Wireless Sensor Network Design via Interacting Particles [*]

Aaron B. Wagner and Venkat Anantharam
Dept. of Electrical Engineering and Computer Sciences
University of California
Berkeley, CA 94720 USA
{awagner,ananth}@eecs.berkeley.edu

Abstract

We consider a design problem in which a one-dimensional array of sensors, connected via a wireless network, is used to track a moving vehicle. Individual sensors occasionally lose the vehicle, but can relocate it with aid from neighboring sensors. Communication between sensors is costly, so we accordingly optimize the communication protocol within a class of broadcast-type protocols to maximize the amount of time that the network as a whole tracks the target, subject to a bound on the total rate of communication between sensors. The solution to the optimization problem is heavily influenced by a phase transition exhibited by the network, which we reveal by appealing to existing results from interacting particle systems. We extend these results in the course of solving our optimization problem to address the novel problem of "designing an interacting particle system."

1 Introduction

Recent technological advances have stimulated interest in *ad-hoc wireless sensor networks* [1]. By a sensor network we mean one employed to measure changes in its surroundings, rather than to relay personal messages between its members. Sensor networks that are wireless and ad-hoc have a special appeal in applications where traditional wired- or cellular-based-wireless networks are impractical: locating vehicles on a battlefield or monitoring the diffusion of radiation or harmful gases after an accident or attack. In these examples, the rapidity with which the network must be deployed prohibits the construction of a wired or cellular infrastructure.

These target applications influence the network design. The nodes must be small, so that they can be deployed into an inhospitable environment via plane or other vehicle. The nodes must be numerous, to provide spatial resolution and to make the network resilient against attack. The networks size, in turn, dictates that it be decentralized.

Minute nodes store minute amounts of energy, which places severe limitations on how they can be networked. Thus communication and sensing must be done sparingly to conserve precious energy. Evidently, there is a trade-off between the amount of communication and sensing that individual nodes do.

---

[*] This research was supported by DARPA Grants F30602-00-2-0538 and N66001-00-C-8062, by Grant N00014-1-0637 from the Office of Naval Research, and by Grant SBR-9873086 and a Graduate Research Fellowship from the National Science Foundation.

The other requirement, that there be many nodes in a decentralized network, affects the design in more subtle ways. Sophisticated network-level sensing procedures are envisioned [1], but how to achieve the desired network-level behavior through the decentralized interaction of a large number of simple sensors is unclear. Extrapolating from local to global dynamics is generally a difficult problem [2, 3].

We consider a concrete problem that incorporates both systems-level issues: the interplay between communication and sensing and the translation from local to global dynamics. Consider the following application of an ad-hoc wireless sensor network. In order to track a vehicle moving in the 2-D plane, we drop an array of N radio- or IF-equipped sensors in a line near the vehicle. Each sensor detects a signal emitted by the vehicle and uses it to estimate the vehicle's relative bearing. Ultimately, this information must be gathered and transmitted outside of the network, although we will not address this here.

Instead, we focus on the tracking problem itself. Suppose that occasionally an individual node ceases from tracking the vehicle, either because the signal that the node uses to track has become too noisy, or because the node needs to conserve power. We assume that once this occurs, the sensor is unable to reacquire the signal on its own. We assume, however, that a neighboring node that is tracking the vehicle can communicate the vehicle's position to it, enabling it to resume tracking. We seek to optimize the protocol that the nodes use to share the tracking information.

To this end we employ the following model. The sensor network is a set of N points in a line, each of which is in one of two states, synchronized (tracking the vehicle) or unsynchronized, at each time. We suppose that node i, once synchronized, remains synchronized for an exponentially distributed amount of time with mean 1, independent of the other nodes, and while it is synchronized it broadcasts the bearing of the vehicle at the times of a Poisson process with rate $2\lambda_i$, a design parameter. Each broadcast is equally likely to be made to the left or right, and causes the neighboring node in that direction to become synchronized if it is not already. The use of exponential distributions and Poisson processes makes the process Markov, where the state is the set of synchronized nodes. This process has a single absorbing state that is reachable from all others, namely the state in which all nodes are unsynchronized. Our design problem is to choose the λ_i's to maximize the hitting time to this state subject to an upper bound on their average.

This Markov process is well-known in the interacting particle systems community as the one-dimensional *contact process* on a finite set. It is usually viewed as a model of population growth: an occupied point i contains a member of a population who survives for a random time that is exponentially distributed with mean 1, and who reproduces at the times of a Poisson process with rate $2\lambda_i$, placing each offspring to its left and to its right (if the points are unoccupied) with equal probability. This description makes it clear that the contact process can be viewed as a "branching process with competition"; the progeny compete for space on the lattice. In the context of population growth, we call the hitting time to the all-unoccupied state the *extinction time*.

This model and its infinite-lattice counterparts have been extensively studied [2, 3]. In particular, in the case that all of the λ_i's are equal, the asymptotic rate of growth of the hitting time to the absorbing state as $N \to \infty$ is known. This result is reproduced in the next section, following a more precise definition of the process. It demonstrates a phase transition: the growth of the extinction time with N is qualitatively different for different values of λ. We obviously prefer to operate our network in the phase that causes the fastest growth of the extinction time with N. When we vary the λ-parameter over the process, however, we will observe different parts of the process evolving in different phases. From a design standpoint, then, the contact process presents an usual problem: the system has multiple phases, one of which is most desirable to us, and we would like as much of the process to operate in the preferred phase as the constraints allow.

To clarify this point, we consider the same optimization problem for the simpler biased voter model. The biased voter model has a similar phase transition to the contact process but its simplicity allows us to more clearly expose how the phase transition affects the solution to the optimization problem.

The balance of the paper is organized as follows. The next section contains the required background on the contact process, including the phase-transition result mentioned above. Section 3 contains our optimization results for the contact process, while Section 4 contains our results for the voter model. We conclude in the final section by suggesting a new paradigm for research in interacting particle systems, based on the design problem considered here.

2 Contact Process Background

The finite, one-dimensional contact process with size N and infection rate λ is defined as the Markov chain with state space $2^{\{1,\dots,N\}}$ and transition rates

$$\begin{aligned} q(A, A\backslash\{j\}) &= 1, \quad \text{if } j \in A, \\ q(A, A \cup \{j\}) &= \lambda|A \cap \{j-1, j+1\}|, \text{ if } j \notin A, \end{aligned} \tag{1}$$

for all $A \subset \{1,\dots,N\}$ and $j \in \{1,\dots,N\}$. Here $|\cdot|$ denotes cardinality. For this process and all others described in this paper, if B is not of the form $A\backslash\{j\}$ or $A \cup \{j\}$ for some integer j then $q(A,B) = 0$. One should interpret the state of the chain as the "set of points that are currently occupied (synchronized)." This process is the finite truncation of the contact process on $\mathbb{Z}$, which is the Markov process whose state space is the power set of $\mathbb{Z}$ and whose transition rates are given by (1) for all $A \subset \mathbb{Z}$ and $j \in \mathbb{Z}$. See Liggett [2] for a construction.

As in the finite case, the empty set is a trap for the infinite process. The infinite process differs, however, in that it can survive forever if the reproduction rate λ is sufficiently large. Specifically, the infinite process exhibits the following phase transition. Let ξ_t^A denote the contact process on $\mathbb{Z}$ with initial state $A \subset \mathbb{Z}$. For the infinite process it is known [3] that there exists $\lambda_c \in (0,\infty)$ such that if $\lambda \le \lambda_c$ then $P(\xi_t^{\{0\}} \neq \emptyset \text{ for all } t) = 0$, but if $\lambda > \lambda_c$ then $P(\xi_t^{\{0\}} \neq \emptyset \text{ for all } t) > 0$. Thus the behavior of the infinite process depends crucially on whether $\lambda \le \lambda_c$ or $\lambda > \lambda_c$. If $\lambda \le \lambda_c$, then the process started with a single occupied point becomes extinct with probability one, and we say that the process is *subcritical* if $\lambda < \lambda_c$ and *critical* if $\lambda = \lambda_c$. If $\lambda > \lambda_c$, then the process started with a single occupied point survives forever with positive probability, and we say that the process is *supercritical*. The exact value of λ_c is not known rigorously; simulations place it around 1.6 [4].

Consider now the finite process of size N, ζ_t, with initial state $\{1,\dots,N\}$. Let σ_N be the extinction time of the process, i.e.,

$$\sigma_N = \inf\{t \ge 0 : \zeta_t = \emptyset\}.$$

For the finite process, clearly $\sigma_N < \infty$ a.s. for all λ. But the following result shows that σ_N still depends crucially on how λ compares with λ_c. This is the hitting time result to which we alluded in the introduction; it is due to Durrett and Liu [5] and Durrett and Schonmann [6].

Theorem 1 ([5, 6]) *If $\lambda < \lambda_c$ then as $N \to \infty$,*

$$\frac{\sigma_N}{\log N} \xrightarrow{P} \frac{1}{\gamma_1(\lambda)}.$$

If $\lambda > \lambda_c$ *then as* $N \to \infty$*,*

$$\frac{\log \sigma_N}{N} \xrightarrow{P} \gamma_2(\lambda).$$

The functions γ_1 and γ_2 are defined in terms of the infinite process:

$$\gamma_1(\lambda) = -\lim_{t\to\infty} \frac{1}{t} \log P\left(\xi_t^{\{0\}} \neq \emptyset\right),$$
$$\gamma_2(\lambda) = -\lim_{N\to\infty} \frac{1}{N} \log P(\tau^N < \infty),$$

where $\tau^N = \inf\{t \geq 0 : \xi_t^{\{1,\ldots,N\}} = \emptyset\}$. Both limits exist for all $\lambda \geq 0$ by subadditivity, but $\gamma_1(\lambda)$ is positive if $\lambda < \lambda_c$, while $\gamma_2(\lambda)$ is positive if $\lambda > \lambda_c$. No closed-form expression is known for either function. In words, Theorem 1 says that if $\lambda < \lambda_c$, then σ_N grows logarithmically with N, whereas if $\lambda > \lambda_c$, then σ_N grows exponentially with N. The results for $\lambda = \lambda_c$ are incomplete and are not important here.

Theorem 1 is best understood by thinking of the contact process as a branching process: individuals in the population die and reproduce. In the subcritical case, the tendency to die is stronger than the tendency to reproduce. For large N, the extinction time is essentially the time it takes for all of the original members of the population to die out; reproduction can be neglected. This yields logarithmic growth with N.

In the supercritical case, the tendency to reproduce is stronger than the tendency to die. Thus the population tends to spread, so it becomes extinct only when all points die out virtually simultaneously. This yields exponential growth with N.

3 Optimization Results

We turn to the question of how to vary the reproduction rate in the contact process to maximize the extinction time's asymptotic growth rate as $N \to \infty$. To facilitate taking $N \to \infty$, we shall consider λ-profiles that are *piecewise constant*, defined as follows. We define a *profile* to be a triple $(K, \vec{\alpha}, \vec{\lambda})$, where K is a natural number and $\vec{\alpha}$ and $\vec{\lambda}$ are K-dimensional vectors with nonnegative elements $(\alpha_1, \ldots, \alpha_K)$ and $(\lambda_1, \ldots, \lambda_K)$, such that $\sum_{j=1}^{K} \alpha_j = 1$. We imagine a profile partitioning the unit interval into K partitions, the first of size α_1 with reproduction rate λ_1, the second of size α_2 with reproduction rate λ_2, etc. The profile concept is useful because a profile induces an inhomogeneous contact process for each N. To see how, fix N, let $\beta_j = \sum_{i=1}^{j} \alpha_i$ for all $j = 1, \ldots, K$, define $i_j = \lfloor \beta_j N \rfloor$, and let $i_0 = 0$. We will assume that N is large enough that $i_j < i_{j+1}$ for all $j = 0, \ldots, K-1$. Our inhomogeneous contact process is then the Markov chain with state space $2^{\{1,\ldots,N\}}$ and transition rates

$$\begin{aligned} q(A, A\backslash\{j\}) &= 1, \text{ if } j \in A \\ q(A, A \cup \{j\}) &= \lambda(j-1)|A \cap \{j-1\}| + \\ &\qquad \lambda(j+1)|A \cap \{j+1\}|, \text{ if } j \notin A, \end{aligned} \tag{2}$$

for all $A \subset \{1, \ldots, N\}$ and $j \in \{1, \ldots, N\}$. Here $\lambda(j) = \lambda_m$ if there is an m with $i_{m-1} < j \leq i_m$, otherwise $\lambda(j) = 0$. We call a process with transition rates (2) a *piecewise-homogeneous contact process*.

Before proceeding with the optimization problem, we must determine how the choice of the reproduction rates affects the growth rate of σ_N. The next three theorems generalize Theorem 1 to piecewise-homogeneous contact processes. Consider a fixed profile and let σ_N be the extinction time of the piecewise-homogeneous process of size N induced by this profile.

Theorem 2 *Let $(K, \vec{\alpha}, \vec{\lambda})$ be a profile such that $\lambda_j < \lambda_c$ for all $j \in \{1, \ldots, K\}$. Then as $N \to \infty$,*

$$\frac{\sigma_N}{\log N} \xrightarrow{P} \frac{1}{\gamma_1(\max(\lambda_1, \ldots, \lambda_K))}.$$

The proof of this theorem and the subsequent ones can be found in the full version of this paper [7]. The intuition is that, in the subcritical case, the populations in each of the partitions die without spreading to the neighboring partitions, so the partitions essentially evolve independently. Then σ_N is determined by the extinction times of the partitions with the maximum rate. In the supercritical case, on the other hand, the partitions interact significantly.

Theorem 3 *Let $(K, \vec{\alpha}, \vec{\lambda})$ be a profile such that $\lambda_j > \lambda_c$ for all $j \in \{1, \ldots, K\}$. Then*

$$\frac{\log \sigma_N}{N} \xrightarrow{P} \sum_{j=1}^{K} \alpha_j \gamma_2(\lambda_j) \text{ and } \frac{\log E[\sigma_N]}{N} \to \sum_{j=1}^{K} \alpha_j \gamma_2(\lambda_j)$$

as $N \to \infty$.

The intuition is that if all of the partitions are supercritical, then any one of them, if left alive, will tend to spread to the others. Then the entire process dies only when all partitions die simultaneously. The time for the jth partition, evolving in isolation, to die is $\exp(\alpha_j \gamma_2(\lambda_j) N + o(N))$ by Theorem 1. Thus σ_N is $\exp\left(\sum_{j=1}^{K} \alpha_j \gamma_2(\lambda_j) N + o(N)\right)$.

We call a profile *mixed* if it contains some subcritical and some supercritical partitions. The exponent for mixed profiles is more difficult to determine, so we will only exhibit bounds on it that are tight enough to allow us to proceed with the optimization problem. Let F be the set of indices j such that i_j separates sub- and supercritical partitions,

$$F = \{i \in \{1, \ldots, K-1\} : (\lambda_i \wedge \lambda_{i+1}) < \lambda_c < (\lambda_i \vee \lambda_{i+1})\}.$$

Now $M = |F| + 1$ is the number of “aggregate partitions”—sets of partitions that are connected, entirely subcritical or supercritical, and maximal in that adding another partition either makes the set not unconnected or mixed sub- and supercritical. We denote these aggregate partitions by C_j for $j \in 1, \ldots, M$:

$$\begin{aligned} C_1 &= \{1, \ldots, \inf F \cup \{K\}\} \\ C_j &= \{\sup C_{j-1} + 1, \ldots, \inf\{i \in F : i > \sup C_{j-1}\} \cup \{K\}\}. \end{aligned}$$

Let L be the number of aggregate partitions that are supercritical, so $L = \lceil M/2 \rceil$ if $\lambda_1 > \lambda_c$, otherwise $L = \lfloor M/2 \rfloor$. We call a C_j consisting of supercritical partitions an *island*, and a C_j consisting of subcritical partitions a *sea*. Let D_j for $j \in \{1, \ldots, L\}$ denote the islands, which are the C_j’s with even or odd indices depending on whether $\lambda_1 < \lambda_c$ or $\lambda_1 > \lambda_c$, respectively. Figure 1 shows an example.

Theorem 4 *Let $(K, \vec{\alpha}, \vec{\lambda})$ be a mixed profile. Then*

$$P\left(\frac{\log \sigma_N}{N} < \max_{i \in \{1, \ldots, L\}} \left(\sum_{j \in D_i} \alpha_j \gamma_2(\lambda_j)\right) - \epsilon\right) \to 0,$$

$$P\left(\frac{\log \sigma_N}{N} > \sum_{i=1}^{L} \sum_{j \in D_i} \alpha_j \gamma_2(\lambda_j) + \epsilon\right) \to 0,$$

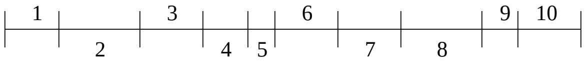

$$C_1 = \{1\},\ C_2 = \{2\},\ C_3 = \{3\},\ C_4 = \{4,5\},\ C_5 = \{6\},\ C_6 = \{7,8\},\ C_7 = \{9,10\},$$
$$D_1 = \{1\},\ D_2 = \{3\},\ D_3 = \{6\},\ D_4 = \{9,10\}.$$

Figure 1: A sample mixed profile. The supercritical partitions have their index placed above the line. The subcritical, below.

for all $\epsilon > 0$ *as* $N \to \infty$, *and*

$$\liminf_{N\to\infty} \frac{\log E[\sigma_N]}{N} \geq \max_{i\in\{1,\dots,L\}} \left(\sum_{j\in D_i} \alpha_j \gamma_2(\lambda_j) \right),$$
$$\limsup_{N\to\infty} \frac{\log E[\sigma_N]}{N} \leq \sum_{i=1}^{L} \sum_{j\in D_i} \alpha_j \gamma_2(\lambda_j).$$

The difficulty is determining when the seas isolate the islands into separate processes. The lower and upper bounds in Theorem 4 correspond to two possible answers to this question, “always” and “never.” If the islands are isolated, the extinction time of the process is the extinction time of its longest-living island, giving an exponent of $\max_{i\in\{1,\dots,L\}} \left(\sum_{j\in D_i} \alpha_j\gamma_2(\lambda_j)\right)$. If the population spreads from one island to another, a process we call *colonizing*, then the process dies only when all of the islands die simultaneously. By the intuition following Theorem 3, this gives an exponent of $\sum_{i=1}^{L}\sum_{j\in D_i} \alpha_j\gamma_2(\lambda_j)$.

We conjecture that the correct answer is “sometimes”: whether a sea prevents two islands from colonizing depends on their sizes and reproduction rates. See the full version of the paper for a statement of our conjecture and support for it [7].

Since $\gamma_2(\lambda) = 0$ if $\lambda \leq \lambda_c$, Theorem 4 asserts that

$$P\left(\frac{\log \sigma_N}{N} > \sum_{j=1}^{K} \alpha_j \gamma_2(\lambda_j) + \epsilon \right) \to 0,$$

and

$$\limsup_{N\to\infty} \frac{\log E[\sigma_N]}{N} \leq \sum_{j=1}^{K} \alpha_j \gamma_2(\lambda_j).$$

Also note that if there is a single island, then Theorem 4 gives the exact exponent.

We return to the optimization problem. We will impose a lower bound on the individual broadcast rates, and an upper bound on their average. That is, we seek a profile that maximizes the asymptotic growth rate of σ_N, subject to the constraint that $\lambda_j \geq \lambda_0$ for all $j = 1, \dots, K$, and $\sum_{j=1}^{K} \alpha_j\lambda_j \leq \lambda_0 + \eta$. Here $\lambda_0 \geq 0$ and $\eta \geq 0$ are the data of the problem. It is clear that as long as $\lambda_0 + \eta > 0$, there are feasible profiles with supercritical partitions: simply make the supercritical partitions sufficiently small. This implies that it is possible to achieve exponential growth of $E[\sigma_N]$ with N. Since the case that $\lambda_0 + \eta = 0$ is not interesting, we shall seek to maximize $\liminf_{N\to\infty} \log E[\sigma_N]/N$. This yields the following optimization problem.

$$\begin{array}{ll} \text{maximize} & \liminf_{N\to\infty} \log(E[\sigma_N])/N \\ \text{over} & K, \vec{\alpha}, \vec{\lambda} \\ \text{subject to} & \sum_{j=1}^{K} \alpha_j\lambda_j \leq \lambda_0 + \eta \\ & \lambda_j \geq \lambda_0 \text{ for all } j \in \{1,\dots,K\}. \end{array} \tag{3}$$

Let $R^*(\lambda_0, \eta)$ denote the supremum of $\liminf \log E[\sigma_N]/N$ over the set of feasible profiles. Theorem 4 shows that for any feasible profile,

$$\liminf_{N\to\infty} \frac{\log E[\sigma_N]}{N} \leq \sum_{j=1}^{K} \alpha_j \gamma_2(\lambda_j),$$

and this exponent can be achieved by rearranging the partitions to yield another feasible profile in which the supercritical region is connected. Thus without loss of generality, we can restrict our attention to profiles with connected supercritical regions, for which the exponents are known. This yields the revised optimization problem

$$\begin{array}{ll} \text{maximize} & \sum_{j=1}^{K} \alpha_j \gamma_2(\lambda_j) \\ \text{over} & K, \vec{\alpha}, \vec{\lambda} \\ \text{subject to} & \sum_{j=1}^{K} \alpha_j \lambda_j \leq \lambda_0 + \eta \\ & \lambda_j \geq \lambda_0 \text{ for all } j \in \{1, \ldots, K\}. \end{array}$$

In this form, it is clear that $R^*(\lambda_0, \eta)$ is simply the *concave hull* of γ_2 over $[\lambda_0, \infty)$ evaluated at $\lambda_0 + \eta$. That is, if we let

$$\hat{\gamma}_2^{\lambda_0}(x) = \sup \left\{ \sum_{j=1}^{n} \alpha_j \gamma_2(\lambda_j) \right\},$$

where the supremum is over $n, \vec{\alpha}$, and $\vec{\lambda}$ such that $\lambda_j \geq \lambda_0$ for all $j = 1, \ldots, n$, and $\sum_{j=1}^{n} \alpha_j \lambda_j = x$, for $x \geq \lambda_0$, then the above observation, combined with Carathéodory's Theorem [8, p. 155] and some continuity arguments that are supplied in the full version of the paper, allow one to prove

Theorem 5 *$R^*(\lambda_0, \eta) = \hat{\gamma}_2^{\lambda_0}(\lambda_0 + \eta)$. Furthermore, $R^*(\lambda_0, \eta)$ is achieved by a profile with two partitions ($K = 2$).*

The intuition behind $K = 2$ is that we would like to make the entire process supercritical, but if $\lambda_0 + \eta < \lambda_c$, then the constraint forbids us from doing so. The optimum growth rate is then obtained by making part of the process supercritical, and leaving the rest at λ_0. To understand this intuition better, suppose hypothetically that γ_2 is concave on (λ_c, ∞). Then if $\lambda_0 > \lambda_c$, so that the points are nominally supercritical, then $\hat{\gamma}_2^{\lambda_0} = \gamma_2$ on (λ_0, ∞) so $R^*(\lambda_0, \eta) = \gamma_2(\lambda_0 + \eta)$, and an optimal profile would be a single partition with rate $\lambda_0 + \eta$. If $\lambda_0 < \lambda_c$, there would exist $\lambda^* > \lambda_c$ such that

$$\hat{\gamma}_2^{\lambda_0}(\lambda) = \begin{cases} \gamma_2(\lambda^*) \frac{\lambda - \lambda_0}{\lambda^* - \lambda_0} & \text{if } \lambda < \lambda^* \\ \gamma_2(\lambda) & \text{if } \lambda \geq \lambda^*. \end{cases}$$

In this case the points are nominally subcritical, so we must actively make them supercritical by increasing their rate. In doing so, we are confronted with a choice of allocating the rate to form a small supercritical region with a very high rate, or a large supercritical region with a relatively small rate. Evidently, λ^* plays a pivotal role in this trade-off: if $\lambda_0 + \eta \geq \lambda^*$, then we would have $\hat{\gamma}_2^{\lambda_0}(\lambda_0 + \eta) = \gamma_2(\lambda_0 + \eta)$, and an optimal profile would be a single partition with rate $\lambda_0 + \eta$. If $\lambda_0 + \eta < \lambda^*$, then the optimal profile would consist of two partitions, one subcritical with rate λ_0 and one supercritical with rate λ^*.

One might expect γ_2 to be concave on (λ_c, ∞) since it is nondecreasing and depends on λ primarily through a comparison to the death rate, which is 1. Thus we expect the effect of increasing λ by Δ to diminish as λ increases. Indeed, γ_2 increases at most logarithmically: if all points in $\{1, \ldots, N\}$ die before reproducing, then $\tau^{\{1,\ldots,N\}} < \infty$ so

$$P\left(\tau^{\{1,\ldots,N\}} < \infty\right) \geq \left(\frac{1}{1+2\lambda}\right)^N,$$

which gives $\gamma_2(\lambda) \leq \log(1+2\lambda)$.

But proving that γ_2 is concave on (λ_c, ∞) is difficult, and in fact it might be false. Drawing upon conjectures from scaling theory, Durrett, Schonmann, and Tanaka [9] conjecture that

$$\lim_{\lambda \downarrow \lambda_c} \frac{\log \gamma_2(\lambda)}{\log(\lambda - \lambda_c)} = \beta$$

for some $\beta > 1$. If this is true, then γ_2 behaves as $(\lambda - \lambda_c)^\beta$ near λ_c, and hence it is not concave on (λ_c, ∞). We are unable to resolve this issue, but note that even if γ_2 is not concave near λ_c, we do not expect it to have more than one inflection point to the right of λ_c, and a single inflection point would not alter the solution to the optimization problem much over the concave case. We conjecture the following.

Conjecture 1 *There exists $\lambda_{c_2} \geq \lambda_c$ such that γ_2 is convex on $[\lambda_c, \lambda_{c_2}]$ and concave on $[\lambda_{c_2}, \infty)$. If $\lambda_0 \geq \lambda_{c_2}$, then $R^*(\lambda_0, \eta) = \gamma_2(\lambda_0 + \eta)$ and $R^*(\lambda_0, \eta)$ is achieved by the profile $(1, 1, \lambda_0 + \eta)$. If $\lambda_0 < \lambda_{c_2}$ then R^* is achieved by a profile with at most two partitions, at most one of whose rate is not λ_0.*

To determine whether the inflection point of γ_2, should it exist, is peculiar to the contact process or common to models of this type, we examine the analogous optimization problem for the simpler biased voter model. The results indicate that the former situation is more likely to be the case.

4 Biased Voter Model Results

The finite, one-dimensional biased voter model with size N and rate λ is the Markov chain whose state space is $2^{\{1,\ldots,N\}}$ and whose transition rates are

$$\begin{aligned} q(A, A\backslash\{j\}) &= |(\mathbb{Z}\backslash A) \cap \{j-1, j+1\}|, \text{ if } j \in A, \\ q(A, A \cup \{j\}) &= \lambda |A \cap \{j-1, j+1\}|, \text{ if } j \notin A, \end{aligned}$$

for $A \subset \{1, \ldots, N\}$ and $j \in \{1, \ldots, N\}$. Each point in $\{1, \ldots, N\}$ is viewed as a person who holds one of two opinions, say 0 and 1; we interpret the state of the process as the set of persons who hold opinion 1. A person with opinion 0 changes to opinion 1 at rate λ times the number of neighbors with opinion 1, as with the contact process. But a person with opinion 1 changes to opinion 0 at rate equal to the number of neighbors with opinion 0, whereas in the contact process the corresponding rate is 1 regardless of the neighbors' state. This process also has a single absorbing state, which is reachable from all other states, namely the one in which all persons have opinion 0.

Due to the symmetry between the two opinions, the biased voter model is much simpler to analyze than the contact process. In particular, note that a point in state 1 with two neighbors in state 1 cannot change states, so if we start the process in state $\{1, \ldots, N\}$, then the set of points in state 1 is always connected, and we only need to keep track of its boundaries. The boundaries, in fact, form independent random walks as long as the process is alive [5], which drift inward at rate 1 and outward at rate λ. This simplifies the analysis and leads to nicer results, such as the following analogue of Theorem 1, due to Durrett and Liu [5]. Consider the biased voter model of size N, ζ_t, with initial state $\{1, \ldots, N\}$, and define σ_N as before.

Theorem 6 ([5]) *If $\lambda < 1$ then as $N \to \infty$,*

$$\frac{\sigma_N}{N} \xrightarrow{P} \frac{1}{2(1-\lambda)}.$$

If $\lambda > 1$ then as $N \to \infty$,

$$\frac{\log \sigma_N}{N} \xrightarrow{P} \log(\lambda).$$

In the subcritical case, the boundaries drift inward at rate $(1 - \lambda)$, and a weak law of large numbers argument shows that σ_N is $N/(2(1-\lambda)) + o(N)$. In the supercritical case, the boundaries drift outward, and hence spend most of their time near 1 and N. But they make excursions toward the center, and the chance that the two boundaries meet in a short time interval is $\lambda^{-N+o(N)}$. Standard arguments then show that σ_N is $\lambda^{N+o(N)}$. Note that for the biased voter model, "λ_c" is 1, and the analogue of γ_2 is $\log$, which is concave.

Our definition of a profile remains the same, but now given the profile $(K, \vec{\alpha}, \vec{\lambda})$, we consider the Markov chain with transition rates

$$\begin{aligned} q(A, A\backslash\{j\}) &= |(\mathbb{Z}\backslash A) \cap \{j-1, j+1\}| \text{ if } j \in A, \\ q(A, A \cup \{j\}) &= \lambda(j-1)|A \cap \{j-1\}| + \\ &\qquad \lambda(j+1)|A \cap \{j+1\}| \text{ if } j \notin A, \end{aligned}$$

for $A \subset \{1, \ldots, N\}$ and $j \in \{1, \ldots, N\}$. The analogue of Theorems 2 and 3 for the biased voter model is given by Theorem 7.

Theorem 7 *If $(K, \vec{\alpha}, \vec{\lambda})$ is a profile such that $\lambda_j < 1$ for all $j \in \{1, \ldots, K\}$, then as $N \to \infty$,*

$$\frac{\sigma_N}{N} \xrightarrow{P} \sum_{j=1}^{K} \frac{\alpha_j}{2(1-\lambda_j)}.$$

If $\lambda_j > 1$ for all $j \in \{1, \ldots, K\}$, then as $N \to \infty$,

$$\frac{\log \sigma_N}{N} \xrightarrow{P} \sum_{j=1}^{K} \alpha_j \log(\lambda_j) \text{ and } \frac{\log E[\sigma_N]}{N} \to \sum_{j=1}^{K} \alpha_j \log(\lambda_j).$$

The intuition behind this result is the same as for the previous one, the only difference being that now the random walks that form the boundaries live in an "inhomogeneous environment"; they drift at different rates in different partitions.

Theorem 8 *Theorem 4 holds for the biased voter model if we replace λ_c with 1 and $\gamma_2(\lambda)$ with* $\log \lambda$.

The solution to the optimization problem for the biased voter model matches the intuition from Section 3.

Theorem 9 *For the biased voter model, if $\lambda_0 > 1$, then $R^*(\lambda_0, \eta) = \log(\lambda_0 + \eta)$ is achieved by the profile $(1, 1, \lambda_0 + \eta)$. If $\lambda_0 < 1$, let λ^* be the unique solution to*

$$1 - \frac{\lambda_0}{\lambda^*} = \log \lambda^*$$

that is greater than 1. If $\lambda_0 + \eta > \lambda^*$, then again $R^*(\lambda_0, \eta) = \log(\lambda_0 + \eta)$. If $\lambda_0 + \eta < \lambda^*$, then $R^*(\lambda_0, \eta) = \eta/\lambda^*$ is achieved by the profile $(2, (\alpha, 1-\alpha), (\lambda^*, \lambda_0))$, where $\alpha = \eta/(\lambda^* - \lambda_0)$.

5 Concluding Remarks

This work illustrates the benefit of employing interacting particle systems to study wireless sensor networks: interacting particle system tools expose behavior like the phase transition in our tracking network that might not be clear otherwise. But the existing interacting particle systems literature is phenomenological; our results on designing a contact process are new. The gap between the existing results and those needed to impact technologies like wireless sensor networks suggests a new paradigm for interacting particle systems research, one that studies models with an engineering eye to answer design- and control-oriented questions.

References

[1] J. M. Kahn, R. H. Katz, and K. S. J. Pister, "Emerging challenges: Mobile networking for smart dust," *J. Commun. and Net.*, vol. 2, no. 3, pp. 188–196, 2000.

[2] T. M. Liggett, *Interacting Particle Systems*, Springer-Verlag, New York, 1985.

[3] T. M. Liggett, *Stochastic Interacting Systems: Contact, Voter and Exclusion Processes*, Springer-Verlag, New York, 1999.

[4] R. Durrett, *Lecture Notes on Particle Systems and Percolation*, Wadsworth, Belmont, 1988.

[5] R. Durrett and X.-F. Liu, "The contact process on a finite set," *Ann. Probab.*, vol. 16, no. 3, pp. 1158–1173, 1988.

[6] R. Durrett and R. H. Schonmann, "The contact process on a finite set II," *Ann. Probab.*, vol. 16, no. 4, pp. 1570–1583, 1988.

[7] A. B. Wagner and V. Anantharam, "Designing a contact process: The piecewise-homogeneous process on a finite set with applications," *Submitted to Ann. Appl. Probab.*, 2002.

[8] R. T. Rockafellar, *Convex Analysis*, Princeton University Press, 1970.

[9] R. Durrett, R. H. Schonmann, and N. I. Tanaka, "The contact process on a finite set III: The critical case," *Ann. Probab.*, vol. 17, pp. 1303–1321, 1989.