Modeling treatment costs associated with a multi-stage pandemic

Michael A. Ekhaus
Gibraltar Analytical
Minneapolis, Minnesota
Email: mekhaus@GibraltarAnalytical.com

Abstract

Abstract: Introduced is an Interacting Particle System (IPS) approach to modeling a multiple stage pandemic. The model population receives differing treatments with possibly differing costs depending on the stage of exposure/infection. The total cost of treatment is a time varying utility function of the ensemble population for which this approach uses particle system methods to simulate the cost of maintaining survival of the population. Interacting Particle Systems are a class of spatial-temporal stochastic processes suitable for studying the spread of infectious diseases and other interaction phenomena.

The paper gives a brief background on particle systems with specific focus on "the contact process", followed by expanding the contact process to represent more realistic state transitions.

1 Introduction

Interacting particle systems (IPS) is a branch of modern probability theory which began in the early 1970's to rigorously study problems motivated from Statistical Physics. Subsequently these methods are being applied to problems in biology, economics, and others fields. In joint work with Sandia National Laboratories, the author has applied particle system methods to study cyber conflicts. Roughly speaking, a computer network will experience catastrophic events if the percentage of computers infected with viruses exceeds a threshold. Since the infected machines may be under the control of malicious forces, understanding how to model and predict network events is crucial. In addition, it is important to understand which strategies are effective in combating attackers while still allowing computers to perform useful work. Specifically, if nefarious interests can cause computer networks to effectively spend all their processing to combat attacks, then the attackers are successful. Similarly, this work begins with questions about how policies and constraints can effect the outcome and effectiveness of containing a pandemic outbreak.

There are three broad categories of particle systems: Contact, Voter, and Exclusion type processes. T. Liggett (see [1], [2]) has written two excellent treatises on IPSs with particular focus on the mathematical intricacy.

This paper deals with with variations on the now classical "Contact Process". The contact process is a spatial-temporal Markov process exhibiting a phase transition and depicts the spread of disease in idealized geographical setting via interactions among the models population, which are modeled as sites on a lattice. The phase transition associated with the contact process is an abrupt change in the probability that an infection can survive and propagate

through out the model. This paper intends to extend particle system methods to more accurately represent factors of pandemics and in general epidemiology, while suggesting a method for analyzing the cost of managing a pandemic. As with all models, these models are still idealizations. This paper takes the viewpoint, and subsequently models, cost as a time-series in which the accumulated cost increases over the evolution of the pandemic. Currently, there is no attempt within to study "optimal cost". Questions of optimality will require further analysis and should use real data. This paper primarily focuses on modeling methodology. Furthermore, the model does not yet include multiple waves within a pandemic nor mutation of the disease.

Although the model is preliminary, this work illustrates issues concerning policy effectiveness in dealing with phenomenon that exhibit phase transitions. Small changes in policy and input parameters may result in large changes in the outcome.

2 Motivation from two dependent Markov Chains

We begin by consider a deceivingly simple model: a 2 state Markov chain. Let the states represent being healthy and being infected. If one is healthy, then with a certain probability one may become infected and vice verse. Such chains are useful and serves as a good starting point. What if we're considering 2 individuals that have some contact with each other, and that the probability of transitioning from being healthy to being infected depends on the other person's health. The individual chains will no longer be Markovian and one is naturally led to studying the Markov Chain describing the pair of individuals. Following this line of reasoning, consider 3 individuals which all have contact with each other. The probability an individual becomes sick should depend on the number of other individuals which this individual has contact that are infected. More precisely, it should depend on the density of possible individuals with which one has contact that are infected. What if we're considering a large, but finite population of individuals (N) and that for each individual there is a list of other individuals with whom this person has contact. A state of this population is represented by whether each individual is healthy or infected, and results in 2^N configurations of the population. Even in cases for which limited contact among individuals is possible, the resulting Markov Chain will become very difficult to analyze by the traditional linear algebra approach used for studying Markov Chains. This is the context for studying the contact process and other interacting particle systems.

2.1 Background: The Contact Process

The classical contact process (see [1]) is a continuous-time Markov process with state space $\{0,1\}^{Z^d}$, where Z^d is the infinite d-dimensional integer lattice. In order to avoid defining the mathematical machinery for infinite particle systems, consider a contact process on a finite square lattice denoted by

$L^2 = \{1, 2, 3, \cdots L\} \times \{1, 2, 3, \cdots, L\}$.

For each $x \in L^2$ there are two possible values 0 or 1 (healthy or infected). An element in $\{0, 1\}^{L^2}$ may be identified with subsets of L^2. Let A be the set of all x in L^2 that are infected. Transitions for the Markov Chain occur as follows:

- A transitions to $A \backslash \{x\}$ for x in A with probability δ, and
- A transitions to $A \cup \{x\}$ for x not in A with probability $\rho_x(I)$. Where $\rho_x(I)$ is the percentage of neighbors (*i.e. density of neighbors*) of x which are infected.

The model may be described by a collection of identical transition diagram denoted in figure 1.

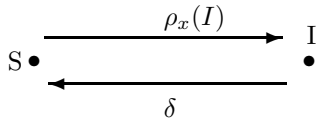

Figure 1: transition diagram for the contact process

The neighborhood structure for the contact process is traditionally the 4 nearest neighbors associated with the regular square lattice, but recently study of the contact process on more general graphical structures has become of interest. Since air travel is so prevalent, the graphical structure should allows long-range connections between the population, whereby one entity can infect another that is not geographically neighboring. The fundamental mathematical questions associated with the contact process involve understanding the long term survivability of the process (i.e. survival of an infectious disease) as a function of the parameter δ. Often the problem is stated for the infinite square lattice and the process is defined in terms of Poisson rates. We'll continue with the finite setting, and concern ourselves with probabilities of events, rather than rates at which these events occurs.

An assignment of 0 or 1 to each $x \in L^2$ is easily visualized as is bitmap. Figure 2 shows the evolution of a contact process on a square lattice with each site having 4 nearest neighbors (wrapping around at the boundaries). There are 22,500 sites in this lattice and scaled time = 1000 indicates that the associated discrete time Markov Chain was run 22,500,000 iterations $(22,500 \times 1000)$ and $\delta = 0.6$. The process was started from an initial configuration of mostly healthy sites (shown in green) with a core of infected sites (shown in red).

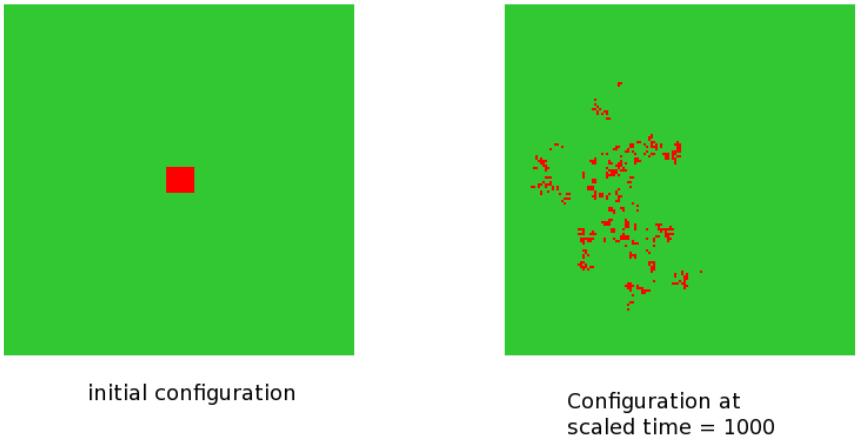

Figure 2: 2-d contact process

3 Long-Range Connections: Watts-Strogatz type graphs

Although the 4 nearest neighbors is a crude idealization of our worlds topology, it is a reasonable starting place for considering local interactions due to its structure and it's always possible to extend the approach to more realistic graphs.

As noted, it is important to allow long-range connections between sites. The topology will be idealized as a Watts-Strogatz type graph (see [3]).

Starting from a finite square lattice, one "rewires" the lattice according to a probability $p \in [0, 1)$ as follows: for each site in the lattice, pick another site uniformly over all other possible sites and with probability p an edge is added to the lattice linking these two sites. Watts and Strogatz also removed an edge in order to maintain a fixed degree for the graph. The graph considered here does not require a fixed degree and long-range connections are only added. The neighboring sites of x are those sites which are linked to x by one edge. Note that the nearest neighbors of x are the original 4 nearest neighbors and any additional sites achieved by the rewiring.

Starting from the same initial configuration as figure 2 and $\delta = 0.6$, but using the additional rewiring of the topology as described above with probability 0.1. The resulting contact process spreads throughout the entire population and reaches equilibrium in under a scaled time of 200 (see figure 3). Compared this with the previous scenario in which the contact process as yet to spread to the entire domain in 5 times the amount of time.

The use of Watts-Strogatz type graphs and other long-range connection graphs is to model the connectedness of people in the world and has been pop-

ularized in the media as: Six degrees of separation. The phenomena is usually described in a positive sense, but when considering a pandemic model this connectedness is a problem. It enhances a virus' ability to branch and spread between hosts. This issue will become crucial in the latter part of this paper when exploring a variety of scenarios.

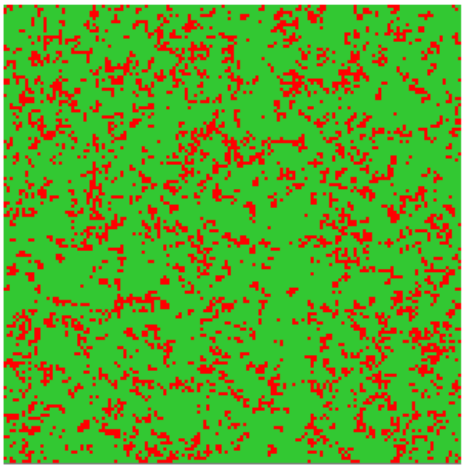

Figure 3: 2-d contact process with long range connections

4 Multi-Stage Pandemic Model

Several variations on the contact process will be discussed. First our model will allow a site in the graph (idealized person) to transition between seven states and allows the process to be non-homogeneous. The additional states will allow for varying costs to be assigned to treatments and the non-homogeneity allows (per site) for varying probabilities that treatment is successful. Furthermore, although seven states is probably not sufficient to model a real world scenario, it suffices to illustrate our methodology. Readers interested in rigorous analysis of contact processes in the multi-stage setting are referred to Krone ([4]). Krone's work considers 3 allowable states per site.

The other extensions are non-traditional and make rigorous analysis of the models, at best, difficult. These extensions impose deterministic policies and constraints on the stochastic process, thus making description of the transitions by way of a transition diagram very difficult to read. For this reason, the traditional extension is presented with a transition diagram followed later by a section describing the policies and constraints.

Associated with the paper, this model was simulated allowing each extension to be independently enforced or not. Whether a policy or constraint significantly contributes to the behavior depends on the models sensitivity to these parameters, and the interactions between policies and constraints.

4.1 state representation

There are 2 categories of states associated with the model, states that receive treatment and have a related cost for a site obtaining the state and the other are non-treatment states.

- **S** - a state in which a site is susceptible to becoming exposed to the pandemic.
- **E** - a state in which a site has been exposed.
- **I** - a state in which a site having been exposed, now has become infected.
- **D** - a state in which the site has become inert to the pandemic. Possibly due immunity or deceased.
- **N** - a state in which the site has changed the risk of exposure. Possibly as a result of an anti-viral. Although not currently modeled in this manner, this state could possibly occur as a result of awareness that an individual may have about the pandemic. For example, due to awareness, one might be more careful about contact with others and hence, lower the probability of exposure. The awareness may be linked back to the percentage of individuals that are infected and deceased.
- **T1** - a state in which the site having been exposed, now receives treatment.
- **T2** - a state in which the site having been infected, now receives treatment.

Although it is arguable that every state has a cost associated with it, the states are tabulated in table 1 as follows

Table 1: Model State Description

Category	Value	Color
Non-treatment states	S	Green
	E	Orange
	I	Red
	D	Black
Treatment states (having cost)	N	DarkBlue
	T1	Yellow
	T2	Magenta

Table 2 shows the allowable transitions that a site can make and the parameters associated with these transitions. Loops are not shown in the diagram. For example,self transitions, such as $S \to S$ are not shown in the diagram.

Table 2: multi-stage contact transition diagram

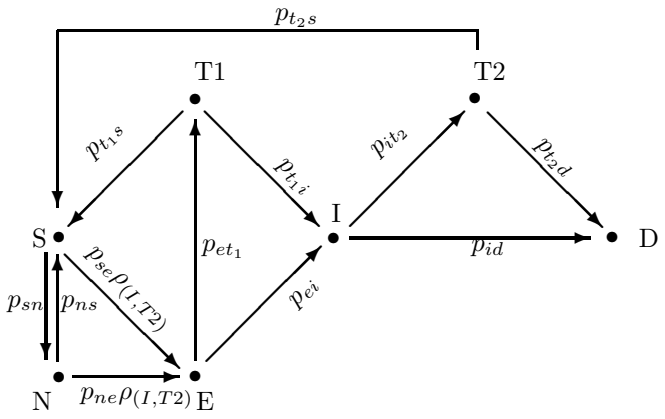

A configuration of this process is an assignment of value to each site x in the graph defined earlier (fixed size L^2 and rewiring probability p). As such, a configuration is a 7 colored bitmap. Note that the long-range connections are not seen in the representation of the configuration. The set of all configurations, or the state space, is $\{S, E, I, D, N, T1, T2\}^{v(\mathbf{G}_L(p))}$, where $v(\mathbf{G}_L(p))$ are the sites of the Watts-Strogatz graph $\mathbf{G}_L(p)$. The dynamics will be a Markov Chain on the state space $\{S, E, I, D, N, T1, T2\}^{v(\mathbf{G})}$. Simulation for subsequent scenarios used lattice with 62,500 sites and a Markov Chain had $7^{62,500}$ states.

4.2 Pandemic dynamics

Throughout the simulation a rewiring probability of 0.01 and 0.15 were used to generate the Watts-Strogatz type graph that the pandemic evolves on.

4.2.1 traditional probabilities

Table 3 categorizes all the traditional Markov transitions for this model as to whether there is a cost associated with the transition and also the probability used in the simulations. For a visual aid, figure 4 presents a snapshot of this 7-stage contact process after approximately 4.5 million iterations of the associated Markov Chain.

Table 3: Transition Table

	transition	transition probability
Transition with Cost	$S \to N$	p_{sn}
	$E \to T1$	p_{et1}
	$I \to T2$	p_{it2}
	$N \to N$	$1 - p_{ns} - p_{ne}$
	$T1 \to T1$	$1 - p_{t1s} - p_{t1i}$
	$T2 \to T2$	$1 - p_{t2s} - p_{t2d}$
Transitions without Cost	$N \to S$	p_{ns}
	$N \to E$	$p_{ne}\rho_{(I,T2)}$
	$S \to E$	$p_{se}\rho_{(I,T2)}$
	$E \to I$	p_{ei}
	$T1 \to S$	p_{t1s}
	$T1 \to I$	p_{t1i}
	$I \to D$	p_{id}
	$T2 \to S$	p_{t2s}
	$T2 \to D$	p_{t2d}
	$S \to S$	$1 - p_{sn} - p_{se}\rho_{(I,T2)}$
	$E \to E$	$1 - p_{et1} - p_{ei}$
	$I \to I$	$1 - p_{it2} - p_{id}$
	$D \to D$	1

Since the stochastic process has only one mechanism for the spread of infection (spread via contact) and that mechanism causes a transition to only state E, it is plausible that the process is "essentially" a contact process if one disregards the state D. The phase portrait for the contact process is completely understood, but understanding the phase portrait for this 7-stage contact process is extremely difficult due to the number of path possibilities. Consequently, understanding the sensitivity associated with parameters and the overall effect on the evolution is also extremely complex. For this exposition, the parameters have been chosen to illustrate a range of possible behaviors resulting from in-

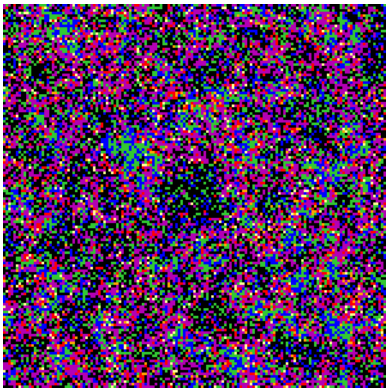

Figure 4: 2-d multi-stage contact process

teractions among the policies and the constraints. The policies and constraints are elaborated next.

4.2.2 policy transitions and constraints

Augmenting a particle system model with policies and constraints is important from many perspectives. First, suppose one had complete knowledge of all events by which a pandemic could evolve. It's unrealistic to assume that we would merely compute our survival probably and let the rest result according to some complicated roll of the "dice". In fact, we are dealing with far less than complete knowledge and would have good reason to control anything that might help the outcome. In this regard, a pandemic is a conflict/war between species in which there is no negotiating.

For the purposes of this work, cost of the various treatments is represented by the occupation density of states associated with treatments and is accumulated over time. There are several policies and constraints which will be placed on this 7-stage contact process and we're interested in exploring the effects on the process behavior and the accumulate cost functions.

In the IPS literature, focus is usually (but not always) placed on infinite graphs and long term asymptotic system behavior. As mentioned earlier, we're interested in large but finite graphs and also in the short-term transient behavior. In particular, non-equilibrium dynamics and convergence to equilibrium. To the extent that state "D" represents both immunity and deceased, it is really to be thought of as the deceased case. As such, it is a goal to converge to distributions with as small a rate of increase for the deceased density.

In order to introduce that sites are not equally likely to have successful treatments, a ranking was assigned to x using a uniform [0,1) random variable (denoted $r(x)$). Although use of the uniform distribution is probably not realistic, it suffices for our purpose. The ranking $r(x)$ indicates the probability that treatment is unsuccessful.

Susceptibility constraint: This constraint modulates the transitions probabilities associated with susceptibility as follows

- replace $p_{t1s} \rightarrow p_{t1s} \times (1 - r(x))$
- replace $p_{t1i} \rightarrow p_{t1i} \times r(x)$
- replace $p_{ei} \rightarrow p_{ei} \times r(x)$
- replace $p_{id} \rightarrow p_{id} \times r(x)$
- replace $p_{t2s} \rightarrow p_{t2s} \times (1 - r(x))$
- replace $p_{t2d} \rightarrow p_{t2d} \times r(x)$

For example, the higher the susceptibility, the more likely a site is to becoming infected and the less likely it is to recovery. This constraint rescales the transition probabilities accordingly.

Resistance constraint:

This constraint modulates the transitions probabilities associated with transitions into the infected state, I. The idea is that a site acquires a resistance to becoming infected as a function of the number of previous times the site has been infected. The number of times a site has been infected is equated to the number of times a site as successfully thwarted infection previously. Let res denote the scaling factor representing resistance. The following probabilities are rescaled as follows

- Replace $p_{ei} \rightarrow p_{ei} \times \{\sharp\text{times previously infected}\}^{-res}$
- Replace $p_{t1i} \rightarrow p_{t1i} \times \{\sharp\text{times previously infected}\}^{-res}$

Treatment policy 1: This policy modulates the transition probabilities associated with receiving treatment based on susceptibility as follows

- replace $p_{sn} \rightarrow p_{sn} \times r(x)$
- replace $p_{et1} \rightarrow p_{et1} \times (1 - r(x))$
- replace $p_{it2} \rightarrow p_{it2} \times (1 - r(x))$

For example, the higher the susceptibility, the more likely a site is to receive treatment and transition to state N, but also the less likely the site is to receive treatments $T1$ and $T2$.

Availability of treatment constraint: This constraint determines the availability of treatment for each possible treatment and if a treatment is unavailable blocks transitions into states associated with the respective treatment. A treatments availability is renewed periodically and the period is an input parameter. Treatment associated with state N is considered a preventive measure, such as use of a anti-viral medication, and this policy blocks excessively applying this treatment. A site x can remain in state N, but the cost function can only

be incremented if the previous increment is more than a prescribed number of scaled time periods. That is, treatment associated with state N is effective for a period of time and one need not apply treatment until after the treatment has expired.

Treatment policy 2: This policy purposes to have sites take preventative measure if a local threat exists. The policy is modeled by overriding the probabilistic nature of receiving treatment associated with state N, providing treatment is available, based on the existence of local threat of possible exposure. Using a local region of the 24 nearest neighbors of site x from only the square lattice. If any of these neighbors are in state I or $T2$, the state x transitions to state N.

Local quarantine policy: This policy replaces $\rho_{(I,T2)} \rightarrow \rho_{(I)}$, effective quarantining sites that are infected and receiving treatment ($i.e$ site in state $T2$). There is a simplifying assumption about the world being modeled: if "one" is knowingly infected, one would seek any available treatment.

Global isolation policy: Since long-range connections allow a pandemic to branch and spread quickly, this policy aims at utilizing the global density of infected sites to switch between allowing and disallowing long range connections. More precisely, if the global density of infected sites exceeds a given threshold, the policy switches to using only the nearest neighbors from the regular square lattice and not the Watts-Strogatz portion of the graph.When the density falls sufficiently, long-range connections are again allowed.

Although this policy is useful in curtailing the branching of infection, it is far to aggressive because, as simulations showed, this policy disallows long-range connects far to often. Remember, long-range connections represent interactions beyond ones immediate connections: air travel, going to a work place, etc.

Augmented global isolation policy: This policy refines the isolation objectives of the "global isolation policy" without entirely eliminating long-range connection, but rather augments the global isolation policy with an intermediate threshold which triggers a random screening process. There is a parameter that indicates the degree (or depth) of screening. If the infection density were to rise beyond the intermediate threshold and exceed the original global threshold, then the original global policy goes into effect and disallows all long-range connections until the infection density is sufficiently reduced.

5 Simulation

For the simulation the lattice was increased from 22,500 sites to 62,500 sites. Simulation of this model was performed by iteratively picking a random site uniformly from among the 62,500 sites. Then transition table 2 is used together with selected policies and constraints to evaluate updating the state at site x. A rescaling of time, #iterations/62,500, has been used, rather than directly using #iterations.

For this model, there is no a priori reason to sample the site updated from other than the uniform distribution, but it is worth noting that one could sample from a different distribution. Furthermore, the state updates could be have been

done synchronously, as opposed to asynchronously as was the case. In some modeling problems synchronized updates are of interest.

5.1 Observations

In general, analysis of such a model will require that one decide which type of scenarios are of interest and plan to analysis portions of the phase portrait. There were several parameter sets used in the following scenarios and some discussion is warranted. Table 4 indicates the wiring probability and the probabilistic parameters associated with traditional transitions indicated in Table 2. Parameter set 1 uses a rather small wiring probability, reminiscent of society that is less connected by long range connection and more localized. The parameter sets with 0.15 wiring probability are to indicated a more modern society with a significantly larger degree of connectivity. The choice for rewiring probabilities was speculatively chosen with the larger being a factor of 15 times with the thought that our current world might reasonably be that much more connected than the world might have been in 1918.

Table 4: Parameter Sets

	wiring prob	p_{se}	p_{sn}	p_{ns}	p_{ne}	p_{et1}	p_{ei}	p_{t1s}	p_{t1i}	p_{it2}	p_{id}	p_{t2s}	p_{t2d}
parameter set 1	0.01	0.80	0.0	0.0	0.0	0.1	0.1	0.1	0.1	0.75	0.002	0.1	0.0001
parameter set 2	0.15	0.80	0.05	0.05	0.20	0.1	0.1	0.1	0.1	0.75	0.002	0.1	0.0001
parameter set 2a	0.15	0.80	0.025	0.05	0.20	0.05	0.1	0.1	0.1	0.375	0.002	0.1	0.0001
parameter set 3	0.15	0.80	0.05	0.05	0.20	0.15	0.1	0.1	0.1	0.75	0.002	0.1	0.0001
parameter set 4	0.15	0.80	0.05	0.05	0.20	0.15	0.1	0.15	0.1	0.75	0.002	0.1	0.0001
parameter set 5	0.15	0.80	0.05	0.05	0.20	0.15	0.1	0.15	0.05	0.75	0.002	0.1	0.0001
parameter set 6	0.15	0.80	0.05	0.05	0.20	0.35	0.1	0.15	0.05	0.75	0.002	0.15	0.0001

Table 5 gives sample outcomes associated with parameter set 1. This parameter set was chosen with the intent of being a benchmark for the later parameter sets. The parameters were chosen so that when approximately 30 percent of the sites had been infected at some time prior, that the mortality was about 0.5 percent and the time period was roughly several hundred periods. In order simplify the tables, sites have been grouped healthy, sick, and deceased/mortality. States $\{S, E, N, T1\}$ are healthy states, $\{I, T2\}$ are sick states, and $\{D\}$ is the deceased state [1].

Table 5 gives sample outcomes mostly associated with parameter set 2. This parameter set was chosen with the intent of representing a society and that has a significantly larger connectivity and also the allows transitions to state $\{N\}$. Recall that state $\{N\}$ represent a preventative treatment such as taking an antiviral medication. Note that the distribution of healthy, sick, and deceased for

[1] the percentage of the combined categories should sum to 1.00, but for the sake of honesty, a error in the code has resulted in the percentages being slightly biased.

Table 5: Summary I

	parameter set	susceptibility	resistance	treatment policy 1	treatment policy 2	quarantine	isolation	max treatments	time period at 30%	healthy at 30%	sick at 30%	mortality at 30%	healthy at 2k	sick at 2k	mortality at 2k
scenario-1	1	-	-	-	-	-	-	-	272	0.787	0.201	0.004	0.372	0.003	0.616
scenario-2	1	✓	-	-	-	-	-	10^6	545	0.802	0.185	0.005	0.223	0.356	0.412
scenario-3	1	✓	✓(-3.0)	-	-	-	-	10^6	500	0.806	0.182	0.004	0.865	0.018	0.109

scenarios 1,2, and 3 have drastically different long term behavior, but are rather similar when the respective samples reached 30 percent coverage of infection. Also, coverage of infection should be understood to mean that the set of sites that have been infected, regardless of whether the sites are currently infected.

Table 6: Summary II

	parameter set	susceptibility	resistance	treatment policy 1	treatment policy 2	quarantine	isolation	max treatments	30%	healthy at 30%	sick at 30%	mortality at 30%	healthy 2k	sick 2k	mortality 2k
scenario-A	2	-	-	-	-	-	-	-	184	0.857	0.134	0.001	0.558	0.293	0.141
scenario-B	2	-	-	-	-	-	-	10^6	154	0.851	0.140	0.001	0.384	0.001	0.607
scenario-C	2	-	✓(-3.0)	-	-	-	-	10^6	154	0.883	0.108	0.001	0.780	0.000	0.192
scenario-D	2	✓	-	-	-	-	-	10^6	223	0.801	0.190	0.001	0.253	0.282	0.457
scenario-E	2	✓	✓(-3.0)	-	-	-	-	10^6	237	0.822	0.169	0.001	0.867	0.013	0.112
scenario-F	2	✓	✓(-3.0)	✓	-	-	-	10^6	203	0.767	0.221	0.004	0.803	0.006	0.183
scenario-G	2	✓	✓(-3.0)	-	✓	✓	✓(4-1)	10^6	245	0.826	0.165	0.001	0.870	0.012	0.109
scenario-H	2	✓	✓(-3.0)	-	✓	✓	✓(4-2)	10^6	262	0.811	0.179	0.002	0.870	0.013	0.109
scenario-Ha	2a	✓	✓(-3.0)	-	✓	✓	✓(4-2)	10^6	213	0.804	0.186	0.002	0.848	0.016	0.128
scenario-I	2	✓	✓(-3.0)	-	✓	✓	✓(5-1)	10^6	275	0.800	0.189	0.003	0.873	0.012	0.107
scenario-J	2	✓	✓(-3.0)	-	✓	✓	✓(10-1)	10^6	263	0.811	0.180	0.002	0.872	0.013	0.107
scenario-K	2	✓	✓(-3.0)	-	✓	✓	✓(10-2)	10^6	247	0.823	0.167	0.001	0.872	0.013	0.107
scenario-L	2	✓	✓(-3.0)	✓	✓	✓	✓(4-2)	10^6	249	0.818	0.170	0.003	0.812	0.007	0.173
scenario-M	2	✓	✓(-4.0)	-	✓	✓	✓(10-2)	10^6	260	0.820	0.170	0.002	0.893	0.011	0.088
scenario-N	2	✓	✓(-5.0)	-	✓	✓	✓(10-2)	10^6	245	0.833	0.158	0.001	0.904	0.008	0.080
scenario-O	3	✓	✓(-5.0)	-	✓	✓	✓(10-2)	10^6	272	0.809	0.181	0.002	0.906	0.009	0.077
scenario-P	4	✓	✓(-5.0)	-	✓	✓	✓(10-2)	10^6	288	0.792	0.198	0.002	0.908	0.008	0.075
scenario-Q	5	✓	✓(-5.0)	-	✓	✓	✓(10-2)	2×10^6	518	0.883	0.108	0.001	0.946	0.010	0.037
scenario-R	6	✓	✓(-5.0)	-	✓	✓	✓(10-2)	2×10^6	970	0.943	0.047	0.002	0.977	0.010	0.005

These scenarios are too numerous to completely describe here, but there are several observation which should be made.

- **Scenario-A**: First, the differences between scenario-1 and scenario-A are slight but important. Scenario A has non-trivial probability of transitioning to state $\{N\}$. The transition to exposed state is 4 times greater for state $\{S\}$ than for $\{N\}$, but the resulting percentage of mortality is greater than 4 times. This is due to the difference in rewiring probabilities.

- **Scenario-E**: Comparing scenario-E with scenario-3, one sees that introducing susceptibility, resistance, and constrained treatments reduces the long term difference.

- **Scenario-{G,H,I,J,K}**: All these scenarios are similar except for differences in the isolation policy. The notation (4-1) indicates 2 parts of the isolation policy. The first number indicates the screening depth and the second number indicated one set of threshold parameters. That "-1" represents one choice of thresholds and "-2" indicates a different choice of thresholds.

- **Scenario-L**: The important difference between scenario-L and scenario-H is the use of **treatment policy 1**. Of course, **treatment policy 1** is not the only imaginable policy that reweighs probability of receiving treatment as a function of susceptibility, but it is interesting to ponder whether reweighing in this manner effects the overall behavior. It would appear that using **treatment policy 1** dramatically raises the long term mortality percentage.

- **Scenario-{M,N}**: These two scenarios increase the screening and resistance parameter. Although these scenarios improve overall mortality, the effect is not a beneficial as one might have hoped.

- **Scenario-{O,P,Q,R}**: These 4 scenario attempt to reduce the overall mortality by adjusting the parameters within parameter set 2 and increasing the available treatments.

5.2 Consideration of Cost

The number of treatment per unit time associated with states $\{N, T1, T2\}$ is a proxy for the accumulated costs of treatment. For each of the various scenarios, cost utility function was compiled during the sample run of the process. Scenario L is worth comparing to scenario H. Scenario L has 6.4 percent larger mortality measured over 2000 time periods of the models. Which corresponds to 125,000,000 iterations of the Markov Chain. The associated cost functions may be viewed respectively in figures 8 and 7.

Table 7 gives the mortality percentages for scenario L and H at time period 750 and 2000. The differences between the percentage change of each scenario over the period of time $[750, 2000]$ is "close enough" for our comparison. If one wanted to have the "same" percent change, then different time intervals need to be considered.

Table 8 summarizes the total number of treatments at time periods 750 and 2000. It is not possible to complete the cost trade-off without knowing/modeling the expected cost per each treatment. The point is: policy treatment 1 may raise the overall mortality but except for percent change associated with treatment T2, scenario L has lower cost growth than scenario H. Noted that the percent

Table 7: Mortality Summary

Mortality Comparison			
scenario	time period 750	time period 2000	percent change
scenario-L	0.149	0.173	16.1 %
scenario-H	0.095	0.109	14.7 %
scenario-Ha	0.117	0.128	9.4 %

change between treatment T2 is rather small in comparison to the other difference, but cost of T2 is probably expected to be the largest of all treatment. **Treatment Policy 1** should probably not be used for a variety of reasons. For the purposes of this paper, the cost of mortality is probably higher than any treatment cost, and as such the 6.4% additional deceased is hard to justify.

Although not directly a treatment cost, another factor needing consideration is the screening and restricting of the connections associated with the isolation policy. For example, scenario L closed all long-range connections 36.2% of the time as opposed to 22% of the time for scenario H. The ratio of the percentage of time all long-range connections were closed to the time connections were only screened is 2.5 for scenario L and 2.7 for scenario H. More analysis is needed to determine the degree to which this increase is due to the use of the **treatment policy 1** or just a statistical fluctuation. If more data on expected cost per treatment were available, then a detailed analysis of the trade-offs could be possible.

Table 8: Treatment Summary

Number of Treatments Comparison					
	scenario	time period 750	time period 2000	difference	percent change
N	scenario-L	2,059,939	6,014,017	3,954,078	192%
	scenario-H	2,457,979	8,406,573	5,948,594	242%
	scenario-Ha	1,919,327	6,380,268	4,460,941	232%
T1	scenario-L	1,039,651	2,945,225	1,905,574	183%
	scenario-H	1,683,311	5,106,018	3,422,707	203%
	scenario-Ha	1,101,328	3,773,227	2,671,899	243%
T2	scenario-L	3,830,001	6,119,557	2,289,556	60%
	scenario-H	5,281,653	8,280,921	2,999,268	57%
	scenario-Ha	5,440,359	8,730,350	3,289,991	60%

Scenario Ha is a bridge for comparing scenario L and scenario H. Rather than using the **treatment policy 1** the average value of the policy's effect is used. Scenario L has 4.5% of its overall deceased that can not be explained simply by appealing to the fact that 3 probabilities have on average been reduced by 50%.

6 Moving toward more realistic models

6.1 The H5N1 virus

The Center for Infectious Disease, Research and Policy at the University of Minnesota (CIDRAP) has given some initial feedback about this model concerning use for the H5N1 virus. The comments have been illuminating and useful. A couple of simple modifications to the software allows easy exploration.

- Concerning the use of the resistance constraint, there is currently no evidence supporting immune boosting in regards to the H5N1 virus.
- Concerning susceptibility of the population: Although there is variability in contracting the H5N1 virus in the cases to date, it seems unlikely to be linked to a genetic disposition or a compromised immune systems.

Resulting from this feedback a few scenarios were run with resistance constraint turned off and with a simple modification to the software. The code that implements the susceptibility policy was changed to not modulate the probabilities associated with treatment. It leaves unchanged the code portions that modify the transition from exposure to infected and infected to deceased. Table 9 tabulates the simulations associated with the modifications relating our discussions concerning the H5N1 virus.

Table 9: Summary of Scenarios from Software Modification for H5N1

	parameter set	modified susceptibility	resistance	treatment policy 1	treatment policy 2	quarantine	isolation	max treatments	time period at 30%	healthy at 30%	sick at 30%	mortality at 30%	healthy at 2k	sick at 2k	mortality at 2k
scenario-S	2	✓	-	-	✓	✓	✓	10^6	266	0.919	0.072	0.001	0.533	0.017	0.443
scenario-T	2	✓	-	✓	✓	✓	✓	10^6	224	0.864	0.125	0.004	0.459	0.029	0.504

6.2 Conclusion

In this paper, a Markov Chain model was presented in which the states of the Markov Chain represent a large collection of sites representing a population. The Markov Chain's transitions represent interactions between sites or entities within the model. The probability parameters are really conditional probability marginals (*i.e* given a site x in a given state and some additional knowledge of the environment, what is the distribution of events). All modeling environments should be simpler than the problem needing modeling. Even in this simplified model the behavior is nontrivial and complex. In order to move toward a large

scale and more realistic model, one first needs to compute (to the best of ones ability) the probability marginals that will contribute to and be the propagation mechanism for updating the state transitions.

In addition, more realistic graphical relationships should be incorporated, because the relationships among people are not static. However, the Watts-Strogatz is a good place to start. Although a complete representation of peoples connectedness is not required, some degree of hierarchy of geography into towns, cities, states, and countries should be incorporated. Then a more realist modeling of long-range connections should be implemented to incorporate how people come within contact of one another.

Currently there is essential no real world time scale set in the model. Part of making the model more realistic should be determining such a time scale. In order to do this, it is preferred to use real data. For example, the World Health Organization (WHO) posts on their website information about the H5N1 virus. As of July 30 2007, there have been 317 laboratory-confirmed cases of the H5N1 virus in humans with a total of 191 deaths (a sobering 60.25 percent). It would be helpful to use the "time to death" for the 191 victims[2]. This statistic would help in setting the time scale for the Markov Chain. Currently, the probabilities are per update within the model. Consider that one could track sites that have become infected and how many transitions are required for the site to transition out of the infected state. One could then fit the data from the observed data of the real world.

Furthermore, such a model should be made as realistic as possible before doing analysis of fluctuations, because such an analysis requires a fair amount of computation.

The model does not currently incorporate any notion of mass hysteria. If the infection and death rates followed these trajectories for even a short time period, most people would not leave their homes. Which in effect is constraining local interactions as well as long range interactions and quarantining the pandemic. Although, one can introduce into the model mass hysteria, it would beneficial to do so after get the graphical relationships more realistic, because this leads to having the branching mechanism of the pandemic more realistic.

Although this paper has not focused entirely on the H5N1 virus, this virus was a motivating factor for the work. Given the number of unanswered questions about this virus, it would be valuable to enhance models that could explore the effectiveness that policies might have on containing a H5N1 pandemic.

Acknowledgment

The author would like to thank Max Rudolph for helpful input during the development of this work and also Aaron Desmond and Nick Kelly from CIDRAP.

[2]CIDRAP has recently directed us on how we may locate some of the information required to make the model more realistic

References

[1] Thomas M. Liggett, *Interacting Particle Systems*, Springer-Verlag, 1985.

[2] Thomas M. Liggett, *Stochastic Interacting Systems: Contact, Voter, and Exclusion Processes*, Springer-Verlag, 1999

[3] Watts D.J. and Strogatz S.H., *Nature*, **393** (1998) 440.

[4] Stephen M. Krone, *The two-stage contact process*, Annals of Applied Probability, **Vol. 9, Number 2** (1999) p.331-351

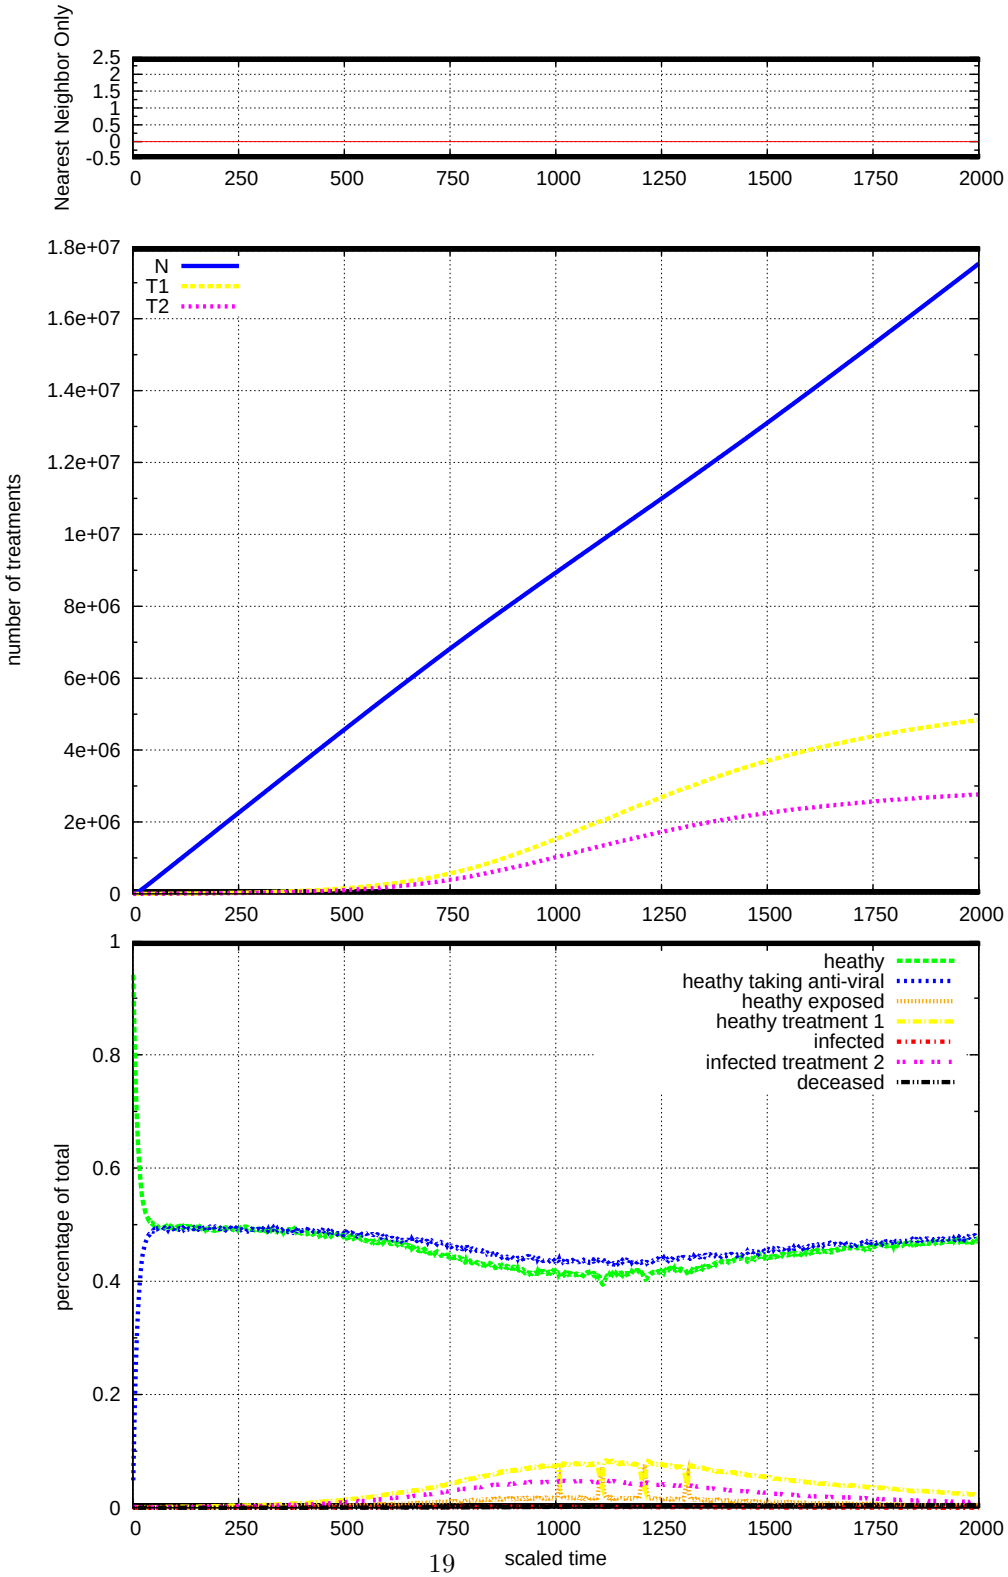

Figure 5: scenario R

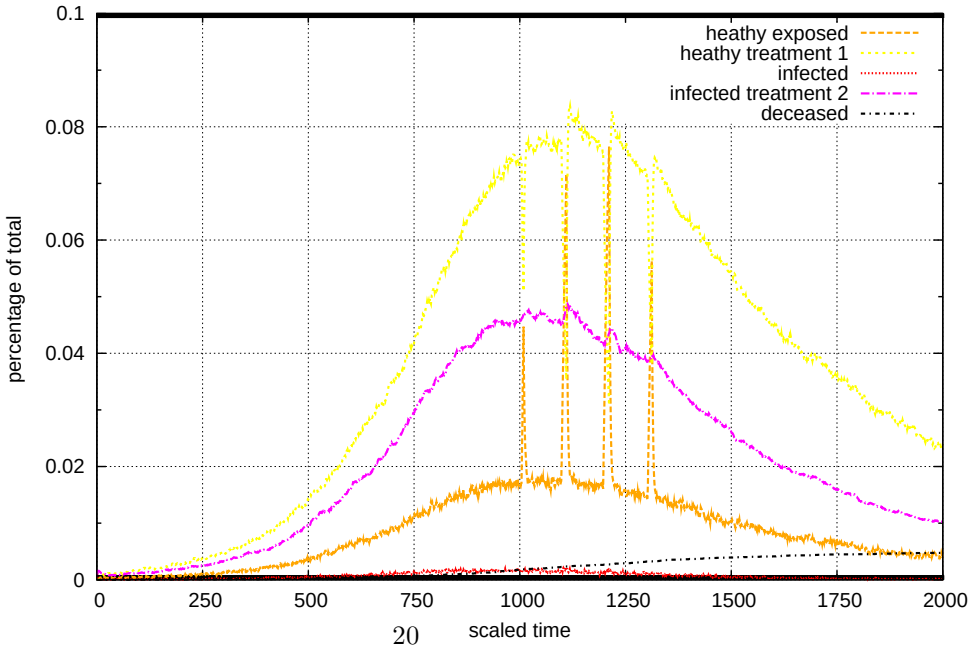

Figure 6: scenario R (expanded)

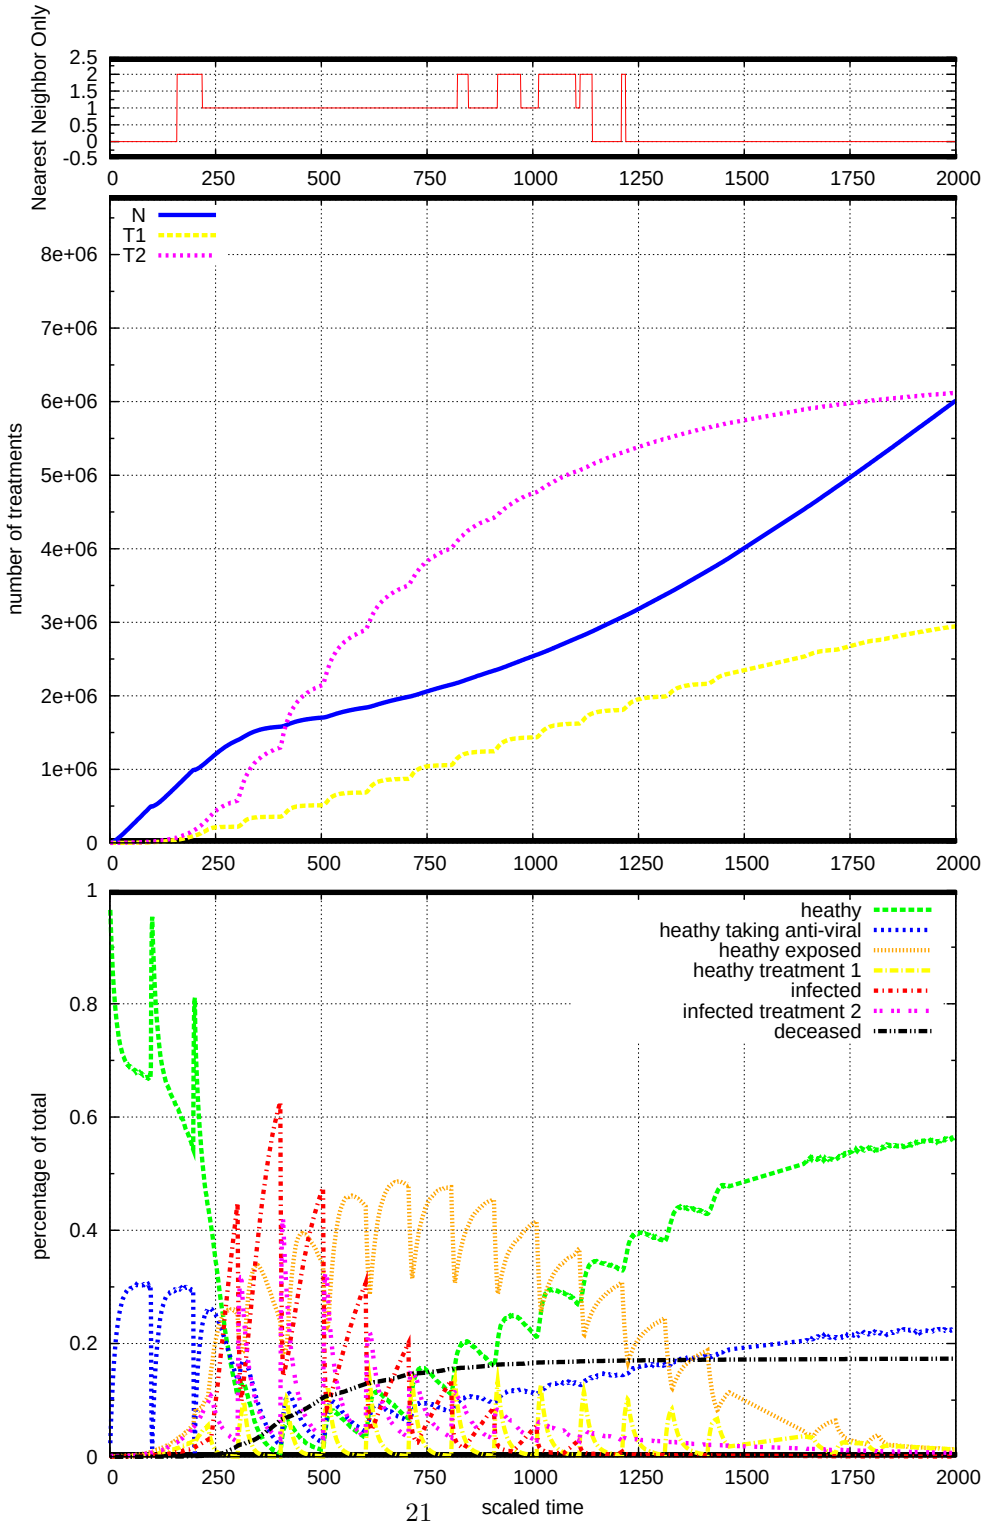

Figure 7: scenario L

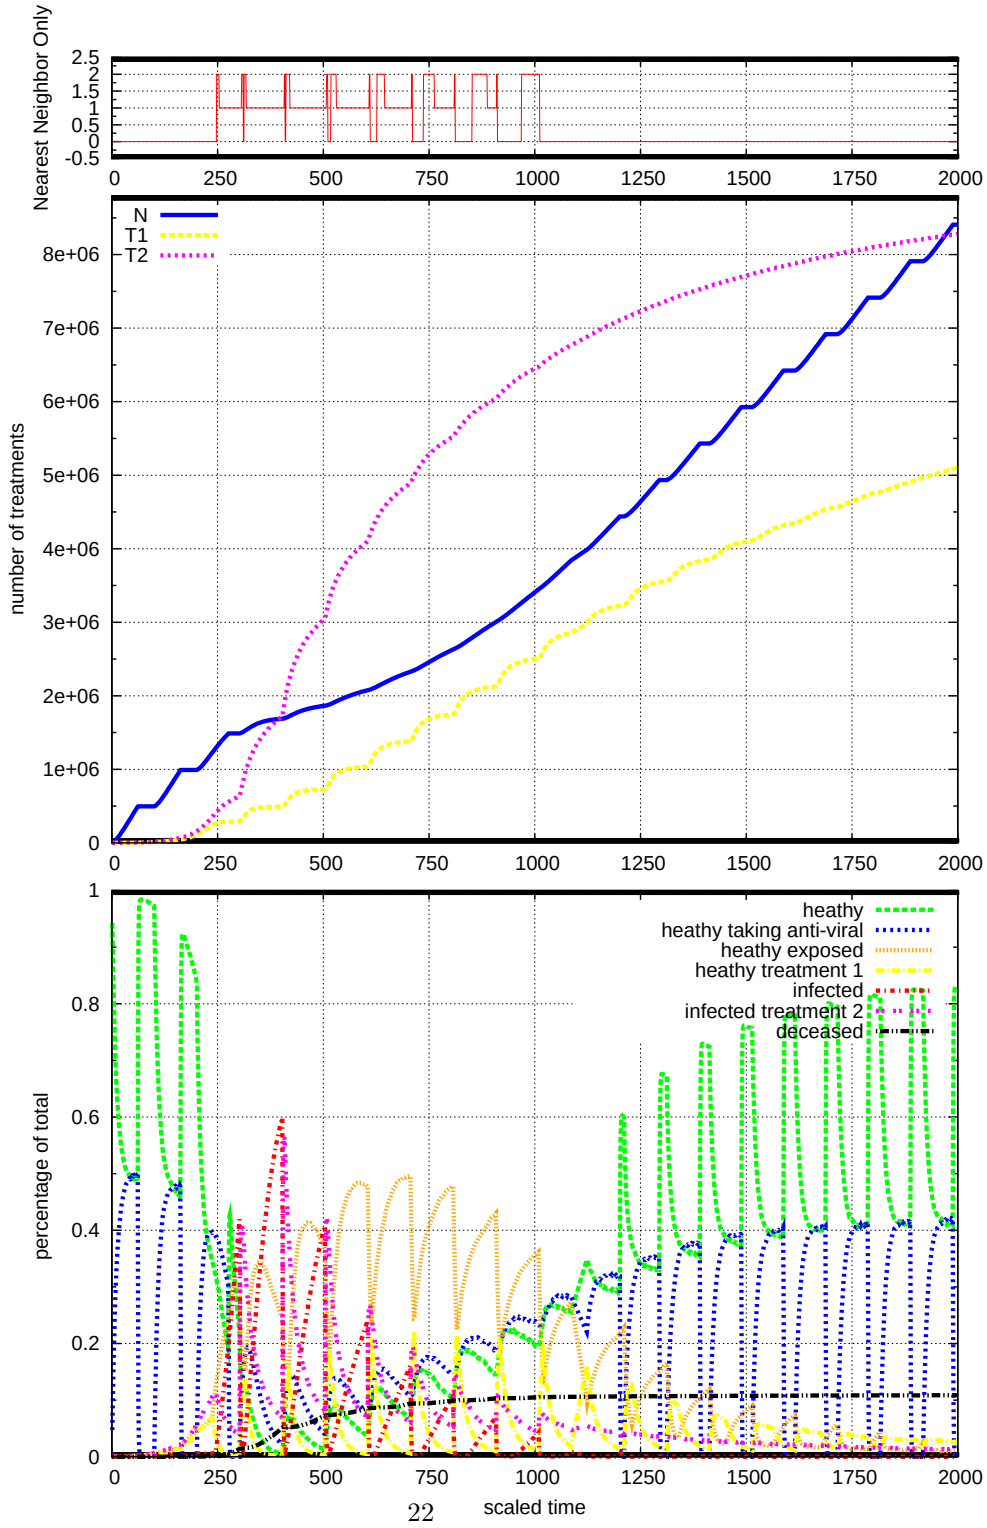

Figure 8: scenario H

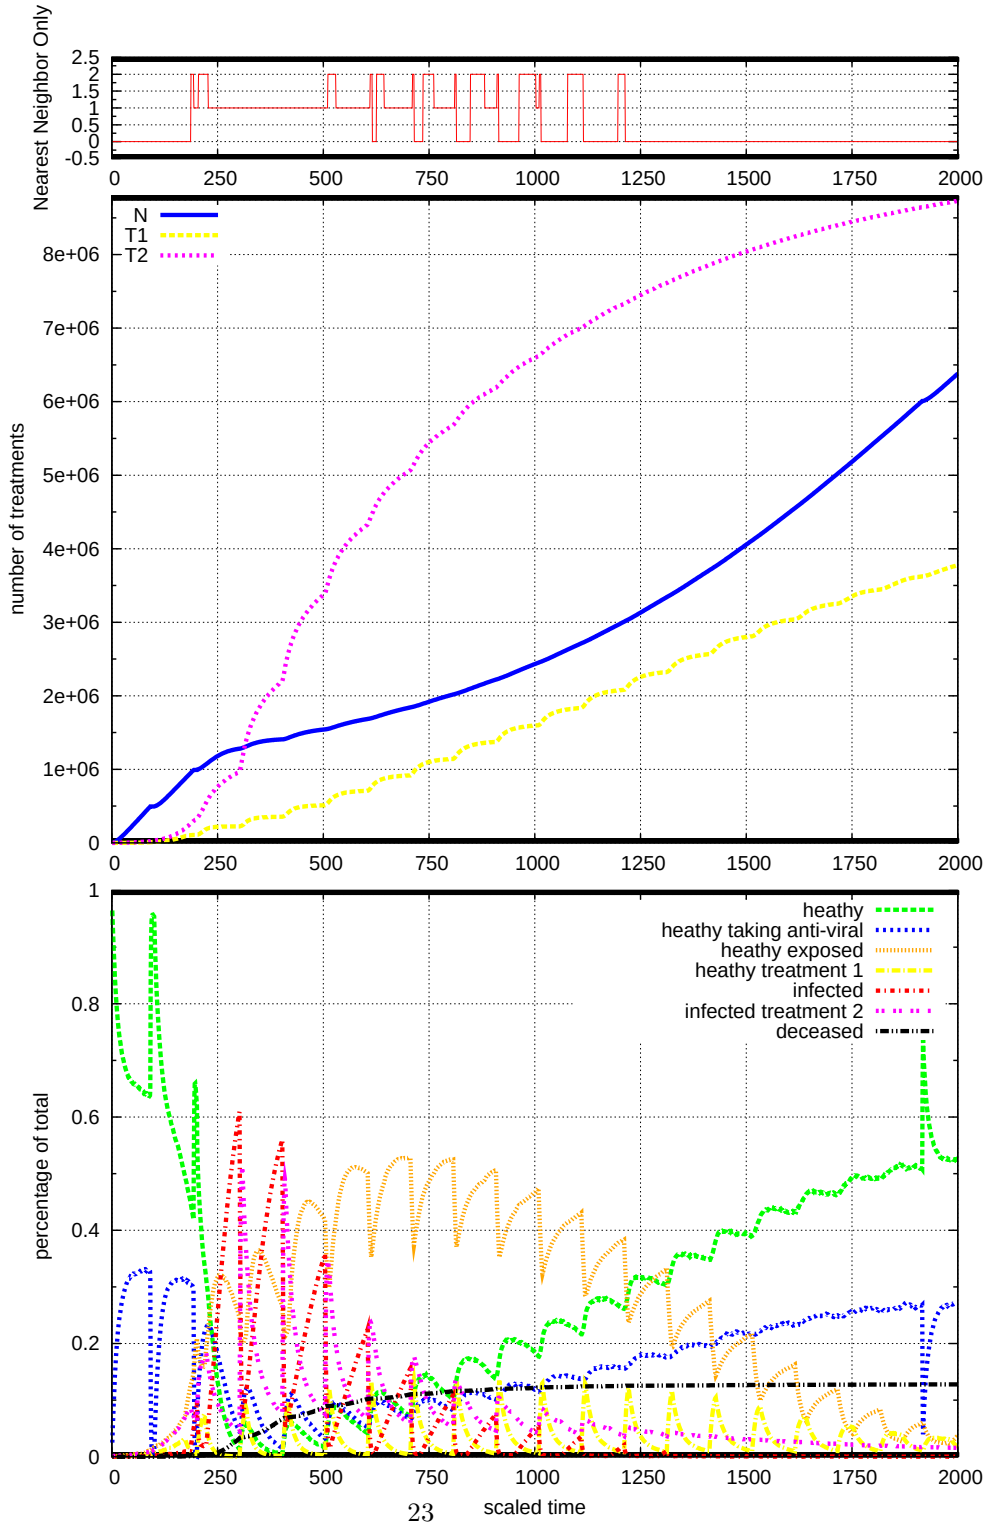

Figure 9: scenario Ha

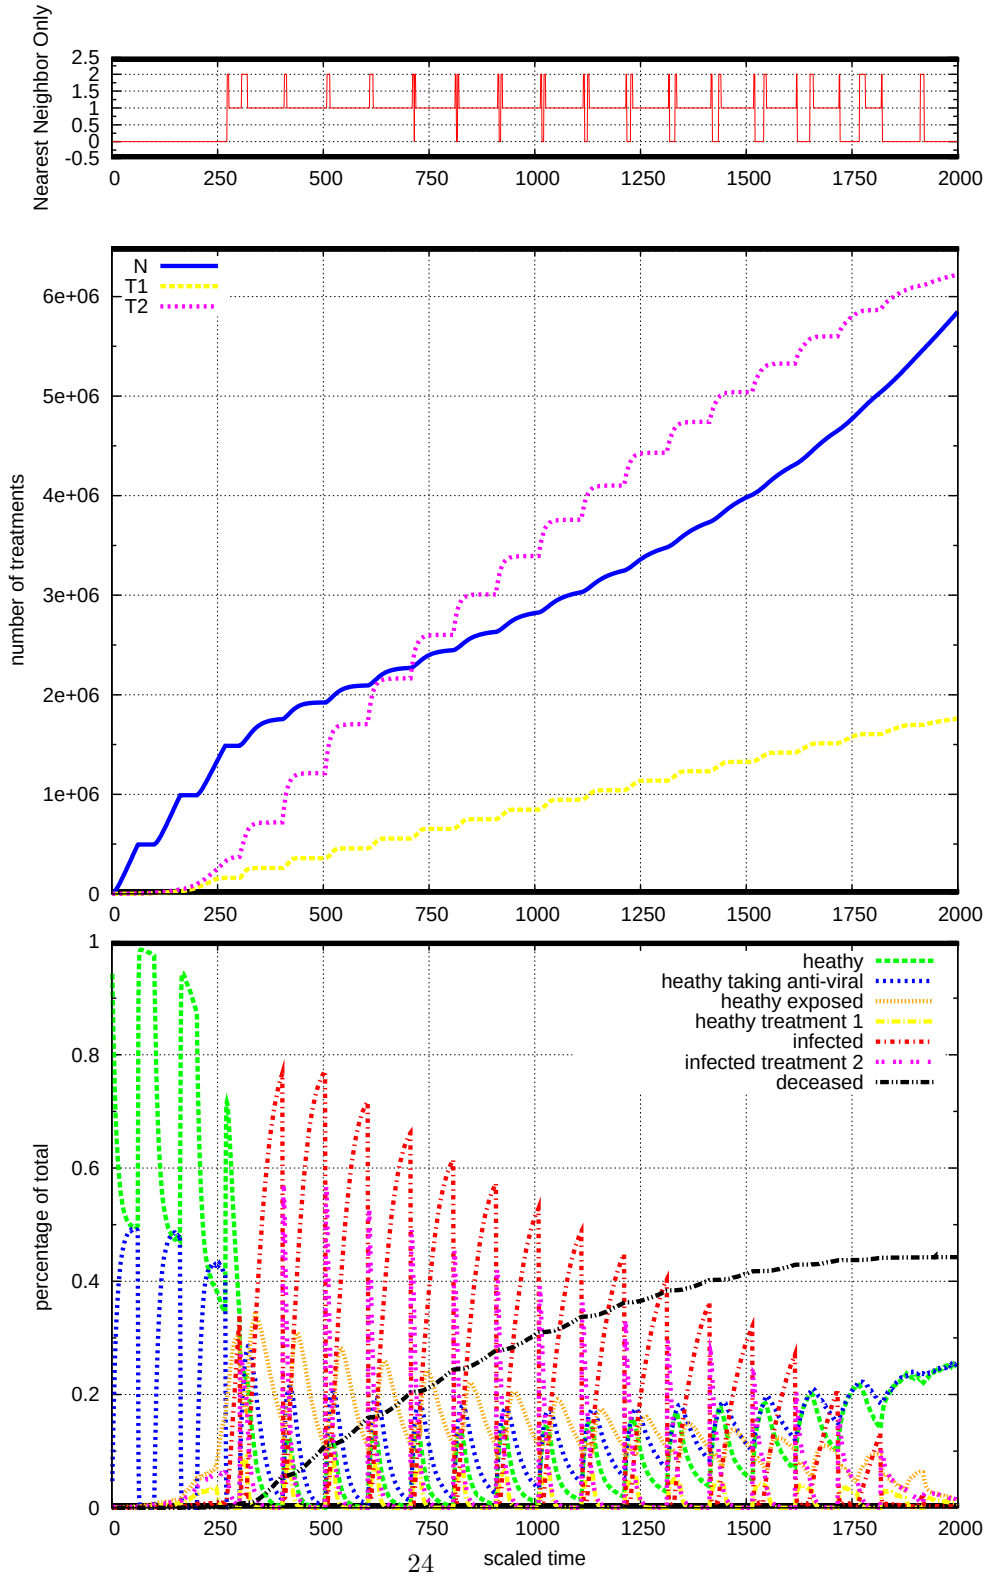

Figure 10: scenario S

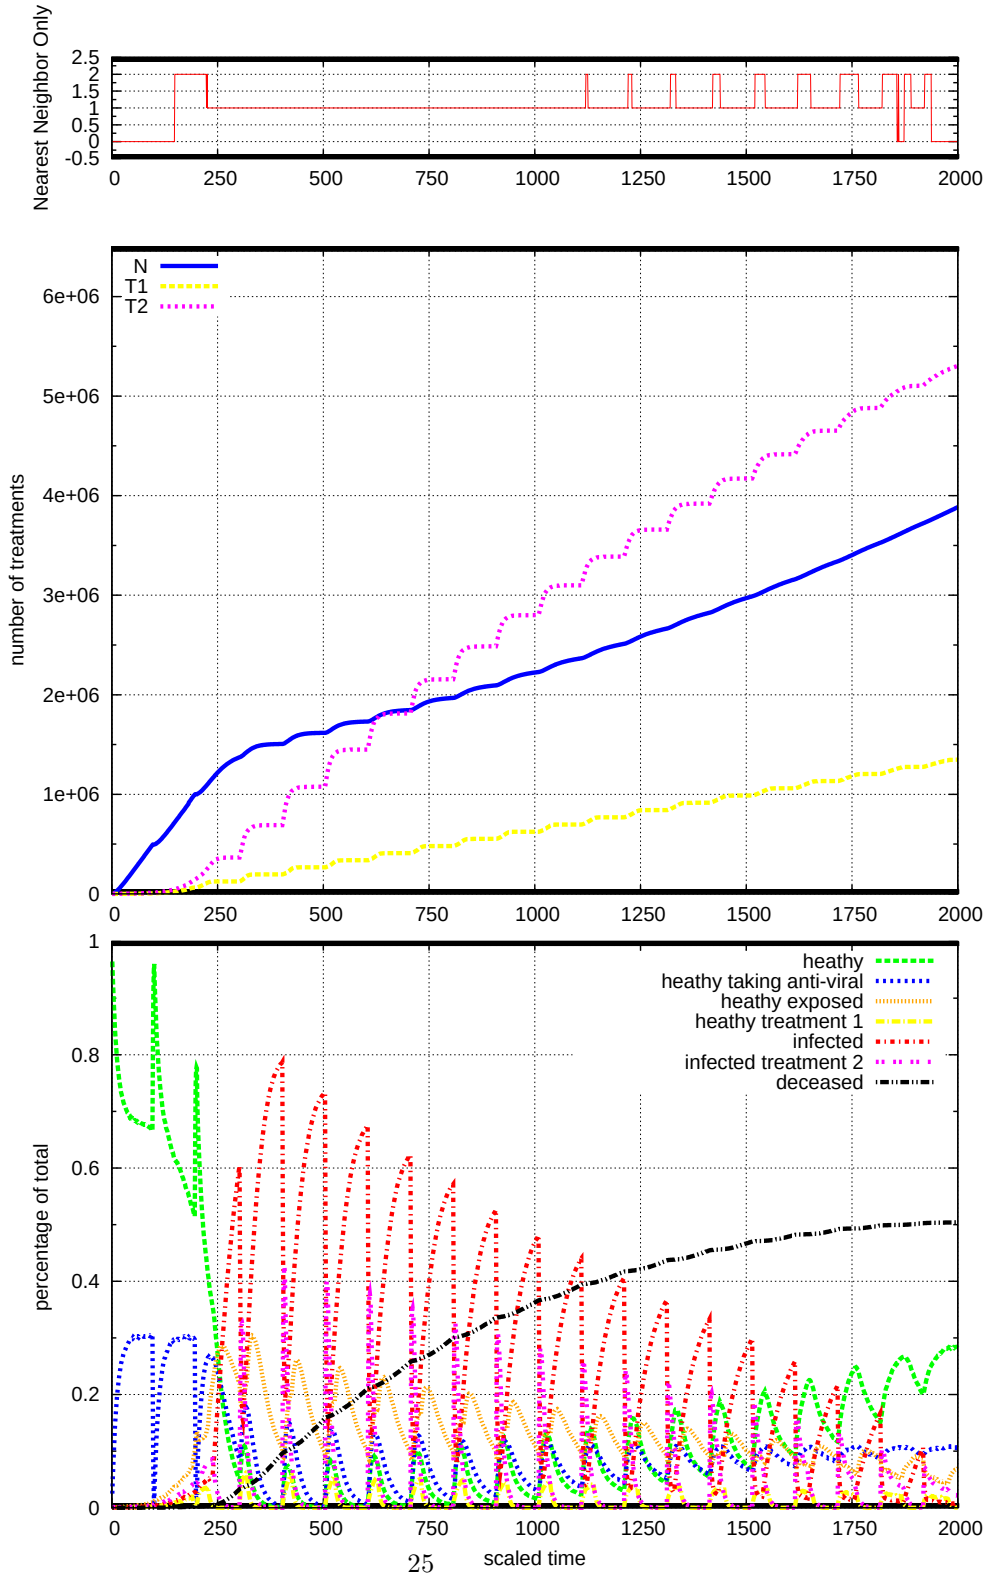

Figure 11: scenario T