Agent-based Modeling as a Decision Making Tool:
How to Halt a Smallpox Epidemic[1]

Zoltán Toroczkai
Center for Nonlinear Studies,,
Theoretical Division, Los Alamos National Laboratory, NM, 87544, USA

Stephen Eubank
Virginia Bioinformatics Institute,
Virginia Tech, Blacksburg, VA, 24061, USA

Science has made considerable inroads in quantitatively characterizing, understanding and controlling non-living systems. We are rather familiar with systems of physics and chemistry, ranging from elementary particles, atoms, molecules, to proteins, polymers, fluids, solids, etc. These are systems of interacting particles with well defined physical interactions, and their properties are described by the known laws of physics and chemistry. Most importantly, their behavior (at least statistically) is *reproducible* given the same initial conditions. There are, however, other types of ubiquitous systems, surrounding us, namely those that involve living entities over which we hardly have any quantitative understanding, neither on the individual nor on the collective level. In the following we will refer to collectives of living entities as "agent-based" or "agent" systems, in order to distinguish them from classical particle systems of inanimate objects. Although there have been intense efforts to study these systems, a generally accepted unifying framework is largely missing. Nevertheless, understanding, and ultimately controlling the behavior of such systems is a subject of extreme importance with applications ranging from biology, through social, to political sciences, and economics.

[1] This paper is based on (Eubank et.al. 2004).

Ultimately, such an understanding can be used to design agent systems like of robots or rovers, which collectively can perform tasks that would normally be prohibitive for humans. Examples include deep-water rescue missions, mine field mapping, distributed sensor networks (including military uses), and rovers for extraterrestrial explorations. In spite that there is no unifying understanding behind these systems, some control over the behavior of these systems can be achieved via ***agent-based modeling*** tools. The idea behind agent-based modeling is rather simple: build a computer model of the agent-system under observation using a ***bottom-up*** approach by trying to mimic as much detail as possible. Building an agent model, however, is a rather expensive task: it involves the phases of 1) *data collection*, 2) *model building*, 3) *exploiting the model, collecting statistics,* and ultimately 4) *validation* which normally means comparing the output of the model with additional observations on the real system. The agent-models I will briefly mention in this paper took about 9 years to develop at Los Alamos National Laboratory. However, once developed, the framework can be used to simulate many similar circumstances and used as a tool with predictive capability.

Some properties of Agent-systems

There are at least two major differences from classical particle systems that make agent-based systems hard to describe and understand, within a unified approach. First, the "particle" or agent is a more complex entity than what can be represented by a simple function such as a Hamiltonian function of a classical system (e.g., a spin system). Secondly, the interaction topology, namely the prescription of which particle interacts with which others, is in general a complex and *dynamic* graph (network), unlike the regular lattices of crystalline solids or the continuous spaces of fluid dynamics. In many cases the notion of "locality" becomes elusive in these networks, such as in social networks, where physical or spatial locality of the agents can have little indication about the social "distances" and social interactions among the agents.

In order to illustrate the more complex structure of the "particle" or agent, and its consequences, in the following we will use traffic, namely people (agents) driving on a highway, as an example, but the statements below are generally applicable.

An agent is an entity with the following set of qualities: *(1)* There is a set of variables $\boldsymbol{x}$ describing the *state of the agent*. Position on the road, speed, health state, etc. The corresponding state space is $\boldsymbol{X}$. *(2)* There is a set of variables $\boldsymbol{z}$, describing the ***perceived state of the environment***, $\boldsymbol{Z}$. The environment includes other agents if there are any. For example, level of congestion, state/quality of the road, weather conditions, etc. *(3)* There is a set of allowable *actions* (output space), $\boldsymbol{A}$. Swerve, brake, accelerate, etc. *(4)* There is a set of *strategies*, which are functions s: $\boldsymbol{(Z\times X)^t} \to \boldsymbol{A}$, that summon an action to a given external perception, current state of the agent and history up to time $\boldsymbol{t}$. These are "ways of reasoning" for the agent. One might think of it also as a *behavioral input space*. For example, depending on the age, background and other factors, some drivers will choose to brake and others to swerve to avoid an accident. Social studies and surveys will supply here valuable statistical inputs, since they can collect data of the type "agents with n years of driving experience who are between ages a_1 and a_2, will swerve *f%* of the time and break *g%* of the time", etc. *(5)* There is a set of *utility variables*, $\boldsymbol{u}\in\boldsymbol{U}$. Time to destination, number of accidents, amount of speeding tickets, etc. *(6)* There is a *multivariate objective function*: $\boldsymbol{F{:}U}$®$\boldsymbol{R}^m$, which might include constraints ("rules"). E.g. the agent has to stay on the road. The analogous version in physics is called action. The agent is *trying is to optimize* this objective function. For example, it is trying to minimize the time to destination, to keep at zero the number of accidents it is involved in, etc. Compared to the particles of classical systems, agents usually have *memory* of the past, which they can use to change/evolve their strategies, a process called *learning*. The other important aspect is that agents can perform *reasoning and planning*, which basically entails a search by the agents of the choice tree and assigning likelihood weights and payoffs given what the other agents might choose. In realistic situations that involves hundreds of agents (such as markets, or traffic) long-term planning and reasoning is impossible to perform, due to the combinatorial explosion of the possibilities and also to the fact that not all information is available to any agent. In this case agents try to follow and exploit patterns in the response of the surrounding environment to their (past) actions, using these patterns to discriminate among their strategies, reinforcing some while diminishing others (*reinforcement learning*). This leads to bounded rationality like behavior, and introduces de-correlations between the strategies, and for that reason, it actually can be make statistical modeling a plausible feat. In the following we will briefly describe two large-scale agent-based models developed at Los Alamos National Laboratory (LANL), a traffic simulator TRANSIMS and an epidemics simulator EPISIMS.

TRANSIMS

The Transportation Analysis and Simulation System (TRANSIMS) is an agent-based model of traffic in a particular urban area (the first model was for Portland, OR, USA). TRANSIMS conceptually decomposes the transportation planning task using three different time scales. A large time-scale associated with land use and demographic distribution is employed to create *activities* for travelers (there several such activity categories such as work, shopping, entertainment, school, etc.). Activity information typically consists of requests that travelers be at a certain location at a specified time, and includes information on travel modes available to the traveler.

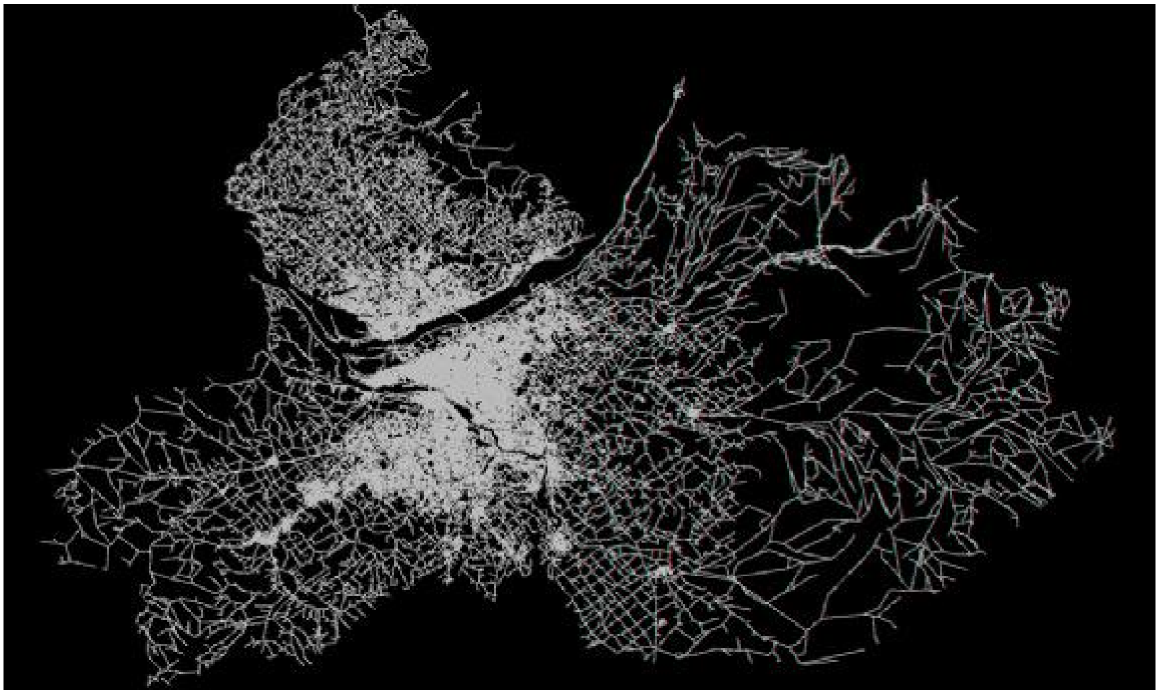

Figure 1. Portland roadway network.

This is achieved by creating a synthetic population and endowing it with demographics matching the joint distributions given in census data. The synthetic households are built by also using survey data from several thousands of households, which are observations made on the daily activity patterns of each individual in the household. These activity patterns are associated with synthetic households with similar demographics. The locations for various activities are estimated taking into account observed land use patterns, travel times and dollar costs of transportation. The intermediate time-scale consists of assigning routes and trip-chains to satisfy the activity requests. To do this, the estimated locations are fed into a routing algorithm to find minimum cost paths through the transportation infrastructure consistent with constraints on mode choice (Barret et.al. 2001, 2002) . An example constraint might be: “walk to a transit stop, take transit to work using no more than 2 transfers and no more than 1 bus”.

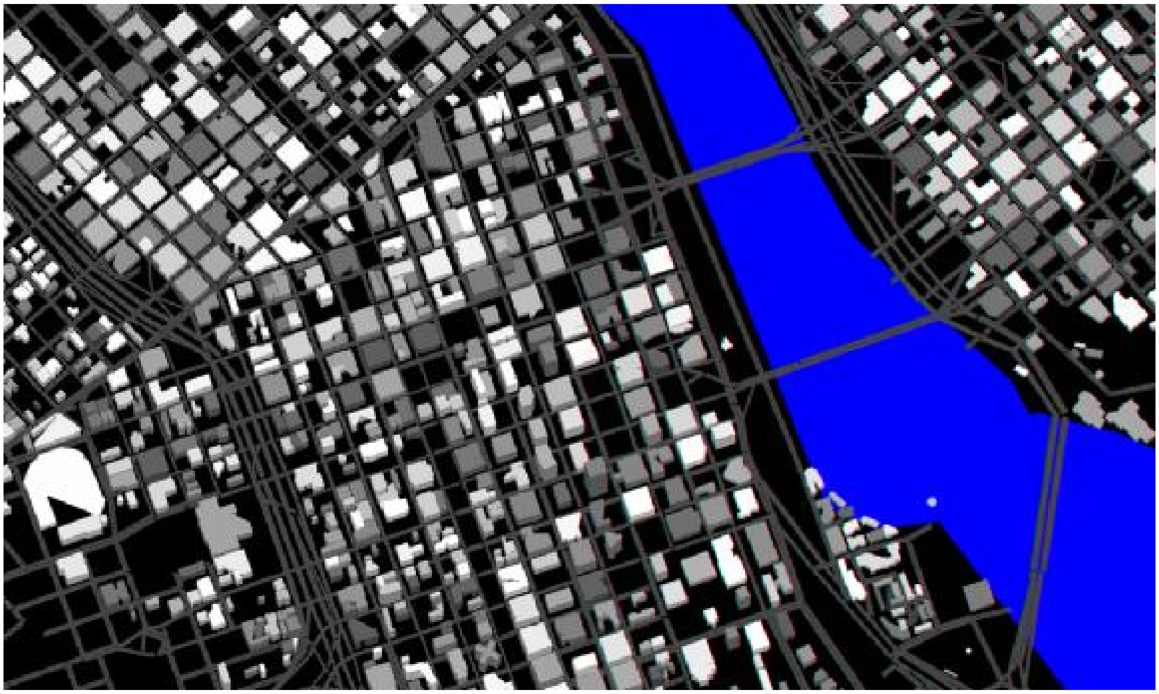

Figure 2. Locations and roads, downtown Portland.

Finally, a very short time-scale is associated with the actual execution of trip plans in the road network. This is done by a cellular automata simulation through a very detailed representation of the urban transportation network. The simulation is in effect a way to resolve the traffic induced when everyone tries to execute their plans simultaneously. The simulation resolves distances down to 7.5 meters and times down to 1 second. It provides an updated estimate of time-dependent travel times for each edge in the network, including the effects of congestion, which it feeds then to the router and location estimation algorithms, which produce new plans. This feedback process continues iteratively until it converges to a “quasi - steady state” in which no one can find a better path in the context of everyone else's decisions. The resulting traffic patterns compare well to observed traffic. The entire process estimates the demand on a transportation network using census data, land usage data, and activity surveys. More information and including availability of the software can be obtained from http://transims.tsasa.lanl.gov/ .

Figure 3. The TRANSIMS microsimulation, downtown Portland.

EPISIMS

Although the TRANSIMS agent-based model is indeed useful for urban planners and traffic analysts, here we would like to focus on one of its applications, namely in the field of epidemics. Diseases such as colds, flu, smallpox or SARS, are transmitted through air between two agents, if they spend long enough time in the proximity of each other, or in a building with closed air ventilation. This means, that we can assume that the majority of the infections will take place in *locations,* like offices, shopping malls, entertainment centers, mass transit units (metros, trams, etc.). Thus, by tracking the people in our TRANSIMS virtual city, we can generate ***a bipartite contact network, or graph***, formed by two types of nodes, namely people nodes and locations nodes. If a person ***p*** enters a location ***l***, then there is an edge drawn between that person and the corresponding location node on this graph. This edge has a time-stamp associated to it representing the union of distinct time intervals that person ***p*** was located at location ***l*** during the day. If two people nodes p_1 and p_2 have an incident edge onto the same location node ***l***, the common intersection of the two time-stamps will tell us the total time the two people spent in the proximity of each other during the day, thus enabling us to determine the possibility of an airborne infection. In the case of Portland, there are about 1.6 million people nodes and 181,000 location nodes and over 6 million edges between them. These are huge graphs, representing considerable challenges for the measurement of its properties. This dynamic contact graph allows us to simulate different disease spread scenarios and test the sensitivities of the epidemics to disease parameters, such as incubation period, person-to-person infection rates, influence of age structure, activity patterns, etc. The epidemiological study tool thus generated is called EPISIMS, which was also developed at LANL, and can be used to aid decision making and planning for example for smallpox outbreaks.

In the following we will briefly summarize some the findings for the case of smallpox spread using the Portland data. For more details on EPISIMS, see (Eubank et.al. 2004).

1) Vaccinating a person means taking it out from the contact graph together with its incident links. An efficient vaccination strategy will remove the smallest subset of nodes

such that the resulting graph is made of many small disconnected pieces, thus forbidding disease spread across the population. The smallpox vaccine is not entirely harmless, it can cause a disease called vaccinia in some people (it can be fatal in some cases), and thus one would like to minimize its impact on the population, making the Mass Vaccination strategy (vaccinate everyone) proposed by (Kaplan et.al. 2002) a last resort. Studying the projection of the bipartite graph onto people nodes, we have found that it is extremely interconnected with very high expansion properties, being able to shatter the graph only if we vaccinated everyone with 10 or more contacts during the day, which effectively meant a Mass Vaccination. Ultimately what we have found to work on such graphs was to perform vaccination of people that frequently took *long-range trips* across the city, corresponding to shortcuts in the network (Watts and Strogatz 1998) making it a more local graph with a larger diameter. This in case of an outbreak allows an effective use of the ring-strategy for quarantining and further vaccinations to stop disease spread.

2) The most crucial parameter in containing epidemics spread is the delay in reaction time. Assume that sensors were developed that can perform an online analysis of the pathogens in the air. Then the question is, in what locations to place them such that they will most effectively (the earliest) capture the onset of the outbreak. Due to a particular, so-called scale-free property (Albert and Barabási 2002) of the locations projection of the bipartite network, one can pinpoint a small set of locations (the so-called dominating set, about 10% of all locations) which would cover a large fraction of people (about 90%) and thus it is the optimal set for detector placement, or for distribution purposes (of prophylactics and supplies). Figure 4 shows the evolution of epidemics after a covert introduction in a particular location (at a university) when the disease is left to spread

(left side) compared to using a targeted contact tracing and quarantining strategy (right).

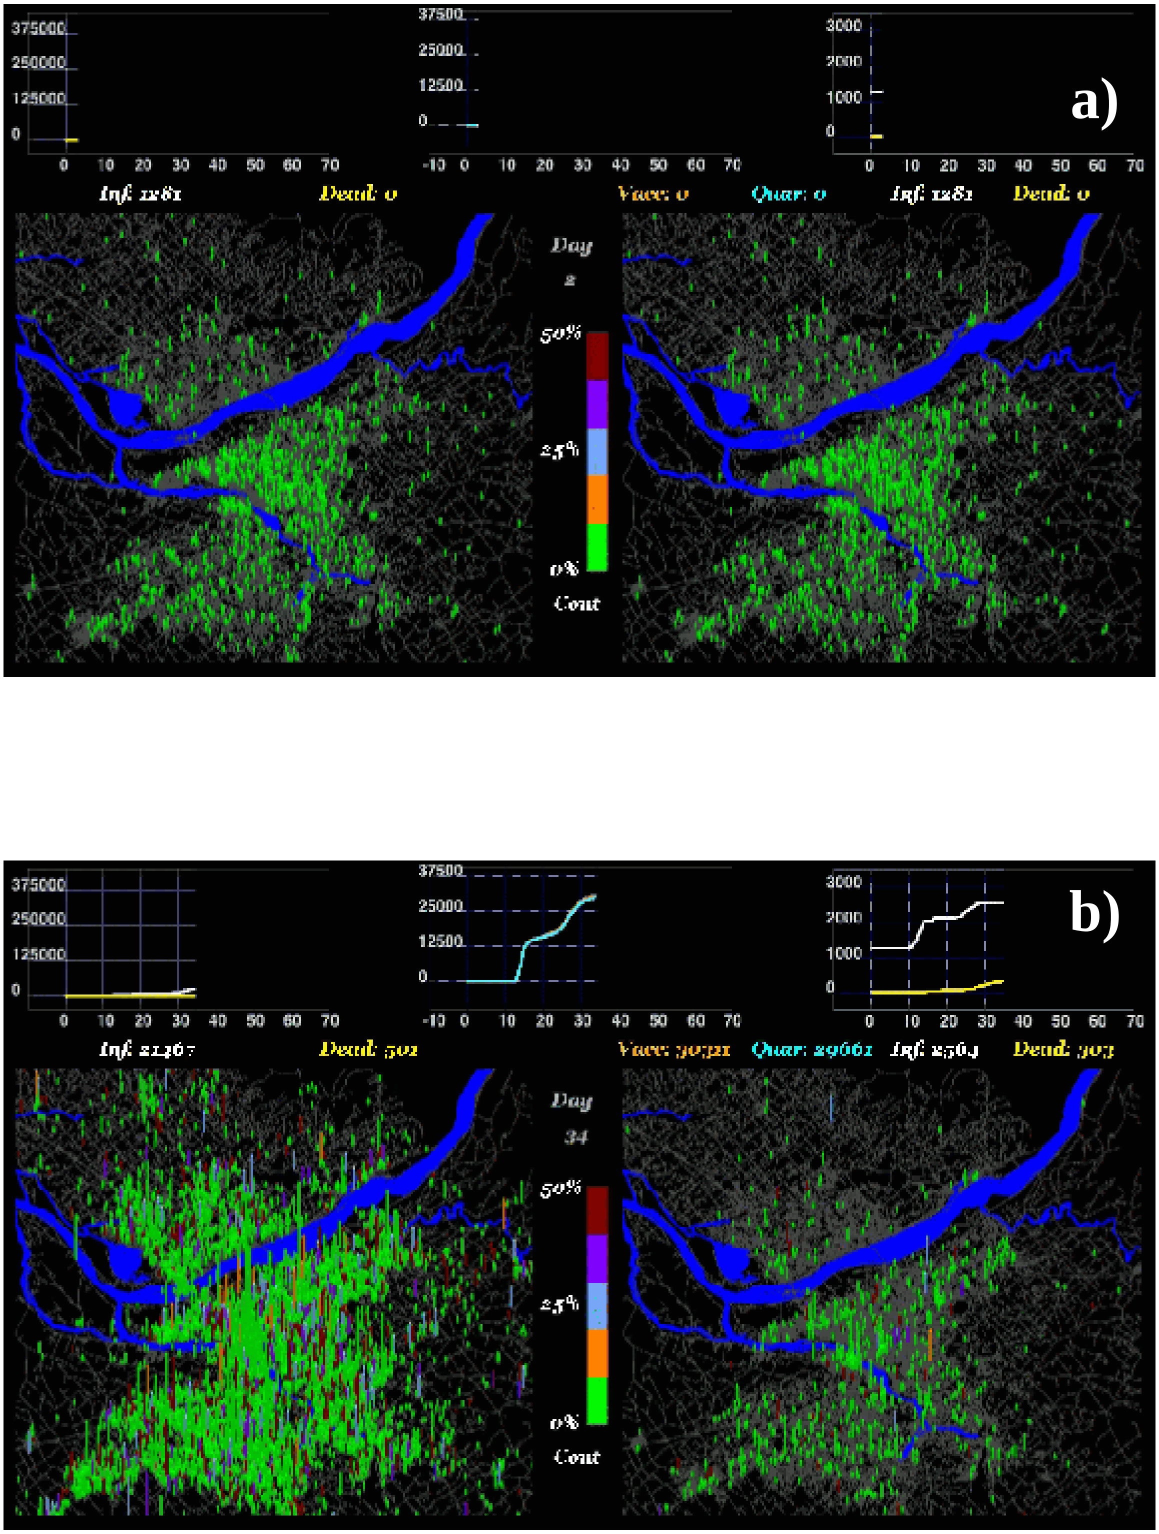

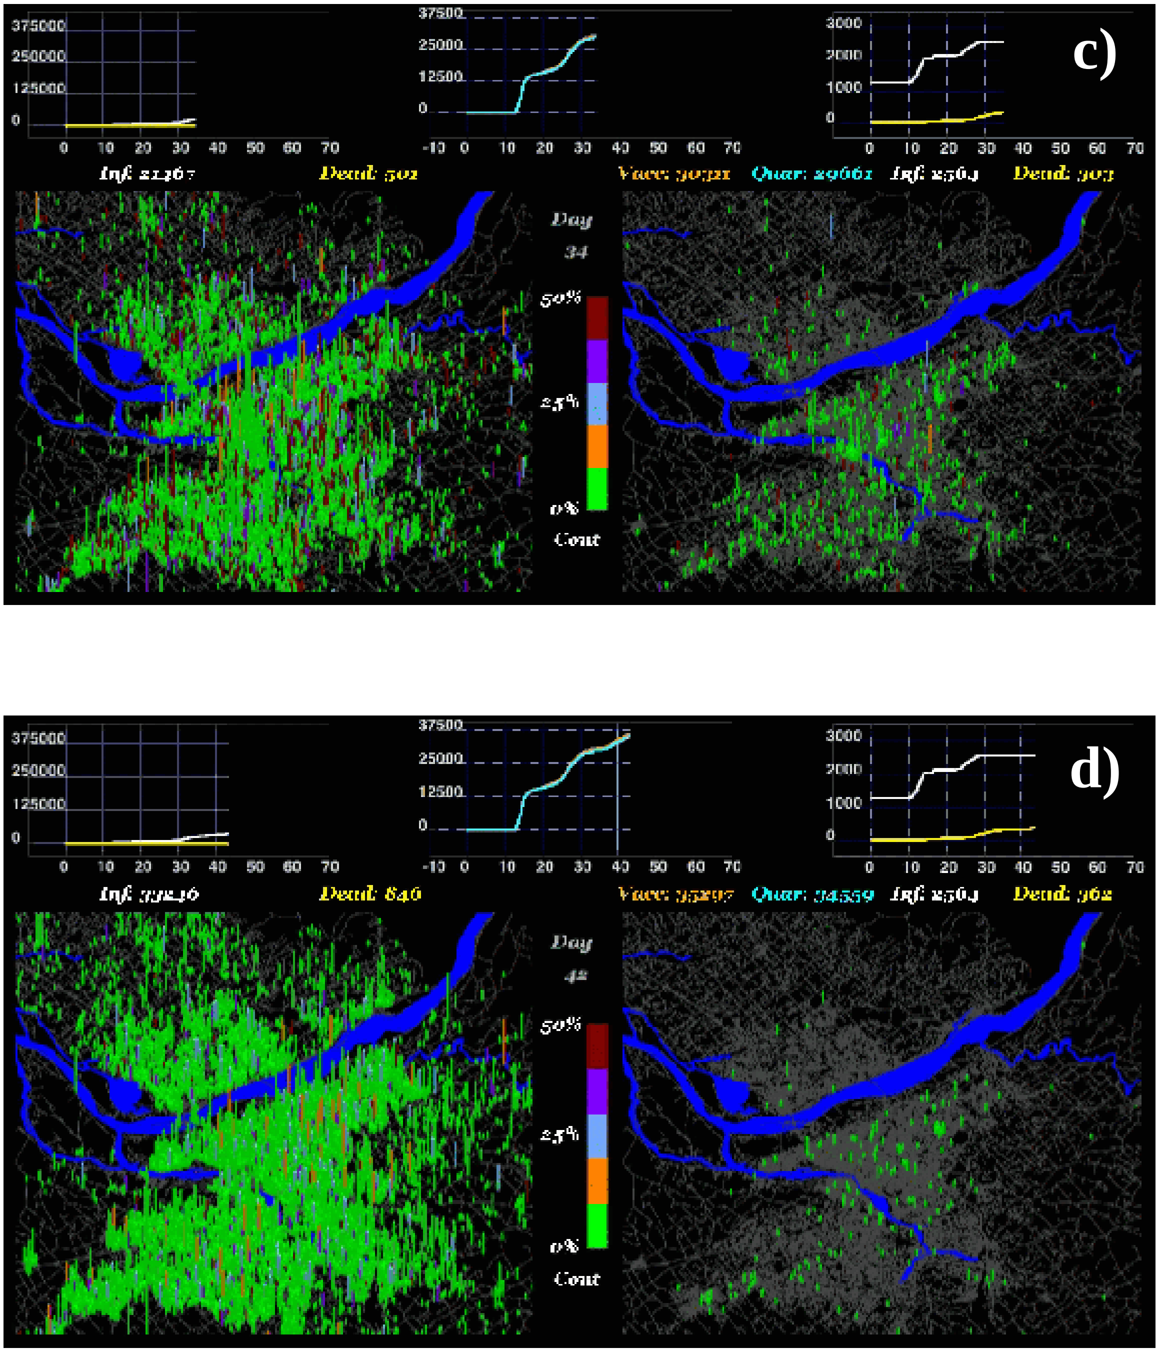

Figure 4. Comparison of a baseline case (on the left), with a targeted vaccination and quarantining strategy (right). The bars represent number infected at each location, and the color represents the fraction of infected people who are infectious. The inserts display the cumulative number of people infected and dead as a function of time, and for the targeted response, the number vaccinated and quarantined. Note the different scales between the leftmost and rightmost inserts. Reprinted from (Eubank et.al. 2004).

ACKNOWLEDGEMENTS

This work was done in collaboration with the TRANSIMS and EPISIMS teams. EPISIMS was supported by the National Infrastructure Simulation and Analysis (NISAC) program at the Department of Homeland Security. The TRANSIMS project was funded by the Department of Transportation. Z.T. was supported by the Department of Energy.

REFERENCES

1. Eubank, S., H. Guclu, V.S.A. Kumar, M.V. Marathe, A. Srinivasan, Z. Toroczkai, N. Wang. 2004. Modelling disease outbreaks in realistic urban social networks. *Nature*, **429**: 180-184.
2. Barrett C., R. Jacob and M.V. Marathe, Formal Language Constrained Path Problems. 2001. *SIAM J. Computing*, **30**(3): 809-837.
3. Barrett C., K. Bisset, R. Jacob, G. Konjevod and M.V. Marathe, An Experimental Analysis of a Routing Algorithm for Realistic Transportation Networks. 2002. in *Proc. European Symposium on Algorithms (ESA)*, Rome; Technical Report No. LA-UR-02-2427, Los Alamos National Laboratory.
4. Kaplan, E., D. Craft, and L. Wein. 2002. Emergency response to a smallpox attack: the case for mass vaccination. *Proc. Natl. Acad. Sci. (USA)* **99**: 10935-10940.
5. Halloran, M., Longini, I. M. Jr., Nizam, A., & Yang, Y. Containing bioterrorist smallpox. *Science* **298**, 1428-1432 (2002).
6. Watts, D. & Strogatz, S. Collective dynamics of small-world networks. *Nature* **393**, 440 (1998).
7. Albert, R. and A.-L. Barabási. Statistical mechanics of complex networks. *Rev. Mod. Phys.* **74**, 47 (2002).