2002 Conference on Information Sciences and Systems, Princeton University, March 20–22, 2002

Designing an Interacting Particle System: Optimization of the Piecewise-Homogeneous Contact Process

Aaron B. Wagner and Venkat Anantharam
Department of Electrical Engineering and Computer Sciences
University of California
Berkeley, CA 94720-1770 USA
{awagner,ananth}@eecs.berkeley.edu

Abstract — **We consider an optimization problem centered around a population-growth model similar to the one-dimensional, finite contact process, but in which the reproduction rate varies spatially over the population. The objective is to choose the reproduction-rate-profile to maximize the time until the population is extinct. We first extend existing results about the contact process that describe the asymptotic growth of the extinction time with the population size to processes in which the reproduction rate is not constant but rather piecewise constant. We then formulate the optimization problem and characterize its solution in terms of statistics of the original contact process. We also examine the analogous problem for the simpler biased voter model, and obtain similar but more lucid results.**

I. Introduction

Interacting particle systems have been successfully employed to model such diverse phenomena as magnetism, population growth, and the propagation of information and opinions [1, 2, 3]. The distinguishing features of these stochastic systems are that they model a large number of particles that live on a fixed lattice or graph, and that each particle is assigned a state that changes randomly at rates determined by the state of nearby particles. Here "particles" need not be atoms or molecules, but may represent such macroscopic entities as plants, computers, and people. Examples of interacting particle systems include the well-known Ising model, and also the contact process and voter model.

This mathematical field of study is attracting the attention of engineers due to the emergence of large systems with "locally-defined dynamics" similar to interacting particle systems, such as packet-switched IP networks like the Internet, and more recently, wireless sensor networks. Traditionally, however, research on interacting particle systems has only answered analysis-type questions, such as determining the transient or steady-state behavior of a prescribed system. Our aim is to study these models with an engineering eye to answer design- and control-oriented questions, such as determining which system among a class of interacting particle systems has behavior that is closest to a specification, or what kind of control laws are appropriate for a particular interacting particle system. Answers to questions of this sort are required in order for interacting particle systems to impact the technologies mentioned above. With this general aim in mind, we consider a concrete optimization problem involving one particular interacting particle system, the contact process.

The contact process is a Markov process used to model the spread of a population. We consider the finite, one-dimensional process, in which there are N points in a line, each of which is in one of two states, occupied and vacant, at each time. An occupied point becomes vacant at rate 1, while a vacant point becomes occupied at rate λ times the number of occupied neighbors (0, 1, or 2), where λ, the *reproduction rate*, is a parameter. The idea is that on each occupied point lives 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 2λ, placing each offspring to its left (if the point is vacant) and to its right (if the point is vacant) 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.

The process starts in the state in which every point is occupied. The state in which every point is vacant is absorbing, and it is reachable from every other state. Thus the primary question of interest is how long the process takes to reach this state. In the terminology of populations, the population eventually becomes extinct and we seek the distribution of the time it takes to do so. This analysis question was answered by Durrett and Liu [4] and Durrett and Schonmann [5] in the limit as N tends to infinity. The answer is reproduced in the next section, following a more precise definition of the process.

We consider an ensuing design question. Suppose the presence of the population is desirable and we wish to maximize its extinction time. To accomplish this, we vary λ from point to point. It is intuitively clear, and easy to prove using coupling techniques [2], that increasing the reproduction rate at a point increases the expected extinction time. We consider the problem of maximizing the extinction time subject to an upper bound on the average rate. In practice we expect N to be large, so our precise objective is to maximize the asymptotic rate of growth of the extinction time with N.

This optimization problem has a natural description in terms of the population growth interpretation of the contact process. Consider a population that lives on a one-dimensional lattice, such as a row of bushes. We seek to maximize the longevity of this population by distributing a fixed quantity of fertilizer over the points. What then is the optimal strategy for distributing the fertilizer?

The optimization problem arises in other contexts as well. Consider, for instance, the design of sensor networks. Suppose that in order to track a vehicle moving in the plane, we drop an array of N radio-equipped sensors in a line near it. Each sensor detects a signal emitted by the vehicle and uses it to estimate the vehicle's bearing relative to the sensor. Periodically, the sensor broadcasts this information to a basestation, which uses the information received from all of the sensors to triangulate the position of the vehicle. The nodes broadcast asynchronously.

Occasionally, the signal received by a sensor becomes too noisy for the sensor to make a meaningful estimate of the vehicle's bearing. We assume that once this occurs, the sensor is unable to reacquire the signal on its own. We assume, however, that a broadcast by one of the neighboring nodes contains enough information about the vehicle's position for the node to reacquire the signal and to continue tracking the vehicle. If we assume that a broadcast enables only one of the node's neighbors to reacquire the signal[1], then we can model the randomness using a contact process, where the "occupied points" refer to the nodes that are tracking the vehicle. Eventually, then, the network will reach the state in which every node has lost the signal; the network designer seeks to maximize the time until this happens. Increasing a node's broadcast rate increases the power it consumes while it tracks the vehicle. Networks of this sort are power limited [7], so a rate constraint is natural. We arrive at our optimization problem.

The most interesting feature of the homogeneous contact process is its phase transition: we will see that the extinction time is highly nonlinear in λ. There exists a constant, λ_c, such that the process survives much longer if $\lambda > \lambda_c$ than if $\lambda \leq \lambda_c$. In our case, since we vary the λ-parameter over the process, we will see that different parts of the process can evolve in different phases. From a design standpoint, then, the contact process presents an unusual 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 finite 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. In the biased voter model, each point in $\{1, \ldots, N\}$ is viewed as a person who holds one of two opinions, say 0 and 1. A person with opinion 0 changes to opinion 1 at rate λ times the number of neighbors with opinion 1, as in 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. When determining the rates at the boundary points, 1 and N, we use the convention that persons 0 and $N+1$ always have opinion 0. This process also has a single absorbing state, which is reachable from all other states, namely the state in which persons 1 through N all have opinion 0. We seek to maximize the hitting time to this state by varying λ from point to point, as with the contact process.

Although we do not consider it here, the Ising model of magnetism also lends itself to an optimization problem of this kind. Consider fixed volumes of N magnetic materials, with varying magnetic strengths. These materials are combined in some way, then magnetized to store one bit of information. If the total volume of the materials is one, how should one arrange the materials within, say, the unit cube to maximize the time until the magnetization is lost?

The effect of design on two-dimensional site percolation has been studied by Carlson and Doyle [8, 9] in the context of power laws in complex systems. Robert, Carlson, and Doyle [10] consider, in the same context, the effect of design on a simple epidemic model in which infection spreads between three cells. Booth *et al.* [11] consider the effect of design on two-dimensional continuum percolation in the context of the connectivity of wireless networks.

The remainder of the paper is organized as follows. The next section contains the required background on the contact process, including the result due to Durrett and Liu [4] and Durrett and Schonmann [5] mentioned above. Section III contains our contact process results, and Section IV contains our voter model results. Some concluding remarks are made in Section V.

II. Background

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

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

for $A \subset \{1,\ldots,N\}$ and $j \in \{1,\ldots,N\}$. Here $|\cdot|$ denotes cardinality. For this process and all others defined in this paper, if $B \subset \{1,\ldots,N\}$ with $B \neq A$ 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." This process is the finite analogue of the contact process on $\mathbb{Z}$, which is the Markov process ξ_t whose state space is the power set of $\mathbb{Z}$ and whose transition rates are given by (1) for $A \subset \mathbb{Z}$ and $j \in \mathbb{Z}$. It turns out that there exists a unique Markov process with these rates, although this is not clear *a priori*. See Liggett [1] for a complete account.

As in the finite case, the empty set is a trap for the infinite process. The process on $\mathbb{Z}$ differs, however, in that it can survive forever if the reproduction rate is sufficiently large. Specifically, the process exhibits the following phase transition [2]. Let ξ_t^A denote the contact process on $\mathbb{Z}$ with initial state $A \subset \mathbb{Z}$. There exists $\lambda_c \in (0,\infty)$ such that if $\lambda \leq \lambda_c$ 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 crucially depends on whether $\lambda \leq \lambda_c$ or $\lambda > \lambda_c$. If $\lambda \leq \lambda_c$, 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, but simulations place it around 1.6 [3]. Loosely speaking, in the supercritical case, the population tends to spread because the tendency of its members to reproduce is stronger than their tendency to die. In the subcritical case, however, the population tends to contract because the tendency to reproduce is weaker than the tendency to die. See [1, p. 283] for a mathematical formulation of this statement.

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

$$\sigma_N = \inf\{t \geq 0 : \zeta_t = \emptyset\}. \tag{2}$$

The result due to Durrett and Liu [4] and Durrett and Schonmann [5], to which we alluded in the introduction, shows that the phase transition also appears in the finite process. Here and throughout, all logarithms have base e.

Theorem 1 ([4, 5]) *If $\lambda < \lambda_c$ then as $N \to \infty$,*

$$\frac{\sigma_N}{\log N} \to \frac{1}{\gamma_1(\lambda)}$$

in probability. If $\lambda > \lambda_c$ then as $N \to \infty$,

$$\frac{\log \sigma_N}{N} \to \gamma_2(\lambda)$$

in probability.

Here γ_1 and γ_2 are deterministic functions of λ, which are defined using the infinite process.

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

where $\tau^N = \inf\{t \geq 0 : \xi_t^{\{1,\ldots,N\}} = \emptyset\}$. Both limits exist by subadditivity for all $\lambda \geq 0$, but $\gamma_1(\lambda)$ is positive if $\lambda < \lambda_c$, while $\gamma_2(\lambda)$

[1]This would be the case if they used directed antennae.

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. When $\lambda = \lambda_c$, Durrett, Schonmann, and Tanaka [6] show that σ_N grows polynomially with N, but the correct power is unknown. We do not consider the critical process here.

The intuition behind Theorem 1 is the following. In the subcritical case, the tendency to reproduce is weaker than the tendency to die. Then 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. That this gives logarithmic growth can be seen by inspecting the case $\lambda = 0$. Here σ_N is just the maximum of N i.i.d. exponential random variables with mean one, so $\log \sigma_N / \log N \to 1$ in probability.

In the supercritical case, the population tends to spread, so it becomes extinct only when all of the points die out simultaneously. The time that one must wait for N events to occur simultaneously during parallel independent trails grows exponentially with N.

III. Optimization Results

We turn to the question posed in the introduction, namely how to vary the reproduction rate of the contact process to maximize the extinction time's asymptotic growth rate. 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, λ, α), where K is a natural number and λ and α are K-dimensional vectors with nonnegative elements $(\lambda_1, \ldots, \lambda_K)$, and $(\alpha_1, \ldots, \alpha_K)$ such that $\sum_{j=1}^{K} \alpha_j = 1$. We think of a profile as 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 $j = 1, \ldots, K$, 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 $j = 0, \ldots, K-1$. Our inhomogeneous contact process is then the Markov chain whose state space is $2^{\{1,\ldots,N\}}$, and whose transition rates are

$$\begin{aligned} q(A, A\backslash\{j\}) &= 1, \text{ if } j \in A \\ q(A, A \cup \{j\}) &= \lambda(j-1)|A \cap \{j-1\}| + \\ &\qquad \lambda(j+1)|A \cap \{j+1\}|, \text{ if } j \notin A, \end{aligned} \tag{3}$$

for $A \subset \{1, \ldots, N\}$ and $j \in \{1, \ldots, N\}$. Here $\lambda(j) = \lambda_m$ where m satisfies $i_{m-1} < j \le i_m$ if $j \in \{1, \ldots, N\}$, and $\lambda(0) = \lambda(N+1) = 0$. We term a process with transition rates given by (3) a *piecewise-homogeneous contact process*. We call a profile in which $\lambda_j < \lambda_c$ for all $j \in \{1, \ldots, K\}$ a *subcritical* profile, one in which $\lambda_j > \lambda_c$ for all $j \in \{1, \ldots, K\}$ a *supercritical* profile, and one in which $\lambda_j \neq \lambda_c$ for all $j \in \{1, \ldots, K\}$ and $\lambda_j > \lambda_c$ for some $j \in \{1, \ldots, K\}$ a *mixed* profile.

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, λ, α) be a subcritical profile. Then*

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

in probability as $N \to \infty$.

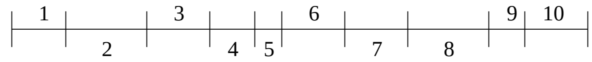

$$C_1 = \{1\}, C_2 = \{2\}, C_3 = \{3\}, C_4 = \{4, 5\}, C_5 = \{6\},$$
$$C_6 = \{7, 8\}, C_7 = \{9, 10\},$$
$$D_1 = \{1\}, D_2 = \{3\}, D_3 = \{6\}, D_4 = \{9, 10\}$$

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

The proof of this theorem and the subsequent ones can be found in the full version of this paper [12]. The intuition is that, in the subcritical case, the populations in each of the partitions tend to die out without spreading to the neighboring partitions, so the partitions essentially evolve independently, and σ_N is determined by the extinction times of the partitions with the maximum rate. In the supercritical case, the partitions interact significantly.

Theorem 3 *Let (K, λ, α) be a supercritical profile. Then*

$$\frac{\log \sigma_N}{N} \to \sum_{j=1}^{K} \alpha_j \gamma_2(\lambda_j)$$

in probability as $N \to \infty$, and

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

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

The exponent for mixed profiles is more difficult to determine, so we will only exhibit bounds on it that are tight enough to enable 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 contiguous, entirely subcritical or supercritical, and maximal in that adding another partition either makes the set not contiguous 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, λ, α) be a mixed profile. Then*

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

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

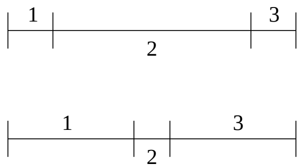

Figure 2: Two three-partition mixed profiles.

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

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

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

The difficulty leading to the incompleteness of this result 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 $\max_{i\in\{1,\ldots,L\}} \left(\sum_{j\in D_i} \alpha_j\gamma_2(\lambda_j) \right)$. If the population can spread from one island to another, a process we call *colonizing*, then the process dies out only when all of the islands die out simultaneously. By the intuition following Theorem 3, this gives an exponent of $\sum_{i=1}^{L} \left(\sum_{j\in D_i} \alpha_j\gamma_2(\lambda_j) \right)$.

We conjecture that the correct answer is "sometimes"; whether a sea prevents two islands from colonizing depends on their sizes and reproduction rates. Consider the two profiles shown in Figure 2. In the first case, a wide subcritical partition separates two small supercritical ones. Our conjecture is that if all of the rates are sufficiently small, then the two supercritical partitions will not be able to colonize and σ_N will have exponent $\max(\alpha_1\gamma_2(\lambda_1), \alpha_3\gamma_2(\lambda_3))$. In the second case, a narrow subcritical partition separates two large supercritical ones. Our conjecture is that if all of the rates are sufficiently large, then the supercritical partitions will colonize, and σ_N will have exponent $\alpha_1\gamma_2(\lambda_1) + \alpha_3\gamma_2(\lambda_3)$. The full version of this paper contains a precise criterion for which of the two scenarios we expect to happen, shows how to extend the conjecture to more than three partitions, and contains some support for the conjecture.

We also show in the full version of this paper that $\gamma_2(\lambda) = 0$ if $\lambda \leq \lambda_c$. Thus 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

$$\limsup_{N\to\infty} \frac{\log E[\sigma_N]}{N} \leq \sum_{j=1}^{K} \alpha_j\gamma_2(\lambda_j). \tag{4}$$

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

Theorems 2 and 4 together show that the piecewise-homogeneous process exhibits a similar phase transition to the homogeneous process, except that the relevant parameter is now the maximum of the K rates. That is, if all K rates are subcritical, then the extinction time grows logarithmically with N, while if any of the rates are supercritical, then the extinction time grows exponentially with N. In the context of our optimization problem, this shows that we can achieve exponential growth of σ_N with N by making part of the process supercritical. This is significant since the constraint may prohibit us from making the entire process supercritical. Specifically, let $\lambda_0 \geq 0$ and $\eta \geq 0$ be given. We interpret λ_0 as the nominal reproduction rate given to each point, and η as the additional reproduction rate that we may distribute over the population as desired. We seek a profile that maximizes the asymptotic growth rate of σ_N, subject to the constraint that $\lambda_j \geq \lambda_0$ for $j = 1, \ldots, K$, and

$$\sum_{j=1}^{K} \alpha_j\lambda_j \leq \lambda_0 + \eta. \tag{5}$$

If $\lambda_0 + \eta > 0$, there are feasible profiles with supercritical partitions: simply make the supercritical partitions sufficiently small. By Theorem 4, it is then 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 gives the following optimization problem.

$$\begin{array}{ll} \text{maximize} & \liminf_{N\to\infty}(\log E[\sigma_N])/N \\ \text{over} & K, \lambda, \alpha \\ \text{subject to} & \sum_{j=1}^{K} \alpha_j\lambda_j \leq \lambda_0 + \eta \\ & \lambda_j \geq \lambda_0 \text{ for } j \in \{1, \ldots, K\}. \end{array} \tag{6}$$

Let $R^*(\lambda_0, \eta)$ denote the supremum of $\liminf(\log E[\sigma_N])/N$ over the set of feasible profiles. Let $\hat{\gamma}_2^{\lambda_0}(\cdot)$ denote the concave hull of $\gamma_2(\cdot)$ on $[\lambda_0, \infty)$, i.e., for $x \geq \lambda_0$,

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

where the supremum is over n, α, and λ such that $\lambda_j \geq \lambda_0$ for all $j = 1, \ldots, n$, and

$$\sum_{j=1}^{n} \alpha_j\lambda_j = x.$$

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* $K = 2$*.*

Here is the idea behind the proof. By (4),

$$\liminf_{N\to\infty}(\log E[\sigma_N])/N \leq \sum_{j=1}^{K} \alpha_j\gamma_2(\lambda_j) \leq \hat{\gamma}_2^{\lambda_0}(\lambda_0 + \eta),$$

which shows that $R^*(\lambda_0, \eta) \leq \hat{\gamma}_2^{\lambda_0}(\lambda_0 + \eta)$. Carathéodory's Theorem [13, p. 155] and some continuity arguments that are supplied in the full version of this paper show that there exists α_1, α_2, λ_1, λ_2, such that $\min(\lambda_1, \lambda_2) \geq \lambda_0$,

$$\alpha_1\lambda_1 + \alpha_2\lambda_2 \leq \lambda_0 + \eta,$$

and

$$\alpha_1\gamma_2(\lambda_1) + \alpha_2\gamma_2(\lambda_2) = \hat{\gamma}_2^{\lambda_0}(\lambda_0 + \eta).$$

Since a profile with $K = 2$ can have at most one supercritical island, Theorem 4 shows that for this profile,

$$\liminf_{N\to\infty}(\log E[\sigma_N])/N = \alpha_1\gamma_2(\lambda_1) + \alpha_2\gamma_2(\lambda_2) = \hat{\gamma}_2^{\lambda_0}(\lambda_0 + \eta),$$

which shows that $R^*(\lambda_0, \eta) \geq \hat{\gamma}_2^{\lambda_0}(\lambda_0 + \eta)$ and that the optimum exponent is achieved by a profile with two rates.

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 obtained, then, 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

$$\hat{\gamma}_2^{\lambda_0}(\lambda) = \begin{cases} \gamma_2(\lambda^*)\frac{\lambda - \lambda_0}{\lambda^* - \lambda_0} & \text{if } \lambda < \lambda^* \\ \gamma_2(\lambda) & \text{if } \lambda \geq \lambda^*. \end{cases} \tag{7}$$

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 that γ_2 is concave on (λ_c, ∞) since it is nondecreasing and $O(\log \lambda)$: if all points in $\{1, \ldots, N\}$ die out before reproducing, then $\tau^{\{1,\ldots,N\}} < \infty$ so

$$P\left(\tau^{\{1,\ldots,N\}} < \infty\right) \geq \left(\frac{1}{1 + 2\lambda}\right)^N, \tag{8}$$

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 [6] conjecture that

$$\lim_{\lambda \downarrow \lambda_c} \frac{\gamma_2(\lambda)}{(\lambda - \lambda_c)^\alpha}$$

exists for some $\alpha > 1$. If this is true, then γ_2 behaves as $(\lambda - \lambda_c)^\alpha$ near λ_c, and hence it is not concave on (λ_c, ∞), since it is convex near λ_c. Of course, it cannot be convex on the entire interval (λ_c, ∞) due to (8). Since we are unable to resolve this conjecture, we cannot assert that the optimal exponent is achieved by profiles with at most one supercritical partition. To validate our intuition, therefore, we turn to the biased voter model.

IV. Biased Voter Model Results

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

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

for $A \subset \{1, \ldots, N\}$ and $j \in \{1, \ldots, N\}$. Here one should interpret the state of the process as the set of persons that currently hold opinion 1. 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 contiguous, and we only need to keep track of the boundaries of this set. The boundaries, in fact, form independent random walks as long as the process is alive [4], and drift inward at rate 1 and outward at rate λ. This fact simplifies the analysis and leads to nicer results, the first of which is the following analogue of Theorem 1, due to Durrett and Liu [4]. Consider the biased voter model of size N, ζ_t, with initial state $\{1, \ldots, N\}$, and define σ_N as before by (2).

Theorem 6 ([4]) *If $\lambda < 1$ then as $N \to \infty$,*

$$\frac{\sigma_N}{N} \to \frac{1}{2(1 - \lambda)}$$

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

$$\frac{\log \sigma_N}{N} \to \log(\lambda)$$

in probability.

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 simply 1, and unlike the contact process, the two limiting functions have simple expressions.

Our definition of a profile remains the same, but now given the profile (K, λ, α), we consider the Markov chain with transitions rates

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

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, λ, α) is a profile such that $\lambda_j < 1$ for all $j \in \{1, \ldots, K\}$, then*

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

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

$$\frac{\log \sigma_N}{N} \to \sum_{j=1}^{K} \alpha_j \log(\lambda_j)$$

in probability as $N \to \infty$, 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 is 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 $\log \lambda$.*

Determining the correct exponents for mixed sub- and supercritical profiles should be relatively easy for the biased voter model. We do not explore this here, however, because our interest in the biased voter model is its solution to the optimization problem, (6).

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 $(1, \lambda_0 + \eta, 1)$. 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 $(2, (\lambda^*, \lambda_0), (\alpha, 1 - \alpha))$, where $\alpha = \eta/(\lambda^* - \lambda_0)$.

The proof is similar to the proof of Theorem 5 for the contact process. Using the approach that we sketched after that theorem, one can show that if $f^{\lambda_0}(\lambda)$ is the concave hull of $\log$ on $[\lambda_0, \infty)$, then $R^*(\lambda_0, \eta) = f^{\lambda_0}(\lambda_0 + \eta)$ and is achieved by a profile with $K = 2$. The remainder of the result follows by observing that the expressions given for the optimal exponent are simply expressions for f^{λ_0}, and the optimal profiles are the ones that achieve the hull.

The intuition behind this result is the same as that following (7). If $\lambda_0 > 1$, or $\lambda_0 < 1$ but $\lambda_0 + \eta$ is large, then the concavity of $\log$ implies that the additional rate has the greatest effect when it is spread uniformly over the population. If $\lambda_0 < 1$ and $\lambda_0 + \eta$ is relatively small, then the trade-off between the size and rate of the supercritical partition becomes significant, and the optimal profile is to split the population into sub- and supercritical parts, and give the subcritical part rate λ_0, and the supercritical part rate λ^*.

V. Concluding Remarks

We considered a natural optimization problem based on the contact process as a first step toward a more general understanding of how to design and control interacting particle systems. This understanding is absent from the interacting particle systems literature, but it is required to put the theory to use in the design of practical systems. Our optimization problem highlighted how a phase transition, a common feature of many interacting particle systems, affects the design process.

Many challenges remain. The theoretical challenges are to answer similar questions for more complicated models, such as the contact process in higher dimensions and the Ising model, and to address closed-loop control questions in place of the open-loop design problem considered here. A simple follow-up question to the one asked here is "which profile minimizes the growth rate of the extinction time?" If the conjecture that a sufficiently large subcritical partition between two supercritical partitions isolates them into independent processes is correct, then minimizing the hitting time would likely involve placing large amounts of rate in small regions that are separated by subcritical "firebreaks."

The practical challenges are related to modeling. One must determine if interacting particle-type models are appropriate for physical systems such as sensor networks and the Internet, and if so what the exact model should be. The potential benefit to this endeavor is the opportunity to use the rich set of tools that have been developed to study interacting particles to study these physical systems.

Acknowledgments

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 authors also wish to thank R. Durrett and R. Schonmann for helpful electronic mail correspondence.

References

[1] T. M. Liggett, *Interacting Particle Systems*, Springer, New York, 1985.

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

[3] R. Durrett, *Lecture Notes on Particle Systems and Percolation*, Wadsworth, Pacific Grove, 1988.

[4] R. Durrett and X.-F. Liu, "The Contact Process on a Finite Set," *The Annals of Probability*, vol. 16, no. 3, 1158–1173, 1988.

[5] R. Durrett and R. H. Schonmann, "The Contact Process on a Finite Set II," *The Annals of Probability*, vol. 16, no. 4, 1570–1583, 1988.

[6] R. Durrett, R. H. Schonmann, and N. I. Tanaka, "The Contact Process on a Finite Set III. The Critical Case," *The Annals of Probability*, vol. 17, no. 4, 1303–1321, 1989.

[7] J. M. Kahn, R. H. Katz, and K. S. J. Pister, "Emerging Challenges: Mobile Networking for Smart Dust," *J. of Commun. and Networks*, vol. 2, no. 3, 188–196, 2000.

[8] J. M. Carlson and J. Doyle, "Highly Optimized Tolerance: A Mechanism for Power Laws in Designed Systems," *Physical Review E*, vol. 60, no. 2, 1412–27, 1999.

[9] J. M. Carlson and J. Doyle, "Highly Optimized Tolerance: Robustness and Design in Complex Systems," *Physical Review Letters*, vol. 84, no. 11, 2529–32, 2000.

[10] C. Robert, J. M. Carlson, and J. Doyle, "Highly Optimized Tolerance in Epidemic Models Incorporating Local Optimization and Regrowth," *Physical Review E*, vol. 63, no. 5, 056122, 2001.

[11] L. Booth, J. Bruck, M. Franceschetti, and R. Meester, "Covering Algorithms, Continuum Percolation, and the Geometry of Wireless Networks," *Preprint.*

[12] A. B. Wagner and V. Anantharam, "Designing a Contact Process: The Piecewise-Homogeneous Process on a Finite Set with Applications," *In Preparation.*

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