

Discrete Interventions in Hawkes Processes with Applications in Invasive Species Management

Amrita Gupta¹, Mehrdad Farajtabar¹, Bistra Dilkina² and Hongyuan Zha¹

¹ School of Computational Science & Engineering, Georgia Institute of Technology

² Department of Computer Science, University of Southern California

agupta375@gatech.edu, mehrdad@gatech.edu, dilkina@usc.edu, zha@cc.gatech.edu

Abstract

The spread of invasive species to new areas threatens the stability of ecosystems and causes major economic losses. We propose a novel approach to minimize the spread of an invasive species given a limited intervention budget. We first model invasive species spread using Hawkes processes, and then derive closed-form expressions for characterizing the effect of an intervention action on the invasion process. We use this to obtain an optimal intervention plan based on an integer programming formulation, and compare the optimal plan against several ecologically-motivated heuristic strategies used in practice. We present an empirical study of two variants of the invasive control problem: minimizing the final rate of invasions, and minimizing the number of invasions at the end of a given time horizon. The optimized intervention achieves nearly the same level of control that would be attained by completely eradicating the species, but at only a fraction of the cost.

1 Introduction

Network diffusion models are a powerful tool for studying processes like the spread of influence and information through social networks [Kempe *et al.*, 2003; Yang and Leskovec, 2010], the dispersal of species through a landscape [Sheldon *et al.*, 2010], and disease contagion in populations [Eames and Keeling, 2002]. The ability to model the dynamics of these diffusion processes enables the development of strategies for steering them towards desirable outcomes. For instance, one might selectively add nodes to an existing network to facilitate diffusion [Sheldon *et al.*, 2010], or one can strategically block transmission along a set of links [Kimura *et al.*, 2008; Khalil *et al.*, 2014] to limit it.

Two of the most studied network diffusion models are the independent cascade (IC) model and the linear threshold (LT) model [Kempe *et al.*, 2003]. In both, the spreading process is modeled as an activation of nodes over discrete time steps. Each node in the network is in a binary state (active or not), and nodes are activated by their active neighbors. In both the IC and LT models, once a node is active it remains so for the rest of the diffusion process, an assumption that is appropriate

for modeling the spread of irreversible phenomena, e.g. the adoption of a product or infection by a disease that confers permanent immunity.

However, many network diffusion processes exhibit *non-progressive* cascades where an active node can become inactive probabilistically at each time step, so that the state of a node fluctuates over time. For example, in species dispersal, a previously occupied habitat patch may become unoccupied [Sheldon *et al.*, 2010], or in infectious disease a patient may recover but be susceptible to reinfection. In this setting, repeated exposure to activation events plays an important role in continuing the diffusion process by reactivating nodes that have become inactive. Sometimes, exposure to multiple activations can also cause a node to become “more” active, e.g. the posting frequency of an individual social media user can increase due to high activity in their network. In these cases, it is more fitting to model the state of a node as a time-varying, real- or continuous-valued function rather than as binary states. Moreover, activation events typically arrive continuously rather than in discrete time steps, warranting the diffusion process to be modeled in continuous time. Other network models like Susceptible-Infected-Recovered (SIR) can encompass some of these characteristics and have been shown to have a close connection to point process models [Rizoio *et al.*, 2018], which are more flexible and intuitive and easily allow the node activities to be modeled according to observed/assumed principles (e.g. mutual-excitation).

Temporal point processes offer a framework for modeling diffusion processes with both continuous activity states and continuous time. The activity of a node can be characterized by a parameter λ representing the rate at which the node stochastically generates events. This λ parameter itself can be responsive to activations arriving at the node, thereby capturing self-exciting behavior in the diffusion process. Temporal point processes have recently been used to model several diffusion processes like the activity of Twitter users [Farajtabar *et al.*, 2014], the spread of avian flu [Kim, 2011], and failures in water pipe networks [Lin *et al.*, 2016]. In the context of invasive species, [Balderama *et al.*, 2012] use a spatiotemporal point process model to characterize the spread of an invasive banana plant, although they do not consider any control.

In terms of controlling diffusion processes, a variety of intervention actions have been analyzed in the discrete-time, binary-state setting, such as selecting source nodes for initi-

ating cascades [Kempe *et al.*, 2003] and modifying network connectivity to guide the diffusion by adding or removing nodes [Sheldon *et al.*, 2010] and edges [Kimura *et al.*, 2008; Khalil *et al.*, 2014] or modifying edge weights [Wu *et al.*, 2015]. In a recent paper, [Bjorck *et al.*, 2018] modeled controlling a diffusion process representing the spread of an invasive prey species by selecting source nodes for a second predator cascade. In contrast, there has been relatively little work on controlling dynamics in network temporal point processes. One possible control action is to manipulate the activity rate parameters λ at specific nodes, e.g. by incentivizing social media users to post more frequently. Steering user activity in this manner was first considered in [Farajtabar *et al.*, 2014], and was used to develop a multistage strategy for mitigating fake news [Farajtabar *et al.*, 2016]. Recent work has also applied methods from stochastic differential equations to find the best intensity for information guiding [Wang *et al.*, 2016] and achieving highest visibility [Zarezade *et al.*, 2017]. In our work, a *discrete* intervention for network point processes is considered for the first time that, unlike the above, modifies the activity rate parameter at select nodes by *deleting the history* of the point process.

Our work is motivated by the invasive species management problem in biodiversity conservation. The spread of non-native species to new areas is a cause of major concern, because they harm native species by disrupting food webs and ecosystem processes. These adverse effects have generated significant interest in limiting their spread. Management is often performed by eradicating invasive species individuals, but their removal can be prohibitively costly. In light of this, a common objective is to optimize the location of control efforts in order to maximize the efficacy of the intervention, e.g. using discrete optimization [Büyüktaktakın *et al.*, 2014]. We derive a novel approach for finding an optimal set of locations at which to remove species given a fixed budget. Our method employs a network temporal point process model that allows us to derive expressions amenable to exact optimization. This is a departure from approaches relying on optimization with computationally intensive simulators [Taleghan *et al.*, 2015] or sample average approximation. Although our work is motivated by a critical problem in environmental sustainability, the novel computational problem it poses appears in other domains that can be modeled using temporal point processes. One such domain is mitigating the spread of infectious diseases using vaccination or active screening programs. Our approach could also be used to model the effect of deleting social media posts on the spread of news through social networks. The computational approach we develop here can be generalized to these broader applications.

2 Invasives Management & Hawkes Processes

2.1 Problem Statement

In the invasive species management problem, the goal is to identify locations at which to eradicate invasive individuals in order to minimize the spread of the species through the landscape. Let L be a set of n distinct land parcels corresponding to basic units of management. An invasive species is observed to be proliferating and dispersing through the landscape until

a given time τ , when an intervention is performed by eliminating all invasive individuals present before τ in a set of land units $U \subseteq L$. Each land unit $i \in L$ has an associated cost c_i reflecting economic land management costs or effort needed to eradicate the invasive individuals, and the total cost of the intervention cannot exceed a given budget \mathcal{B} . A feasible intervention plan is therefore a set of land units U with total intervention cost within \mathcal{B} . After the intervention, the invasive species continues to spread until time $T > \tau$, but without the proliferative influence of the individuals eradicated at time τ . Our goal is to find a feasible intervention plan that minimizes the degree to which the landscape is affected by the invasion.

Invasives Spread Dynamics

Traditionally, the spread of invasive species is modeled using a combination of differential equations describing the population dynamics and a dispersal kernel describing the displacement of individuals [Hastings *et al.*, 2005; Büyüktaktakın *et al.*, 2014], or their stochastic counterparts [Fitzpatrick *et al.*, 2012]. More recently, studies have demonstrated the potential of point process models [Balderama *et al.*, 2012; Pauchard *et al.*, 2016] for characterizing the spatial distribution of invasive plant species.

We present a multidimensional Hawkes process model for the dynamics of invasive species spread. To formulate it as a network diffusion process, we model the n land units in landscape L as nodes V in a directed graph, with edges between nodes that are close enough for dispersal to occur between the corresponding land units. The appearance of a new invasive individual in node i at time s is denoted by the invasion event (i, s) . Indexing invasion events by e , the history of the network diffusion process up to immediately before some time t is the set of events $\mathcal{H}_{t-} := \{(i_e, s_e) | s_e < t\}$.

Invasive species can be introduced at any time by carriers like wind, animals or humans. These arrivals are called *exogenous* invasions, and their rate can vary spatially depending on landscape features or human activity. The instantaneous rate at which individuals are introduced to node i at time t is denoted by $\mu_i(t)$, and represents the probability of an exogenous invasion event in an infinitesimally small time window $[t, t + dt)$. Once an invasive individual has become established, it survives for an average lifetime β . Since many invasive species mature early and have short life expectancy [Sakai *et al.*, 2001], we assume an individual born at s_e faces a constant risk of death $\omega = \frac{1}{\beta}$, so that the probability of the individual surviving until time t is given by the survival function $e^{-\omega(t-s_e)}$. While the individual survives, it initiates *endogenous* invasions, e.g. by releasing offspring. The likelihood of an individual appearing at location i due to the dispersal of the offspring of an individual at location j depends on edge weight a_{ij} between the two nodes, which can be, e.g., a decaying function of distance [Arim *et al.*, 2006].

All these effects together influence the rate at which new individuals appear in a given node i at time t , or the intensity $\lambda_i(t)$. This represents the conditional probability of observing an invasion event in a small time window $[t, t + dt)$ given the history \mathcal{H}_{t-} .

$$\lambda_i(t) = \mu_i(t) + \sum_{(j_e, s_e) \in \mathcal{H}_{t-}} a_{ij_e} \cdot e^{-\omega(t-s_e)} \quad (1)$$

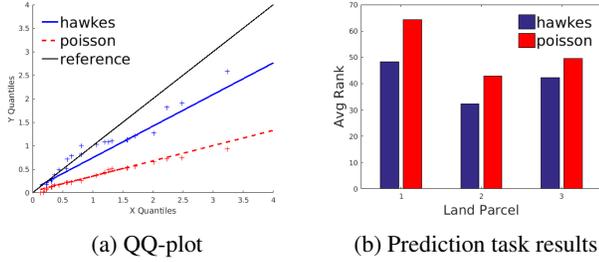


Figure 1: (a) Residual analysis for randomly chosen location in Bunchgrass Ridge. The events are better fit by a Hawkes process (slope closer to 1) than a Poisson process. (b) Hawkes process model has lower error when predicting the location of the next invasion.

The first term $\mu_i(t)$ is the rate of exogenous invasion events at node i , and the summation term captures the contribution of past invasion events (j_e, s_e) in the network towards endogenous invasions in node i at time t .

We have verified the applicability of this modeling framework using data about the encroachment of *A. grandis* trees into montane meadows at Bunchgrass Ridge in Oregon [Halpern, 2012]. We performed two goodness-of-fit tests to compare how well the spatiotemporal spread of the trees could be explained by a Hawkes process model versus a homogeneous Poisson process. The first was an established statistical test based on residual analysis for event data (Figure 1a), while the second was based on performance of each model on a prediction task (Figure 1b). Notably, we find that the Hawkes model explains the observed data better than a Poisson process (Figure 1), suggesting that there are indeed excitatory interactions in the invasion process.

Control Objectives

Given the graph representing our landscape and the invasion process dynamics described above, we can quantify the degree to which the landscape is affected by the invasive species spread at the end of our planning horizon T in a number of ways. One reasonable goal is to minimize the rate of invasions at time T subject to an intervention u at time τ , captured by $\lambda(t; u)$. Since $\lambda(t; u)$ depends on events that will stochastically occur between τ and t , it will vary across different realizations of the stochastic process, so instead we aim to minimize the total *expected* intensity at time T . Let $\eta_i(t; u) = \mathbb{E}[\lambda_i(t; u)]$, where the expectation is taken over all possible realizations of the process.

Given: A graph $G(V, E)$ representing landscape L , edge weights A with $a_{ij} = 0$ for $(i, j) \notin E$ and $a_{ij} > 0$ for $(i, j) \in E$, intervention time τ and finite time horizon T , intervention costs c_i for each node $i \in V$ and budget \mathcal{B} .

Find: A feasible intervention plan comprised of nodes $U \subseteq V$ such that $\sum_{i \in U} c_i \leq \mathcal{B}$, that minimizes $\sum_{i \in V} \eta_i(T; u)$.

Another plausible goal is to minimize the *total expected number* of invasions that occur until time T , since the ecological damage resulting from invasions is often a function of the population size. We cannot affect the process until τ , so this amounts to minimizing the number of invasions in the interval $[\tau, T)$. We store the number of invasion events at each node over time using an n -dimensional counting process

where $\mathcal{N}_i(t; u)$ represents the number of invasive species individuals that have appeared in cell i by time t . Then, given the same inputs as before,

Find: A feasible intervention plan $U \subseteq V$ such that $\sum_{i \in U} c_i \leq \mathcal{B}$, that minimizes $\sum_{i \in V} \mathbb{E}[\mathcal{N}_i(T)]$.

2.2 Hawkes Processes

A multidimensional Hawkes process can be thought of as a spatiotemporal point process, which is a random collection of points representing the time and location of events. An n -dimensional point process can be described by a counting process $\mathcal{N}(t) = (\mathcal{N}_1(t), \dots, \mathcal{N}_n(t))^\top$ where $\mathcal{N}_i(t)$ is the number of events occurring at location i before time t . The behavior of the process is characterized by the conditional intensity $\lambda(t)$. Given the history of the process up to time t , \mathcal{H}_{t-} , the expected number of events in a small time window $[t, t + dt)$ is given by $\mathbb{E}[d\mathcal{N}(t) | \mathcal{H}_{t-}] = \lambda(t)dt$.

Hawkes processes model self-exciting phenomena in which the occurrence of events causes additional events to be more likely, such as social media posts spurring reposts [Farajtabar *et al.*, 2014], earthquake aftershocks inducing further aftershocks [Ogata and Zhuang, 2006], and in this work, an invasive species individual causing another individual to appear at the same or other nodes. This self-exciting behavior is modeled using a history-dependent intensity of the form:

$$\lambda_i(t) = \mu_i(t) + \sum_{e: t_e < t} \phi_{ij_e}(t, t_e) \quad (2)$$

$$= \mu_i(t) + \sum_{j=1}^n \int_0^t \phi_{ij}(t, s) d\mathcal{N}_j(s) \quad (3)$$

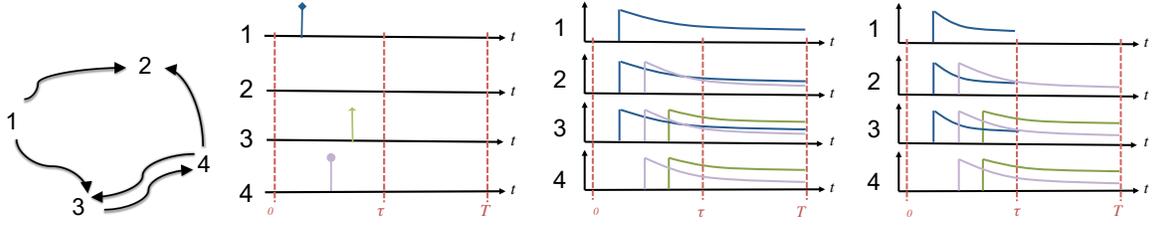
$\phi_{ij}(t, s)$ is called the impact function and captures the temporal influence of an event at location j at time s on the occurrence of events at location i at time $t \geq s$. The first term $\mu_i(t)$ is the exogenous event intensity, from outside the network and independent of the history, and the second term $\sum_{e: t_e < t} \phi_{ij_e}(t, t_e)$ is the endogenous event intensity, modeling influence and interaction within the network. Defining $\Phi(t, s) = [\phi_{ij}(t, s)]_{i,j=1 \dots n}$, $\lambda(t) = (\lambda_1(t), \dots, \lambda_n(t))^\top$, and $\mu(t) = (\mu_1(t), \dots, \mu_n(t))^\top$, we can compactly rewrite Eq (2) in matrix form:

$$\lambda(t) = \mu(t) + \int_0^t \Phi(t, s) d\mathcal{N}(s) \quad (4)$$

A common choice of impact function is the truncated exponential function $\Phi(t, s) = Ae^{-\omega(t-s)} \cdot \mathbf{1}_{\geq 0}(t-s)$ where $\phi_{ij}(t, s) = a_{ij}e^{-\omega(t-s)} \cdot \mathbf{1}_{\geq 0}(t-s)$. The coefficient a_{ij} represents the strength of the influence of j on i , and the influence of an event that occurs at time s is 0 before s and decays off after s (e.g. a social media post becomes less relevant, an infected person becomes less contagious, or an invasive species becomes less likely to survive and reproduce).

3 Discrete Interventions in Hawkes Processes

Given a network diffusion process starting at time $t_0 = 0$, suppose we plan to perform a management action at time



(a) A sample network. (b) Events before τ (c) Intensity from events (d) Effect of event removal at 1

Figure 2: A sample network and event history up to τ . Each event contributes to the intensities at the event's node and its neighbors. If the event at node 1 is deleted at τ , its contribution to intensities for $t > \tau$ disappears.

$\tau > t_0$ to steer the diffusion process over the network towards some objective at an arbitrary time $T > \tau$.

Since the intensity at any time t' only depends on the history of events up to time t' , we can also define the *state* at any time t' as $y(t') := \int_0^{t'} e^{-\omega(t'-s)} d\mathcal{N}(s)$, capturing the current effect of all events that have happened at each node up to time t' . Then considering the time of intervention τ ,

$$\begin{aligned} \lambda(t) &= \mu(t) + Ay(t) = \mu(t) + \int_0^t Ae^{-\omega(t-s)} d\mathcal{N}(s) \\ &= \mu(t) + \underbrace{\int_0^\tau Ae^{-\omega(t-s)} d\mathcal{N}(s)}_{\text{events before } \tau} + \underbrace{\int_\tau^t Ae^{-\omega(t-s)} d\mathcal{N}(s)}_{\text{events after } \tau} \\ &= \mu(t) + Ae^{-\omega(t-\tau)}y(\tau) + \int_\tau^t Ae^{-\omega(t-s)} d\mathcal{N}(s) \end{aligned}$$

Our management action entails the deletion of all events at a given set of nodes U (see Figure 2). This can alternatively be thought of as resetting the state of those locations to 0 at time τ . Therefore, for $t > \tau$ we have the intervention-dependent intensity:

$$\begin{aligned} \lambda(t; u) &= \mu(t) + Ae^{-\omega(t-\tau)}(u \circ y(\tau)) \\ &\quad + \int_\tau^t Ae^{-\omega(t-s)} d\mathcal{N}(s; u) \end{aligned} \quad (5)$$

where \circ denotes element-wise product. Vector u encodes our management action where $u_i = 0$ indicates removing the history at node i and $u_i = 1$ means not intervening at i .

3.1 Expected Behavior After Intervention

We now derive closed-form expressions for our control objectives in terms of the expected intervention-dependent intensity $\eta(t; u)$. The first objective of interest is to minimize the sum of expected rate of events at our target time: $\sum_i \eta_i(T; u)$.

By the superposition theorem of point processes, the process $\mathcal{N}(t; u)$ is decomposed into two independent processes:

$$\mathcal{N}(t; u) = \mathcal{N}_e(t; u) + \mathcal{N}_h(t; u)$$

$\mathcal{N}_e(t, u)$ is the counting process for events caused by the exogenous intensity from τ to t , and $\mathcal{N}_h(t; u)$ comprises the events generated due to the effect of previous events (history) before τ . Each of these processes have associated intensities

$\lambda_e(t; u)$ and $\lambda_h(t, u)$:

$$\lambda_e(t; u) = \mu + \underbrace{\int_\tau^t Ae^{-\omega(t-s)} d\mathcal{N}_e(s; u)}_{\text{from new events generated by } \mu} \quad (6)$$

$$\lambda_h(t; u) = \underbrace{Ae^{-\omega(t-\tau)}(u \circ y)}_{\text{from events before } \tau} + \underbrace{\int_\tau^t Ae^{-\omega(t-s)} d\mathcal{N}_h(s; u)}_{\text{from new events generated by history}} \quad (7)$$

Correspondingly, we have their expected values $\eta_e(t; u) = \mathbb{E}[\lambda_e(t; u)]$ and $\eta_h(t; u) = \mathbb{E}[\lambda_h(t; u)]$. For $\eta_e(t; u)$, we can write:

$$\eta_e(t; u) = \mu + \mathbb{E} \left[\int_\tau^t Ae^{-\omega(t-s)} d\mathcal{N}_e(s; u) \right] \quad (8)$$

$$= \mu + \int_\tau^t Ae^{-\omega(t-s)} \eta_e(s; u) ds \quad (9)$$

Using Theorem 1 from [Farajtabar *et al.*, 2016], $\eta_e(t; u) = \Psi(t)\mu$ is a solution to Equation 9 if and only if $\Psi(t) = I + \int_0^t Ae^{-\omega(t-s)}\Psi(s)ds$. For our choice of impact function:

$$\Psi(t) = I + A(A - \omega I)^{-1}(e^{(A - \omega I)t} - I) \quad (10)$$

Intuitively, $\Psi(t)$ is a matrix function indexed by i, j which are nodes. $\Psi_{i,j}(t)$ can be interpreted as the total contribution of possible events at node i at time t from events at j at any time before t (directly and indirectly).

Additionally, according to Theorem 3 in [Farajtabar *et al.*, 2016], by using integration by parts and the Laplace transform of point processes from [Farajtabar *et al.*, 2014], we can show that $\eta_h(t; u) = \Xi(t - \tau)A(u \circ y)$ where $\Xi(t) = e^{(A - \omega I)t}$. Putting these two together we have the analytical form for our first objective:

$$\mathbb{E}[\lambda(t; u)] = \Psi(t)\mu + \Xi(t - \tau)A(u \circ y) \quad (11)$$

For the second objective we aim to minimize the total expected number of events in all nodes, $\sum_i \mathbb{E}[\mathcal{N}_i(T; u)]$:

$$\mathbb{E}[\mathcal{N}_i(T; u)] = \mathbb{E} \left[\int_0^T d\mathcal{N}_i(s; u) \right] = \int_0^T \eta(s; u) ds \quad (12)$$

If we define $\Gamma(t) = \int_0^t \Psi(s)ds$ and $\Upsilon(t) = \int_0^t \Xi(s)ds$ we have:

$$\mathbb{E}[\mathcal{N}(t; u)] = \Gamma(t)\mu + \Upsilon(t - \tau)A(u \circ y). \quad (13)$$

Intuitively, $\Gamma_{i,j}(t)$ is the cumulative influence from i to j up to time t .

In summary we have;

$$\mathbb{E}[\lambda(T; u)] = \Psi(T)\mu + \Xi(T - \tau)A(u \circ y) \quad (14)$$

$$\mathbb{E}[\mathcal{N}(T; u)] = \Gamma(T)\mu + \Upsilon(T - \tau)A(u \circ y) \quad (15)$$

where

$$\Xi(t) = e^{(A - \omega I)t} \quad (16)$$

$$\Psi(t) = I + A(A - \omega I)^{-1}(e^{(A - \omega I)t} - I) \quad (17)$$

$$\Upsilon(t) = (A - \omega I)^{-1}(e^{(A - \omega I)t} - I) \quad (18)$$

$$\Gamma(t) = It + A(A - \omega I)^{-1}(\Upsilon(t) - It) \quad (19)$$

3.2 Optimization Formulations

Given the closed forms we have derived for the expected behavior of the network diffusion process after intervention, we can define our first optimization problem as:

$$\begin{aligned} & \underset{u}{\text{minimize}} \quad \sum_i \Psi(T)\mu + \Xi(T - \tau)A(u \circ y) \\ & \text{subject to:} \quad \sum_i (1 - u_i)c_i \leq \mathcal{B}, \\ & \quad \quad \quad u_i = \{0, 1\} \forall i \in \{1, 2, \dots, n\} \end{aligned} \quad (20)$$

where c_i and \mathcal{B} are defined as before.

Similarly, our second objective is:

$$\begin{aligned} & \underset{u}{\text{minimize}} \quad \sum_i \Gamma(T)\mu + \Upsilon(T - \tau)A(u \circ y) \\ & \text{subject to:} \quad \sum_i (1 - u_i)c_i \leq \mathcal{B}, \\ & \quad \quad \quad u_i = \{0, 1\} \forall i \in \{1, 2, \dots, n\} \end{aligned} \quad (21)$$

Note that the dependence of the objectives on our control variable, u , is linear since the effect of multiple events on the intensity at any given node is additive (Equation 2) and due to the linearity of expectations. If we recast our problem as maximizing the *reduction* in expected invasion rate or number of invasions (monotone modular functions) subject to a knapsack constraint, we get the classic knapsack problem. Submodular function maximization methods and approximation algorithms can thus be applied to approximately solve these optimization problems, but we can employ straightforward yet powerful binary optimization techniques to find guaranteed optimal intervention plans rather than approximations for realistically-sized instances. We used the mixed integer linear programming solver offered through the `intlinprog` function in MATLAB 2016b.

3.3 Heuristic Interventions

Besides the optimized recommendations for intervention nodes, it is also possible to choose nodes on the basis of a number of heuristics that are very natural to multidimensional Hawkes process models. In each case, nodes are considered in decreasing order of a heuristic criterion, and we greedily build a set of intervention nodes U by adding each successive node as long as there are events to remove there and the cost of intervening at the node can be covered with our remaining budget. We consider the following heuristic criteria:

- **Exogenous intensity (μ_i):** nodes with the highest rate of exogenous events.
- **Number of events until τ ($\mathcal{N}_i(\tau)$):** nodes with the highest number of events in the observation window.
- **Intensity due to global events at τ ($\lambda_i(\tau)$):** nodes with the highest intensity at the intervention time due to both local events and events in neighboring nodes.
- **Intensity due to local events (state) ($y_i(\tau)$):** nodes with the highest intensity due to only local events.

Interestingly, many of these criteria have analogs in invasion biology and control strategies based on them have been proposed. For example, exogenous intensity is comparable to “propagule pressure”, which is believed to be an important determinant of whether non-native species successfully invade new habitats [Wittmann *et al.*, 2014]. Density-based eradication (related to $\mathcal{N}_i(\tau)$) and strategies that balance the density and fecundity of the population (related to $y_i(\tau)$) have also been widely studied [Taylor and Hastings, 2004].

4 Experiments

We validate our discrete intervention methodology by testing the behavior of our analytical expressions on simulated ground-truth Hawkes process cascades, and analyze the runtime performance of our intervention optimization approach. We then apply our methodology to compare strategies for controlling the spread of a real invasive species through a landscape.

4.1 Validation of Expressions & Scalability

To verify our derived closed-form expressions against some ground-truth network diffusion processes and to test the scalability of our optimization approach, we generate synthetic networks with known parameters as described in Appendix A. We can empirically evaluate the closed-form expressions for our intervention objectives $\mathbb{E}[\lambda(T)]$ and $\mathbb{E}[\mathcal{N}(T)]$. To do this, we simulate a single realization of an invasion cascade up to time $\tau = 50$, implement a fixed intervention u and simulate many realizations of the subsequent cascade from time τ to T with which we compute the empirical mean intensity and number of invasions at each time $\tau < t \leq T$. We compare these to the theoretical expected intensity and number of invasions computed using Equations 14 and 15, following the same intervention u . Figure 3 shows how the invasion process evolves over time and after intervention, and how the theoretically computed values closely match the observed empirical mean values for both quantities, even though the actual stochastic realizations can deviate from this average.

Next, we test our discrete intervention optimization approach on networks of increasing size. The intervention cost c_i at each node is set to a fixed unit cost plus a cost proportional to the number of events there at time τ . We simulate multiple realizations of event cascades from $t = 0$ to τ and in each case compute \mathcal{B}_{all} the cost of removing all events that have appeared in the network by $t = \tau$. We set the intervention budgets \mathcal{B} as fixed percentages of \mathcal{B}_{all} to allow comparisons between the different realizations. Due to the closed-form expressions for the expected effect of an intervention

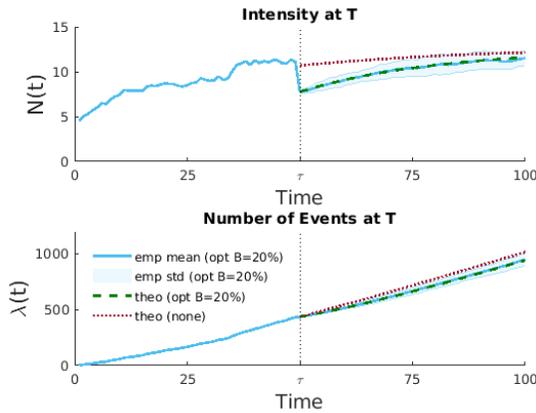


Figure 3: Simulated event cascades after an optimal intervention with $\mathcal{B} = 0.20 \cdot \mathcal{B}_{all}$ on a synthetic network compared to analytically computed intensity and number of events for the optimal plans.

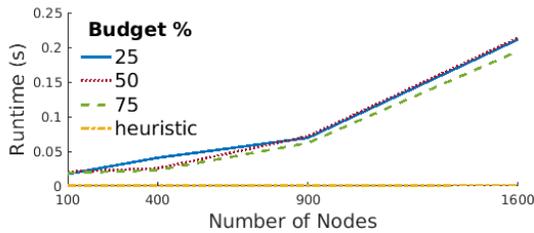


Figure 4: Runtime of ILP versus network size, compared to the time taken to compute intervention locations using a heuristic. The reported time is the average over 20 randomly simulated cascades.

action, the optimization problems in Equations 20 and 21 take the form of knapsack problems, for which MIP solvers are highly effective. Figure 4 shows that these integer programs are surprisingly fast to solve and scale relatively well for the network sizes we consider here.

4.2 Invasives Management with a Limited Budget

We apply our discrete intervention optimization for multidimensional Hawkes processes in a real-world setting to model the potential impact of different management strategies on *A. grandis* encroachment in Bunchgrass Ridge. As reported in Section 2.1, we find that a multivariate Hawkes process is able to explain the invasive spread of this species better than a homogeneous Poisson process model. We use a Hawkes process model fit to these data using the EM approach utilized in [Farajtabar *et al.*, 2016] whose efficiency in small scale datasets has been studied [Lewis and Mohler, 2011]. We assign uniform intervention costs for each land parcel to compare alternatives for budget-constrained invasive species control.

In order to study the impact of budget restrictions on the effectiveness of invasive species management, we vary the intervention budget \mathcal{B} available from 20% to 80% of \mathcal{B}_{all} . Results are shown in Figure 5. We observe that the optimized plans attain close to the same level of control as eradicating all individuals at time τ , but at only 60-80% of the cost, indicating that our method for optimal invasive management has the potential to deliver significant cost savings.

Our analytical expressions allow us to compute the expected benefit of different hypothetical management actions,

including the ones proposed by the heuristic strategies. Having our optimal solutions as a baseline, the results in Figure 5 indicate that the heuristic approaches considered here are highly sub-optimal, especially in low-budget settings. Figure 6 shows the % optimality gap of the heuristic management strategies relative to the optimal plan. The best performing heuristic approach for minimizing both the rate and the number of new invasions was the one based on the exogenous intensity μ_i , suggesting that monitoring the processes by which *A. grandis* is introduced to the Bunchgrass Ridge region could be an effective approach to curb its invasive spread.

By visualizing the spatial distribution of intervention locations (Figure 7), we find that the optimal intervention plans for minimizing each objective are different from one another. In particular, it appears that minimizing intensity focuses intervention effort at relatively few core invasion hotspots whereas minimizing the total number of invasions targets more peripheral locations. This suggests there are possible trade-offs that may be of interest to conservation planners developing long-term strategies for invasives management.

5 Conclusions and Future Work

We develop a method for characterizing the effect of a novel discrete intervention in the context of spatiotemporal Hawkes processes. This intervention consists of removing the effect of selected events in the history of the Hawkes process. We derive closed-form expressions for the effect of such an intervention action on the expected intensity and expected number of events in the network at the end of a finite time horizon T . We build on this framework to obtain optimal intervention plans employing this type of intervention mechanism, and apply this approach to the problem of controlling the spread of invasive species through a landscape. Our results on a real-world tree invasion dataset suggest that optimized intervention plans obtained using our approach achieve cost-effective control, and also provide a benchmark against which other intervention plans that are used in practice can be evaluated.

The proposed model showcases the potential of Hawkes process in managing invasive species propagation, and it can be extended in many ways to further cope with realistic settings. It can be easily generalized to handle multiple intervention points and combined with Markov Decision Processes for sequential decision making. As a first step towards planning interventions over multiple time points, one can repeatedly apply our one-step approach with a rolling horizon, observing the stochastic diffusion process between each time, which is indeed often the approach used in practice. Moreover, in practice all species may not be successfully removed after an intervention, with success probability depending potentially on the location, the species itself, budget, or time. This will add a thinning process layer to the Hawkes process. In addition, the methodologies we have developed for the invasive species management problem can be applied to other problem domains, such as in limiting the spread of disease or misinformation in social networks.

Acknowledgments

Gupta and Dilkina were partially supported by National Science Foundation Grant No. CCF-1522054 (COMPUSTNET):

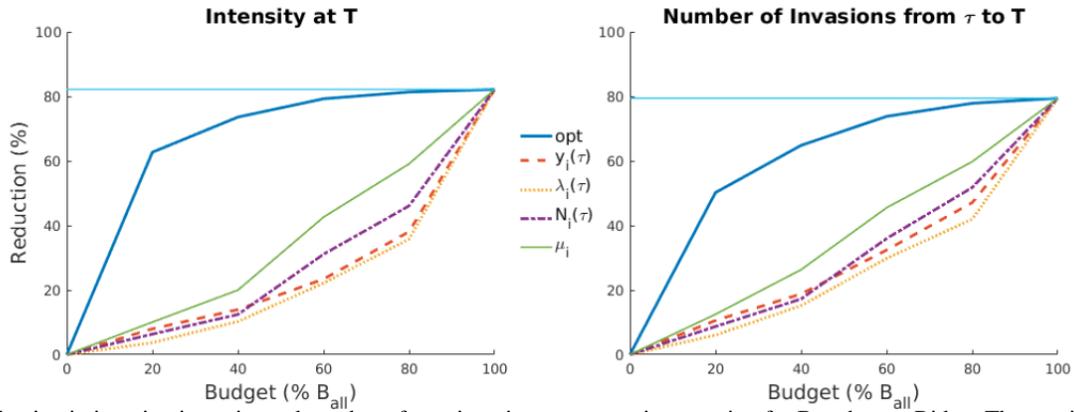


Figure 5: Reduction in invasion intensity and number of new invasions wrt to no intervention for Bunchgrass Ridge. The maximum possible reduction is obtained by intervening at all locations with invasions at time τ .

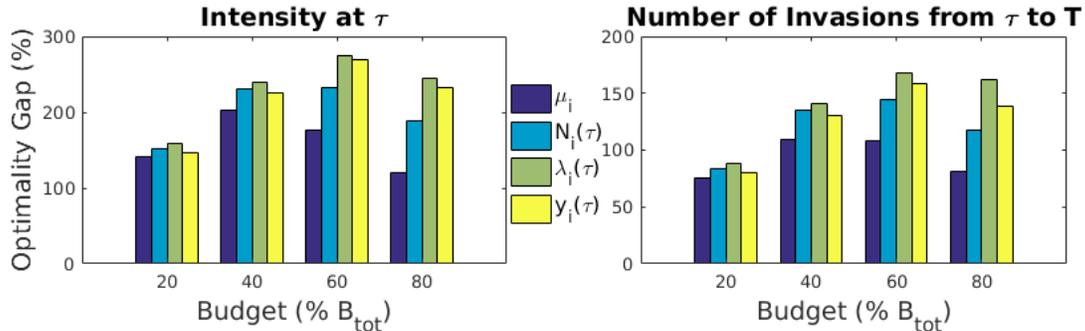


Figure 6: Mean optimality gap % for the heuristic strategies.

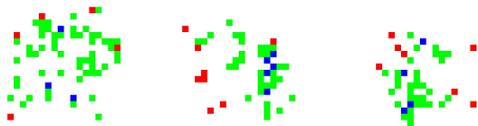


Figure 7: Overlap between optimal intervention plans for minimizing the total invasion rate versus the total number of invasions at T with $20\%B_{all}$ for 3 synthetic landscapes. Locations minimizing intensity only (blue), number of invasions only (red) and both (green).

Expanding Horizons of Computational Sustainability). We thank Charles Halpern and Julia Jones for sharing the Bunchgrass Meadows dataset. Data were provided by the HJ Andrews Experimental Forest and Long Term Ecological Research program, administered cooperatively by the USDA Forest Service Pacific Northwest Research Station, Oregon State University, and the Willamette National Forest. This material is based upon work supported by the National Science Foundation under Grant No. DEB-1440409.

A Synthetic Network Generation

We generate synthetic landscapes to test our invasive species control strategy. Each landscape consists of an N -by- N grid of cells. The exogenous intensity at each node μ_i is constant over time and is a uniformly distributed random variable in $[0, \mu_{max}]$ with $\mu_{max} = 0.02$. A small number of nodes have a higher exogenous intensity and act as seed points in the diffusion process. We assign the mutually-exciting parameters a_{ij} to be values in the range $[0, a_{max}]$, with $a_{max} = 0.05$. Finally, we set the influence decay rate $\omega = 0.15$ and a finite

time horizon of $T = 100$.

The a_{ij} parameters take the form $a_{max} \cdot e^{-d_{ij}^2}$. a_{max} can be thought of as the establishment success rate of offspring of the invasive species in ideal conditions. We set the establishment success rate of the invasives to $a_{max} = 0.05$. This value can be scaled by habitat suitability at the destination (assumed uniform in our networks, but can be substituted with a habitat suitability matrix generated with a mixture of 2D Gaussian functions, or using expert opinion and maps of environmental covariates). We assume dispersal occurs only between each cell and its 8-cell neighborhood. We can also generate networks with links between a small number of pairs of non-adjacent cells, to model the effect of occasional long-range dispersal events on invasive species spread.

References

[Arim *et al.*, 2006] Matías Arim, Sebastián Abades, Paula Neill, Mauricio Lima, and Pablo Marquet. Spread dynamics of invasive species. *PNAS*, 103(2):374–378, 2006.

[Balderama *et al.*, 2012] Earvin Balderama, Frederic Paik Schoenberg, Erin Murray, and Philip Rundel. Application of branching models in the study of invasive species. *J Am Stat Assoc*, 107(498):467–476, 2012.

[Bjorck *et al.*, 2018] Johan Bjorck, Yiwei Bai, Xiaojian Wu, Yexiang Xue, Mark C Whitmore, and Carla Gomes. Scalable relaxations of sparse packing constraints: Optimal biocontrol in predator-prey network. In *Proc. of the 32nd AAAI Conference on Artificial Intelligence (AAAI-18)*. AAAI, 2018.

- [Büyükahtakin *et al.*, 2014] İ Esra Büyükahtakin, Zhuo Feng, and Ferenc Szidarovszky. A multi-objective optimization approach for invasive species control. *J Oper Res Soc*, 65(11):1625–1635, 2014.
- [Eames and Keeling, 2002] Ken Eames and Matt Keeling. Modeling dynamic and network heterogeneities in the spread of sexually transmitted diseases. *PNAS*, 99(20):13330–13335, 2002.
- [Farajtabar *et al.*, 2014] Mehrdad Farajtabar, Nan Du, Manuel Gomez Rodriguez, Isabel Valera, Hongyuan Zha, and Le Song. Shaping social activity by incentivizing users. In *NIPS*, pages 2474–2482, 2014.
- [Farajtabar *et al.*, 2016] Mehrdad Farajtabar, Xiaojing Ye, Sahar Harati, Le Song, and Hongyuan Zha. Multistage campaigning in social networks. In *NIPS*, pages 4718–4726, 2016.
- [Fitzpatrick *et al.*, 2012] Matthew Fitzpatrick, Evan Preisser, Adam Porter, Joseph Elkinton, and Aaron Ellison. Modeling range dynamics in heterogeneous landscapes: invasion of the hemlock woolly adelgid in eastern north america. *Ecol Appl*, 22(2):472–486, 2012.
- [Halpern, 2012] Charles Halpern. Ecology and restoration of montane meadows at bunchgrass ridge near the andrews experimental forest, 1999–2009. bunchgrass study., 2012.
- [Hastings *et al.*, 2005] Alan Hastings, Kim Cuddington, Kendi Davies, Christopher Dugaw, Sarah Elmendorf, Amy Freestone, Susan Harrison, et al. The spatial spread of invasions: new developments in theory and evidence. *Ecology Letters*, 8(1):91–101, 2005.
- [Kempe *et al.*, 2003] David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proc. of the th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 137–146. ACM, 2003.
- [Khalil *et al.*, 2014] Elias Boutros Khalil, Bistra Dilkina, and Le Song. Scalable diffusion-aware optimization of network topology. In *Proc. of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1226–1235. ACM, 2014.
- [Kim, 2011] Harry Kim. *Spatio-temporal point process models for the spread of avian influenza virus (H5N1)*. UC Berkeley, 2011.
- [Kimura *et al.*, 2008] Masahiro Kimura, Kazumi Saito, and Hiroshi Motoda. Minimizing the spread of contamination by blocking links in a network. In *AAAI*, volume 8, pages 1175–1180, 2008.
- [Lewis and Mohler, 2011] Erik Lewis and George Mohler. A nonparametric em algorithm for multiscale hawkes processes. *J Nonparametric Stat*, 1(1):1–20, 2011.
- [Lin *et al.*, 2016] Peng Lin, Bang Zhang, Ting Guo, Yang Wang, and Fang Chen. Interaction point processes via infinite branching model. In *AAAI*, pages 1853–1859, 2016.
- [Ogata and Zhuang, 2006] Yosihiko Ogata and Jiancang Zhuang. Space–time etas models and an improved extension. *Tectonophysics*, 413(1):13–23, 2006.
- [Pauchard *et al.*, 2016] Aníbal Pauchard, Adrián Escudero, Rafael A García, Marcelino Cruz, Bárbara Langdon, Lohengrin A Cavieres, and Jocelyn Esquivel. Pine invasions in treeless environments: dispersal overruns microsite heterogeneity. *Ecology and Evolution*, 6(2):447–459, 2016.
- [Rizoiu *et al.*, 2018] Marian-Andrei Rizoiu, Swapnil Mishra, Quyu Kong, Mark Carman, and Lexing Xie. Sir-hawkes: Linking epidemic models and hawkes processes to model diffusions in finite populations. In *Proc. of the 2018 World Wide Web Conference on WWW*, pages 419–428. International World Wide Web Conferences Steering Committee, 2018.
- [Sakai *et al.*, 2001] Ann Sakai, Fred Allendorf, Jodie Holt, David Lodge, Jane Molofsky, Kimberly With, Syndal-las Baughman, et al. The population biology of invasive species. *Annu Rev Ecol Syst*, 32(1):305–332, 2001.
- [Sheldon *et al.*, 2010] Daniel Sheldon, Bistra Dilkina, Adam Elmachtoub, Ryan Finseth, Ashish Sabharwal, Jon Conrad, Carla Gomes, David Shmoys, William Allen, Ole Amundsen, et al. Maximizing the spread of cascades using network design. In *Proc. of the 26th Conference on Uncertainty in Artificial Intelligence*, pages 517–526. AUAI Press, 2010.
- [Taleghan *et al.*, 2015] Majid Alkaee Taleghan, Thomas G Dietterich, Mark Crowley, Kim Hall, and H Jo Albers. Pac optimal mdp planning with application to invasive species management. *JMLR*, 16(1):3877–3903, 2015.
- [Taylor and Hastings, 2004] Caz Taylor and Alan Hastings. Finding optimal control strategies for invasive species: A density-structured model for spartina alterniflora. *J Appl Ecol*, 41(6):1049–1057, 2004.
- [Wang *et al.*, 2016] Yichen Wang, Evangelos Theodorou, Apurv Verma, and Le Song. Steering opinion dynamics in information diffusion networks. *arXiv preprint arXiv:1603.09021*, 2016.
- [Wittmann *et al.*, 2014] Meike Wittmann, Dirk Metzler, Wilfried Gabriel, and Jonathan Jeschke. Decomposing propagule pressure: the effects of propagule size and propagule frequency on invasion success. *Oikos*, 123(4):441–450, 2014.
- [Wu *et al.*, 2015] Xiaojian Wu, Daniel Sheldon, and Shlomo Zilberstein. Efficient algorithms to optimize diffusion processes under the independent cascade model. *NIPS Work. on Networks in the Social and Information Sciences*, 2015.
- [Yang and Leskovec, 2010] Jaewon Yang and Jure Leskovec. Modeling information diffusion in implicit networks. In *ICDM*, pages 599–608. IEEE, 2010.
- [Zarezade *et al.*, 2017] Ali Zarezade, Utkarsh Upadhyay, Hamid Rabiee, and Manuel Gomez-Rodriguez. Redqueen: An online algorithm for smart broadcasting in social networks. In *Proc. of the 10th ACM International Conference on Web Search and Data Mining*, pages 51–60. ACM, 2017.