K. S. Chong, R. Cowan, and L. Holst, *Adv. Appl. Probab.* **32**, 177 (2000).
- [54] B. C. Ross, *PLOS ONE* **9**, e87357 (2014).