

The Life of the Party: Diffusion in Networks under Sequential Participation

ECON284/CS360: Simplicity and Complexity in Economic Theory

Lyle Goodyear

Annika Younge

Abstract

We study a model of party attendance in which guests arrive in a uniformly random order and irrevocably decide whether to stay based on how many friends are already present. The final party is determined by the interaction between the social network, the arrival order, and a threshold rule. For threshold $k = 1$, we show the final party equals the reachable set from the hosts in the permutation-induced directed acyclic graph. This equivalence yields a path-counting framework that produces exact, closed-form expressions for the expected party size over time on trees, rings, and complete bipartite graphs. We show that party success depends not on the number of connections but on their structure. For the seed selection problem (sending initial invites), we prove that when $k = 1$, the expected party size is submodular in the seed set, so greedy host selection would achieve at least a $(1 - 1/e)$ -approximation, but when $k \geq 2$, submodularity fails in this diffusion model.

Keywords

Diffusion on networks, Graph orientations, Collective action

1 Introduction

What makes a party successful? A party forms over the course of the night; each guest decides whether to stay based on who else is there, but who is there depends on who arrived earlier and stayed. The life of the party therefore is determined by the interaction between the social network (who knows who), the arrival order (who shows up when), and a decision to attend (you stay if you see enough friends).

We formalize this interaction as a threshold diffusion process on a fixed graph with an independent arrival order for nodes. We first fix a set of host nodes. Guests arrive in a uniformly random order and make an irrevocable decision to stay or to leave. If each guest stays as soon as one friend is present, then on a connected network the best arrival order fills the entire room, while the worst arrival order leaves a party just with the hosts and their immediate neighbors. The expected final size of the party depends heavily on the structure of social connections in the network.

The paper makes three contributions. First, we establish an equivalence between the party process and reachability in a directed acyclic graph with random edge orientations, and we use this equivalence to derive exact equations for the expected party size on simple graph families. Second, we show through simulation that network structure, specifically the amount of clustering, is the main determinant of expected party success, even with fixed-degree. Third, we show that the seed-selection (invite) problem has guarantees that depend on individual decision-thresholds; when a guest needs only one friend present, $k = 1$, expected party size

is submodular in invites and greedy host selection is near-optimal. But for $k \geq 2$, we provide a counterexample to prove that the expected party size is not necessarily submodular, and sending invites can make initial hosts complements rather than substitutes. We conclude with potential extensions to this model. We additionally note that although application is deliberately simple, a model of diffusion with deterministic arrivals and departures may be relevant to other wide-ranging applications.

2 Model

Let $G = (V, E)$ be a simple undirected graph on n nodes representing a social network: an edge $\{u, v\}$ means u and v know each other. Let $S_0 \subseteq V$ denote the *seed set* (the hosts), with $|S_0| = m$. Let $k \geq 1$ be a uniform threshold. The arrival order is a permutation σ drawn uniformly at random from all $(n - m)!$ orderings of $V \setminus S_0$. The party evolves as follows. Set $P(0) = S_0$. At each step $t = 1, \dots, n - m$:

- (i) Node $\sigma(t)$ arrives and observes the current party $P(t - 1)$.
- (ii) If $|N(\sigma(t)) \cap P(t - 1)| \geq k$, then $\sigma(t)$ *stays*: $P(t) = P(t - 1) \cup \{\sigma(t)\}$.
- (iii) Otherwise, $\sigma(t)$ *leaves permanently*: $P(t) = P(t - 1)$.

The final party is $P^* = P(n - m)$, and the party size $|P^*|$ is a random variable over the permutation σ . The key features of this model are: the arrival order is uniformly random and independent of the network; the stay-or-leave decision is deterministic given the current party; and departure is *irrevocable*. A guest who leaves never reconsiders, even if friends arrive later.

3 Reachability Equivalence

The arrival order induces a directed orientation of G : orient each edge $\{u, v\}$ from its earlier-arriving endpoint to its later-arriving one, with hosts (present from the start) pointing to all of their non-host neighbors. The result is a directed acyclic graph (DAG), since the arrival times give a topological order.

PROPOSITION 3.1 (REACHABILITY EQUIVALENCE). *For $k = 1$: node $v \in P^*$ if and only if v is reachable from S_0 in the permutation-induced DAG.*

PROOF. ($\Rightarrow$) If v stayed, then when v arrived some neighbor u was already present, so u arrived before v and $u \in P^*$. The edge $\{u, v\}$ is therefore oriented $u \rightarrow v$, and by induction u is reachable from S_0 , so v is as well.

($\Leftarrow$) Suppose there is a directed path $s = w_0 \rightarrow w_1 \rightarrow \dots \rightarrow w_d = v$ with $s \in S_0$, so that $w_0, w_1, \dots, w_d$ arrive in increasing order. When w_1 arrived, s was present (hosts are always present), so w_1 stayed. When w_2 arrived, w_1 was present, so w_2 stayed. Continuing, $v = w_d$ stayed. $\square$

The party is therefore the set of nodes reachable from the seed set via paths whose nodes arrive in increasing order. This transforms the stochastic process into a combinatorial question about reachability in a random directed acyclic graph (DAG).

4 Expected Party Size

We now focus on the problem of computing the expected party size given a fixed network and seed set. By linearity of expectation:

$$\mathbb{E}[|P^*|] = |S_0| + \sum_{v \notin S_0} \Pr(v \in P^*). \quad (1)$$

Computing the expected party size thus reduces to computing, for each node v , the probability (over the random ordering) that v is reachable from S_0 . By the reachability equivalence, this is the probability that at least one simple path from S_0 to v is directed in the random DAG.

A single path of length d (using d edges) is directed if and only if its d non-seed nodes arrive in increasing order. In a uniform random permutation, the relative ordering of any d specific elements is uniform over all $d!$ possibilities, so:

$$\Pr(\text{path of length } d \text{ is directed}) = \frac{1}{d!}. \quad (2)$$

This depends only on the path length, not on n or the rest of the graph. When there are multiple paths from S_0 to v , node v is reachable if *any* one of them is directed. By inclusion–exclusion over all r simple paths $P_1, \dots, P_r$ from S_0 to v :

$$\Pr(v \in P^*) = \sum_i \frac{1}{d_i!} - \sum_{i < j} \Pr(P_i, P_j \text{ both directed}) + \dots \quad (3)$$

This is intractable to compute in general: the number of simple paths between two nodes can be exponential in dense graphs, and each overlap term requires computing the probability that multiple paths are simultaneously directed. The next section shows that when the paths are internally disjoint, every term has a clean closed form.

5 Counting Paths

Disjoint paths. We say that two s – v paths are *internally disjoint* if they share only the endpoints s and v . When all paths under consideration are internally disjoint, the inclusion–exclusion terms factorize cleanly.

THEOREM 5.1 (DISJOINT PATH FORMULA). *Let $P_1, \dots, P_r$ be internally disjoint paths from S_0 to v with lengths $d_1, \dots, d_r$. Then for any nonempty subset $S \subseteq [r]$ of paths,*

$$\Pr(\text{all paths in } S \text{ are simultaneously directed}) = \frac{1}{D_S} \cdot \prod_{i \in S} \frac{1}{(d_i - 1)!},$$

where $D_S = 1 + \sum_{i \in S} (d_i - 1)$ is the total number of distinct non-seed nodes across all paths in S .

Consequently,

$$\Pr(v \in P^*) = \sum_{\substack{S \subseteq [r] \\ S \neq \emptyset}} (-1)^{|S|+1} \cdot \frac{1}{D_S} \cdot \prod_{i \in S} \frac{1}{(d_i - 1)!}. \quad (4)$$

PROOF. Consider a subset S of paths from S_0 to v . The non-seed nodes appearing on these paths are: for each $i \in S$, the $d_i - 1$ internal nodes of path P_i , plus the shared endpoint v . Since the paths are internally disjoint, these are all distinct, totaling $D_S =$

$1 + \sum_{i \in S} (d_i - 1)$ nodes. The relative ordering of these D_S nodes is uniform over all $D_S!$ possibilities, regardless of n . The ordering constraints for all paths in S to be simultaneously directed are: (1) v must arrive last among all D_S nodes (v must be at the end of every path), and (2) each path P_i 's internal nodes must arrive in increasing order among themselves.

The probability that v arrives last in every path is $1/D_S$. Given v is last, the remaining $D_S - 1$ positions must be filled by $|S|$ groups of internal nodes, each in increasing order but free to interleave. The number of such interleavings is the multinomial coefficient $\binom{D_S - 1}{d_1 - 1, \dots, d_{|S|} - 1}$, and dividing by $(D_S - 1)!$ total arrangements gives $\prod_{i \in S} 1/(d_i - 1)!$. The inclusion–exclusion formula (4) follows immediately. $\square$

Time-indexed formula. The formula above characterizes the final party P^* , when all non-seed nodes have arrived. We can modify it to track the party over time. For v to be at the party after t arrivals via a specific path of length d from some $s \in S_0$ to v , the d non-seed nodes on the path must all arrive within the first t slots and in the correct order. These conditions are independent:

$$\Pr(\text{length-}d \text{ path directed by time } t) = \frac{\binom{t}{d}}{\binom{n - |S_0|}{d}} \cdot \frac{1}{d!}. \quad (5)$$

The first factor is the probability that d specific elements of a random permutation of $n - |S_0|$ non-seed nodes all land in the first t positions: there are $\binom{n - |S_0|}{d}$ equally likely sets of d positions, and $\binom{t}{d}$ of those lie within the first t . The second factor is the probability of the correct ordering. At $t = n - |S_0|$ (everyone has arrived), the first factor equals 1 and we recover $1/d!$. The full time-indexed inclusion–exclusion for disjoint paths is obtained by replacing each term's probability with the corresponding time-indexed version.

6 Results for Simple Graphs

The time-indexed inclusion–exclusion formula gives exact curves for the expected party size over time on any graph family in which the paths from a node to the seed set are internally disjoint.

Trees. On a tree where the root is the only seed node, there is exactly one simple path from the seed s to each node v , of length $d = d_T(s, v)$. The inclusion–exclusion formula has a single term.

THEOREM 6.1 (TREE FORMULA). *Let T be a tree with root seed s and threshold $k = 1$. For every node v ,*

$$\Pr(v \in P^*) = \frac{1}{d_T(s, v)!}.$$

Consequently, $\mathbb{E}[|P^*|] = \sum_{v \in V} 1/d_T(s, v)!$.

PROOF. On a tree, the unique path from s to v has length $d = d_T(s, v)$. The d non-seed nodes on this path must arrive in increasing order, which occurs with probability $1/d!$. $\square$

COROLLARY 6.2 (TREES CANNOT SUPPORT LARGE PARTIES). *On a tree where every node has b children, $\mathbb{E}[|P^*|] = \sum_{d=0}^D b^d/d!$, which converges to e^b as the depth $D \rightarrow \infty$. More generally, for any sequence of trees (T_n, s_n) with bounded maximum degree Δ and $n \rightarrow \infty$,*

$$\frac{\mathbb{E}[|P^*|]}{n} \rightarrow 0.$$

PROOF. On a tree with maximum degree Δ , the number of nodes at distance d is at most $\Delta(\Delta - 1)^{d-1}$. Therefore $\mathbb{E}[|P^*|] \leq 1 + \sum_{d \geq 1} \Delta(\Delta - 1)^{d-1}/d! = 1 + \frac{\Delta}{\Delta-1}(e^{\Delta-1} - 1)$, a constant depending only on Δ . Dividing by $n \rightarrow \infty$ gives the result. $\square$

Distance from the host carries a factorial penalty in the likelihood of staying. A node at distance d joins only if d intermediaries arrive in a specific order, which occurs with probability $1/d!$. This decays faster than any exponential, so even large trees with high branching factors produce bounded parties. A binary tree with a million nodes produces a party of about $e^2 \approx 7.4$ people in expectation.

Ring networks. On the ring C_n with seed at node 0, each node d has exactly two internally disjoint paths: the clockwise path of length d and the counterclockwise path of length $n - d$. The inclusion-exclusion has three terms.

PROPOSITION 6.3 (RING FORMULA). *On C_n with seed at node 0,*

$$\Pr(d \in P^*) = \frac{1}{d!} + \frac{1}{(n-d)!} - \frac{1}{(n-1)(d-1)!(n-d-1)!}.$$

PROOF. Apply Theorem 5.1 with two paths of lengths $d_1 = d$ and $d_2 = n - d$. The single-path terms give $1/d!$ and $1/(n-d)!$. For the pair term, $D_{\{1,2\}} = 1 + (d-1) + (n-d-1) = n-1$, giving $1/((n-1)(d-1)!(n-d-1)!)$. Substituting into (4) yields the result. $\square$

As $n \rightarrow \infty$, the counterclockwise path has length $n - d \rightarrow \infty$, so its contribution $1/(n-d)!$ and the correction term both vanish. The expected party size converges to $1 + 2 \sum_{d=1}^{\infty} 1/d! = 2e - 1$. The ring therefore adds $e - 1$ people over the path baseline of e .

Complete bipartite graphs. On $K_{a,b}$ with the seed on side A , every node on side B is adjacent to the seed and stays regardless of arrival order permutation. For a non-seed node u on side A , there are b internally disjoint paths of length 2, each through a different B -node.

PROPOSITION 6.4 (COMPLETE BIPARTITE FORMULA). *On $K_{a,b}$ with the seed on side A , each non-seed A -node satisfies*

$$\Pr(u \in P^*) = \frac{b}{b+1}.$$

PROOF. Inclusion-exclusion over b paths of length 2 gives $\Pr(u \in P^*) = \sum_{k=1}^b (-1)^{k+1} \binom{b}{k} \cdot 1/(k+1)$, since each subset of k paths has $D_S = k+1$ and $\prod (d_i - 1)! = 1$. Using the integral identity $1/(k+1) = \int_0^1 x^k dx$ and the binomial theorem:

$$\Pr(u \in P^*) = \int_0^1 [1 - (1-x)^b] dx = 1 - \frac{1}{b+1} = \frac{b}{b+1}. \quad \square$$

The expected party size is $\mathbb{E}[|P^*|] = 1 + b + (a-1)b/(b+1)$. As $b \rightarrow \infty$, the reachability probability approaches 1, such that many short disjoint paths accordingly make attendance nearly certain.

Stars. We further consider a network in a star structure; a star $K_{1,n-1}$ has a single center c and $n-1$ leaves. If the center is seeded, every leaf is adjacent to the seed and $\mathbb{E}[|P^*|] = n$. If a leaf ℓ is seeded, the center stays with probability 1 (it is adjacent to ℓ), and each

remaining leaf stays if and only if the center arrives before it, with probability $1/2$. Therefore $\mathbb{E}[|P^*|] = 1 + 1 + (n-2)/2 = (n+2)/2$. The ratio between center-seeded and leaf-seeded outcomes is nearly 2 for large n .

Simulated trajectories. We conclude our exploration of specific graph structures by a simple simulation of the party outcomes. Figure 1 plots the expected party size $\mathbb{E}[|P(t)|]$ as a function of the fraction of people arrived, computed using the time-indexed inclusion-exclusion formula (5) for each graph family. All four graphs have $n = 100$ nodes and a single seed. On the path, the party barely grows beyond the seed's immediate neighbor, finishing at e . The ring improves only slightly, to $2e - 1$. The binary and ternary trees do better (about 7.3 and 15.5, just short of the infinite-depth limits e^2 , about 7.4, and e^3 , about 20.1, because the trees are truncated at $n = 100$). In any case, the party size is just a small fraction of the total network. The trajectory is convex for each case: most of the party's growth comes late, once the last few vertices on a path finally arrive in order.




Figure 1: Expected party size over time from the exact inclusion-exclusion formula, for four graph families with $n = 100$ and $k = 1$. The horizontal axis is the fraction of non-seed vertices that have arrived. All four families produce parties that are small relative to the network.

7 Network Structure and Party Outcomes

The exact results above cover only a few highly structured families. We present an initial exploration of party success on general networks with simulation with a small-world model. We use the Watts-Strogatz small-world model, which interpolates between a ring lattice ($\beta = 0$) and a random graph ($\beta = 1$) as a rewiring probability β increases. Every graph in the family has the same number of nodes ($n = 100$) and the same average degree (6); only the arrangement of the connections between nodes changes as we increase β .

Figure 2 shows the expected party size as a function of β . At $\beta = 0$ (ring lattice), the network is locally tree-like with high clustering (0.60), and the expected party size is about 15 out of 100. As β increases toward 1, clustering drops, and the party size triples to about 48. Note that this increase occurs with no change in average degree.

The mechanism is the same as in our path framework: on a ring lattice, most paths from the seed to distant nodes pass through the same local neighbors, so a node has redundant access to the same




Figure 2: Expected party size and clustering coefficient as a function of Watts–Strogatz rewiring probability β , with $n = 100$ and degree 6 fixed. Each point averages over 80 random graph instances and 2000 arrival permutations per instance. The party size increases monotonically as local structure grows more random.

small region. In this case, rewiring an edge into a random shortcut creates a new path to a distant part of the network. The factorial penalty $1/d!$ over graph distance means that a single short path to a new region generally improves party size much more than many long paths within an existing region.

8 Sending Invites

We last turn to the problem of sending initial invites to the party. Given a graph G and a budget of m invitees, we wish to characterize the seed set S_0 that maximizes the expected party size $f_G(S_0) := \mathbb{E}_\sigma[|P^*(S_0)|]$. As per standard influence maximization results (e.g. Kempe, Kleinberg, and Tardos [4]), the expected influence spread is submodular for all thresholds, so greedy seed selection is at least a $(1 - 1/e)$ -approximation. In our model, we similarly show submodularity when guests decide to stay/leave with a threshold of $k = 1$. However, we then show that submodularity can fail when $k \geq 1$, such that invitees may act as complements rather than substitutes.

8.1 Submodularity for $k = 1$

THEOREM 8.1 ($k = 1$: SUBMODULARITY). *For $k = 1$, the function $f_G(S) = \mathbb{E}_\sigma[|P^*(S)|]$ is monotone submodular.*

PROOF. It suffices to show that for each node v the function $g_v(S) = \Pr(v \in P^*(S))$ is monotone submodular, since $f_G(S) = |S| + \sum_{v \notin S} g_v(S)$ and a nonnegative sum of monotone submodular functions is monotone submodular.

Fix an arrival order σ , and for each node v let

$$B_v(\sigma) = \{s \in V : v \text{ is reachable from } s \text{ in the } \sigma\text{-induced DAG}\}$$

be its set of witnesses (with $v \in B_v(\sigma)$ via the empty path). For $k = 1$ the reachable set is a union over seeds, $\text{Reach}(S, \sigma) = \bigcup_{s \in S} \text{Reach}(s, \sigma)$, so the Reachability Equivalence (Proposition 3.1), applied with seed set S , gives

$$v \in P^*(S) \iff S \cap B_v(\sigma) \neq \emptyset.$$

Hence, for every fixed σ ,

$$|P^*(S)| = \sum_{v \in V} \mathbf{1}\{S \cap B_v(\sigma) \neq \emptyset\},$$

a coverage function of S : each term rewards S for intersecting the target set $B_v(\sigma)$. We check that each term is monotone submodular. Write $h(S) = \mathbf{1}\{S \cap B_v(\sigma) \neq \emptyset\}$ and fix $A \subseteq T$ and $x \notin T$. The marginal $h(A \cup \{x\}) - h(A)$ equals 1 when $x \in B_v(\sigma)$ and $A \cap B_v(\sigma) = \emptyset$, while $h(T \cup \{x\}) - h(T)$ equals 1 when $x \in B_v(\sigma)$ and $T \cap B_v(\sigma) = \emptyset$. Since $A \subseteq T$, the second condition implies the first, so

$$h(A \cup \{x\}) - h(A) \geq h(T \cup \{x\}) - h(T).$$

Thus h is monotone submodular, and so is $|P^*(\cdot)|$, the party size, for every arrival order σ . Finally, $f_G(S) = \mathbb{E}_\sigma[|P^*(S)|]$ is an average of monotone submodular functions and is therefore itself monotone submodular. $\square$

COROLLARY 8.2 (GREEDY GUARANTEE). *For $k = 1$, the greedy algorithm (repeatedly adding the node with the largest marginal increase in f_G) achieves $f_G(S_{\text{greedy}}) \geq (1 - 1/e) \max_{|S|=m} f_G(S)$ for any budget m .*

PROOF. f_G is nonnegative, monotone, and submodular. The result follows from the Nemhauser–Wolsey–Fisher [5]. $\square$

In this setting, each guest has a set of witness hosts – the seeds from which an order-respecting path reaches them – so expected attendance is effectively a coverage objective. Hosts thus act as substitutes when $k = 1$: a host added to a larger seed set can help at most as much as when added to a smaller set.

8.2 Failure of Submodularity for $k \geq 2$

THEOREM 8.3 ($k \geq 2$: SUBMODULARITY FAILS). *For every $k \geq 2$ there is a graph G on which f_G is not submodular. In particular, let G be the star $K_{1,k}$ with center c and leaves $u_1, \dots, u_k$, run under threshold k . Taking $A = \emptyset$, $B = \{u_1, \dots, u_{k-1}\}$, and $x = u_k$,*

$$f_G(A \cup \{x\}) - f_G(A) = 1 < 2 = f_G(B \cup \{x\}) - f_G(B),$$

so f_G violates diminishing-returns.

PROOF. Under threshold k , every leaf has degree $1 < k$, so a non-seed leaf can never reach threshold and never joins. The only node that can be activated is the center c , whose neighbors are exactly the leaves; since non-seed leaves never activate, c joins if and only if all k leaves are seeded. Hence, for any seed set S of leaves,

$$f_G(S) = |S| + \mathbf{1}\{S = \{u_1, \dots, u_k\}\},$$

a value that does not depend on the arrival order. Therefore $f_G(\emptyset) = 0$, $f_G(\{u_k\}) = 1$, $f_G(\{u_1, \dots, u_{k-1}\}) = k - 1$, and $f_G(\{u_1, \dots, u_k\}) = k + 1$, which give the marginals above. $\square$

Intuitively, when the decision-threshold requires more than a single seed, $k \geq 2$, hosts can be complements rather than substitutes. A guest who needs k present friends may be impossible to activate until the last of those friends is seeded, so the marginal value of a host can rise as other hosts are added.

9 Model Extensions

We present a relatively simple model of sequential, irrevocable, and conditional participation under a network-independent arrival order. We impose a number of initial assumptions that could more interestingly be relaxed.

Network-correlated arrivals. A uniform choice of order explicitly manages all guests symmetrically. In practice, arrivals may be correlated with network information. For example, arrival order may plausibly be correlated with node popularity (degree). Give each non-seed vertex v an independent exponential clock $T_v \sim \text{Exp}(\lambda_v)$ and sort arrivals by T_v - if all nodes have equal rates, we recover the uniform order results. On a tree with a single seed, the unique seed-to- v path $s = u_0, u_1, \dots, u_d = v$ is order-respecting exactly when $T_{u_1} < \dots < T_{u_d}$, and given the memoryless property of the arrival clocks

$$\Pr(v \in P^*) = \prod_{i=1}^{d-1} \frac{\lambda_{u_i}}{\lambda_{u_i} + \lambda_{u_{i+1}} + \dots + \lambda_{u_d}}.$$

Again note that setting equal rates collapses this to $1/d!$, recovering Theorem 6.1. With degree-biased rates $\lambda_v \propto \deg(v)$ the product is largest if the path runs from high-degree to low-degree vertices.

Beyond internally disjoint paths. Theorem 5.1 assumes the seed-to- v paths are internally disjoint. When paths overlap their order-respecting events are correlated and simple counting and inclusion-exclusion can no longer be applied directly. Evaluating $\Pr(v \in P^*)$, for example, then amounts to counting linear extensions of the partial order generated by the competing path constraints, and counting linear extensions is #P-complete [1]. However, it is feasible to construct bounds on the inclusion probabilities. The shortest path alone gives the lower bound $\Pr(v \in P^*) \geq 1/d_G(s, v)!$. An upper bound would follow if the order-respecting path events were positively associated; independent-paths for example form $1 - \prod_p (1 - 1/|p|!)$, where p is over all simple s - v paths.

10 Related Work

Threshold models and influence maximization. Granovetter [3] examined threshold models of collective behavior, and Watts [6] related them to network structure. Kempe, Kleinberg, and Tardos [4] showed that expected influence under the linear threshold and independent cascade models is monotone submodular, so greedy seeding gives a $(1 - 1/e)$ -approximation. However, note that their submodularity relies on randomized thresholds (each node draws its threshold uniformly at random), whereas here with a fixed deterministic threshold, submodularity no longer holds if $k \geq 2$. The $k = 1$ case maintains submodularity and reduces to a coverage problem. Additionally, the final active set is not an order-independent monotone closure of the seed set - we instead have irrevocable participation, such that a guest who arrives before enough friends are present leaves and will not return. The final party therefore depends on the arrival order.

Collective action on networks. Chwe [2] studied how network structure supports collective action through common knowledge. Our model shares a similar setting (eg with a fixed network, a seed set, and threshold participation), but instead of agents that coordinate through common knowledge, the agents of this model make irrevocable sequential decisions under a random arrival order. The tree formula $\Pr(v \in P^*) = 1/d!$ could be viewed as an example of the qualitative observation that recruitment (analogous to the probability of a node joining the party) becomes harder with distance.

Random orientations of graphs. The reachability equivalence connects our model to the study of randomly oriented graphs. However, in most of these settings, the edges are oriented *independently*, whereas our edge directions are coupled upon selection of the single shared arrival order.

11 Conclusion

We introduced a model of sequential, irrevocable party attendance. For a one-friend threshold the model is solvable directly, and the final party is defined by the set of guests reachable from the hosts along arrival-increasing paths. Counting paths additionally gives closed-form results for trees, rings, and complete bipartite graphs. We provide accompanying simulations and demonstrate the importance of graph structure by rewiring a fixed-degree lattice in a small-world graph. When $k = 1$, expected attendance is a coverage function of the host set and is submodular. When $k \geq 2$, however, submodularity no longer holds. Of course, this model of sequential participation on networks may extend beyond parties to more interesting applications.

References

- [1] Graham Brightwell and Peter Winkler. 1991. Counting linear extensions. *Order* 8, 3 (1991), 225–242.
- [2] Michael Suk-Young Chwe. 1999. Structure and strategy in collective action. *Amer. J. Sociology* 105, 1 (1999), 128–156.
- [3] Mark Granovetter. 1978. Threshold models of collective behavior. *Amer. J. Sociology* 83, 6 (1978), 1420–1443.
- [4] David Kempe, Jon Kleinberg, and Éva Tardos. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 137–146.
- [5] George L. Nemhauser, Laurence A. Wolsey, and Marshall L. Fisher. 1978. An analysis of approximations for maximizing submodular set functions—I. *Mathematical Programming* 14, 1 (1978), 265–294.
- [6] Duncan J. Watts. 2002. A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences* 99, 9 (2002), 5766–5771.