Simulation of Complex Systems

Homework 1: Disease Spreading

Assessment date: November 14 2014

In this exercise we are going to implement a simple agent-based model for studying disease spreading. Agent-based models are a natural next step from the simple ODE approach. If you want a glimpse of what is possible, read the Nature opinion "Modelling to contain pandemics" by Epstein.

We will consider the three-compartment model known as SIR, where each individual is either Susceptible to the disease, Infected, or has Recovered and is immune. Infected individuals infect susceptibles they meet with some rate β and recover with some rate γ. In a simple ODE or PDE-version of the model only the ratio $R_0 = \beta/\gamma$ matters for the behavior of the model. This is not the case in our agent-based one.

To make the movement of the agents as simple as possible, we will use a random walk on a square grid (lattice), where every time step each agent either sits still with some probability $1-d$, or moves to a random neighboring tile (site) with probability d, where d sets the diffusion rate.

Check for infection when agents of the susceptible and infected types land on the same lattice site. Each infected should have a probability β of infecting each susceptible at its current site and a probability γ of recovering, both for each time step. To make the simulation scalable (we want to be able to look at at least 1000 agents), don't check the position of every agent against everyone elses (this scales as the square of the number of agents). Instead, do something along the lines of keeping a list, corresponding to the lattice, that keeps track of which agents (if any) are at a given site. This scales linearly in the size of the lattice and the density of agents.

Examination: Work in your assigned groups. During lab hour, either 14/11 or the extra lab 17/12, you should together demonstrate your results to a tutor in the way indicated at the respective exercise. (We might be a bit flexible here, in the sense that you can have your work assessed also in the other lab sessions if you have a compelling reason for doing so. But sticking to the schedule is the preferred options as that will make things run more smoothly). Also email one set of code per group to kolbjorn@chalmers.se with "SoCS HP1" and your group number in the subject.

Make sure you go through your demonstration by yourself before so that everything works. Everyone involved will appreciate the reduced queue times. Feel free to show just a subset of the exercises if you haven't done them all (whether you plan to do so later or not).

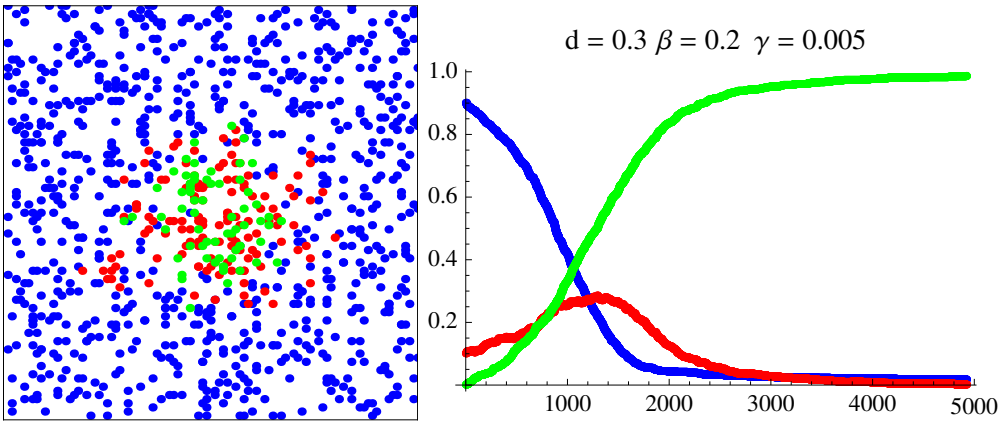

Figure 1: SIR-agents on a lattice and a plot of the proportions of individuals in each state over time.

Exercises:

1. Implement the basic model and visualize it. Start with just a single agent and make sure it performs the random walk correctly. Then test some small number of agents to check that the disease dynamics seems reasonable, and then scale it up to, e.g., a thousand agents on a 100×100–lattice. To demonstrate: real-time or recorded visualizations of each of these sanity checks. **(6p)**

2. Show that the model contains both regimes, i.e., that there are param-

eter values for which the disease spreads to a large proportion of the population and values for which it doesn't. To demonstrate: visualization of the agents and plots of the proportions of agents in each group over time at one set of parameter values for each regime. **(6p)**

3. Show that, in contrast with the ODE:s discussed in the lecture, the epidemic threshold R_c depends on not just the ratio β/γ but on the parameters themselves. Use a small initial number of infectives ($\approx 1\%$), fix a value for β, run the model for each of several values of γ and record the final proportion of recovered $R_\infty = R(\infty)$. Plot these values as a function of β/γ, similarly to Fig. 2. You should clearly see the two regimes. Note that it might be beneficial (and good practice) to do several runs at each setting and average the results to avoid noisy results. Then repeat the process for another value of β and compare the results. To demonstrate: a plot comparing the two data sets, showing clearly respective threshold. **(7p)**

4. Map out the whole critical line in the β-R_0-plane, i.e., repeat the above process for enough values of β that you can determine the important features of the phase diagram. To demonstrate: the resulting 3D phase diagram. **(6p)**

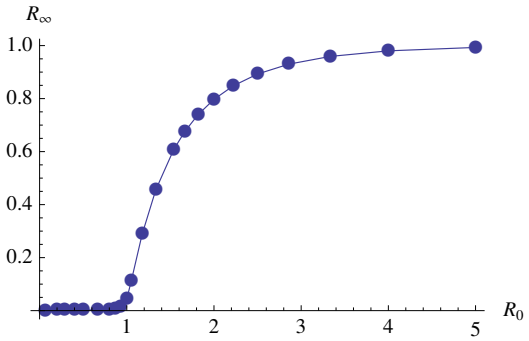

Figure 2: The epidemic threshold from an ODE version of the SIR model. The total proportion of the population afflicted by the disease, R_∞, is negligible under the critical value $R_c = 1$ of the basic reproduction number and quickly increases towards 1 above it.