

Epidemic Modelling Using Cellular Automata

Shih Ching Fu and George Milne

School of Computer Science & Software Engineering
Faculty of Engineering, Computing and Mathematics
The University of Western Australia
35 Stirling Highway, Crawley, Western Australia, 6009.
Phone: +61 8 9380 2281
Fax: +61 8 9380 1089
{scfu,george}@csse.uwa.edu.au

Abstract. This paper presents some of the results of our cellular automaton (CA) epidemic model. Of particular interest is how a CA can produce different patterns of spread according to geographical conditions. The model described here, although unvalidated, produces patterns of emergent behaviour which intuitively mimics what we observe in nature. This leads into the possibility of using CA-based simulations for ‘what if?’ game playing and hence improve our understanding, analysis, prediction, and possible containment of epidemic spread as influenced by geographical factors.

Keywords: epidemic, spatial models, cellular automata.

CR Classification: I.6.5, J.3.

1 Introduction

Public health issues are seeing greater visibility in the media; of particular concern is virus spread in populated areas. Decreased worker productivity as a result of illness costs industry millions of dollars every year [1]. The recent virus epidemics such as foot and mouth disease in the United Kingdom [2], Avian influenza in the Hong Kong and the Netherlands [3], and of course the Severe Acute Respiratory Syndrome outbreak in Asia [4], have meant that the monitoring of outbreaks is gaining importance for governments and public health officials. Hence it is desirable to predict patterns of viral infection given certain environmental conditions. It is hoped that modelling geographically dependent features of a phenomenon such as virus spread will help us better understand, predict, and ultimately control that phenomenon’s behaviour.

1.1 Non-homogeneous Populations and Disease Spread

The majority of existing epidemic models utilize differential equations [5, 6] and do not take into account spatial factors such as variable population density and

population dynamics [7]. These models, such as those using mean field type equations [8] and Markov chains, assume populations are closed and well mixed; that is, host numbers are constant and individuals are free to move wherever they wish. When trying to devise more realistic models we should attempt to incorporate spatial parameters to better reflect the heterogeneous environment found in nature. An alternative to using deterministic differential equations is to use a two-dimensional cellular automaton to model the location specific characteristics of a susceptible host population. Further to that, it is easy to include stochasticity in a CA model and capture the probabilistic nature of disease transmission.

1.2 Cellular Automata

Cellular automata (CA) are characterized by their discretization of space and time [9]. Typically a cellular automaton consists of a graph where each node is a finite state state automaton (FSA) or *cell*. This graph is usually in the form of a two-dimensional lattice whose cells evolve according to a global update function applied uniformly over all the cells. As arguments, this update function takes the cell's present state and the states of the cells in its *interaction neighbourhood* as shown in Fig. 1. The interaction neighbourhood is a collection of nearby cells which the update function interrogates for state information. The size and shape of neighbourhood varies from application to application.

As the CA evolves, the update function will determine how microscopic (or local) interactions influence the overall macroscopic (or global) behaviour of the complete system.




Fig. 1. State transitions depend on neighbourhood states.

1.3 What is a virus?

Viruses are primitive organic entities, so primitive that scientists are reluctant to class them as living organisms. On their own, viruses are lifeless and unable to reproduce [10]. Viruses themselves comprise two main components: genetic code and protective protein coating. The primary goal of a virus is to replicate by penetrating living cells, rewriting their DNA, and programming them to create hundreds of copies of the virus. Consequently, they can replicate rapidly.

Whilst a single infected host might not be significant, a virus that has spread through a large population of hosts presents serious health and economic threats. The study of the health of entire *populations*, particularly those exposed to an infectious disease, is known as epidemiology; this paper reports the outcome of a study of cellular automata as the basis for a new class of models applicable to infectious disease epidemiology.

1.4 What is an epidemic?

The term ‘epidemic’ is often misused; most people actually mean ‘outbreak’. In either case, it is important to remember that ‘epidemic’ and ‘outbreak’ are relative terms whose meanings vary between diseases and environments. For instance, during any one time there are thousands of cases of influenza across Australia: influenza is *endemic*. An *outbreak* of a hundred cases would not be cause for alarm. However, if there was just a single case of smallpox, its high mortality rate and the lack of immunity in our population means health authorities would be placed on high alert and severe containment measures put in place. When an outbreak reaches many times the endemic level, it is classed as an epidemic. Large scale epidemics such as the influenza outbreak in 1918 are deemed *pandemics* [11].

Epidemics are an integral part of nature and the possibility of major health, economic, and military impacts of epidemics makes them the subject of extensive modelling research.

2 Epidemic Models

Models of epidemic spread all share one property: the virtual world in which they run is an idealized and approximate one. This arises from the difficulty of incorporating all the variables we see in nature into a simulation that is both accurate and tractable. When modelling a complex system there is a trade-off between a model’s degree of abstraction and its usefulness; that is, which details can be excluded without devaluing the results?

In the SIR epidemic representation framework [7], a host can be Susceptible, Infective, or Recovered. Susceptible individuals are those who are healthy and do not carry the virus but can contract it from infective hosts. Infective hosts carry the contagion and are able to pass it on to another host. Recovered hosts are those who are no longer infective and have acquired immunity from future infection. Note that this immunity is not necessarily everlasting and that the *transition* between the S, I, and R *states* is probabilistic, with probabilities being determined by the observed characteristics of specific diseases.

3 CA Model for Disease Dynamics

The cell evolution in a cellular automaton follows an update function that takes the state of a particular cell and its neighbours and determines the next state.




Fig. 2. The progression of symptomatic phases (below the line) and their relationship with infection life cycle (above the line).

Some previous models have used deterministic update rules [12], but probabilistic rules appear to reflect nature more accurately [13–15]. Here we describe the cell and update rule definitions adopted in our probabilistic epidemic model.

A CA cell represents an equal sized area of landscape containing a specific population. Different cells will have different populations, that is, differing densities and possibly different ‘across cell’ traversal or mobility properties.

3.1 Cell Definition

The basic behavioural unit in our model is a cell. Here, ‘cells’ are *automata* cells and not *biological* cells. The user is able to define the physical area of landscape which is modelled by a single cell by adjusting various epidemic spread parameters. For example, if for a particular execution of the simulation the host mobility parameter is set very small, the cells can be interpreted as being physically large compared with a later simulation where the mobility is set high. This is because in the first case, the low mobility (or high traversal difficulty) can be equated to the large distances hosts need to cover to reach adjacent cells and so pass the disease on to adjacent cells.

For our model, care must be taken to differentiate *cells* from the *hosts* that reside inside the cells. Each cell has the following attributes:

- carrying capacity,
- total population,
- susceptible subpopulation,
- infective subpopulation, and
- recovered subpopulation.

These attributes are used to encode information about population densities and host mobility,. Typically, these attributes will be determined by by geographical (spatial) features such as urban versus rural and free-movement versus limited village-to-village movement.

The carrying capacity of a cell is used as a mechanism to limit the movement of hosts between cells. It prevents crowding within a particular cell, likened

to an abstract measure of surface area. The number of newborns per epoch is also dependent on whether a cell has reached its carrying capacity. Although the effect of the land's carrying capacity is not directly enforced in nature, for the purposes of simulation, carrying capacities are a straightforward way to encourage or attenuate the motion of individuals between cells.

Traditionally, as in the CA of von Neumann [16], a single 'object' occupied each of the cells that constituted the larger cellular automaton. Rather than stipulate this, our model allows multiple individuals to dwell in one cell up to the above-mentioned carrying capacity. Allowing variable cell population provides two main advantages: first it reduces the total number of cells and hence reduces computation time; secondly it provides generality. For instance, if we set the carrying capacity to 1 we revert back to a 'classic' cellular automaton. Note that each cell is considered *well-mixed* in the sense that during each time step, all the individuals in a particular cell will have come in contact with one another.

3.2 World Definition

A two-dimensional array of cells and the epidemic spread parameters that govern their evolution constitute the world that the hosts 'live' in. The cells are arranged in a rectangular grid comprising square cells with external dimensions that may or may not be square. The world boundaries serve as impenetrable barriers to host movement, which conceptually could be oceans or political boundaries allowing immigration into and out of this world into adjacent worlds. The adjustable epidemic parameters that control cell evolution and hence the emergent and often unpredictable behaviour of the overall system are described in the next section.


3.3 Adjustable Simulation Parameters

The following is a composite list of epidemic spread parameters as used by existing epidemic models [5, 17, 18], particularly those models examining virus pathogens that can survive outside the bodies of hosts. This is not an exhaustive list, but it contains what are believed to be the most significant factors that account for the behaviour of an epidemic and consequently implemented in our model. Apart from the interaction radii, all the following parameters are modelled using probabilities that directly impact the update rules applied over the CA lattice. By changing the probabilities, the specific characteristics of distinct diseases may be modelled. This paper only discusses those parameters which are a function of population density and dynamics, that is, neighbourhood size, motion, vectored infection, and contact infection. The other parameters are mentioned for completeness.

- neighbourhood radius
- motion probability
- immigration probability
- birth rate

- death rate
- virus morbidity
- vectored infection probability
- contact infection probability
- spontaneous infection probability
- recovery probability
- re-susceptible probability

Neighbourhood radius This parameter determines the size of the interaction neighbourhood that a cell interrogates for state information. Our model uses the square Moore neighbourhood whose area, n , is determined by an interaction radius, r , as shown in Equation 1.



$$n = (2r + 1)^2 \tag{1}$$

There are two distinct interaction radii: motion and infection. The motion radius defines the greatest distance, measured in cells, a host can move during a time step. The infection radius is slightly different to the motion radius in that it does not relate to hosts but to the virus pathogen. The infection radius defines the greatest distance the virus pathogen can travel outside the body of a host on its own. This quantity is used to model the spread of a virus via natural vectors such as airborne droplets in influenza or vermin as in bubonic plague.

Motion probability The individuals in the world are permitted to move between cells. The motion probability determines the frequency of this motion. Our model assumes homogeneous mixing *within* each cell, but the motion of hosts *between* cells is governed by the motion probability parameter. For example, if the motion probability, p_{move} , is set to 0.4, you would expect roughly two in five hosts to shift from the cell they currently reside in to another cell within its motion neighbourhood. The success of a host's inter-cell movement is dependent on whether the destination cell has reached its carrying capacity. In our model, the destination cell is selected randomly from the local interaction neighbourhood.

Vectored infection rate Apart from entering an otherwise ‘clean’ cell inside a mobile infective host, the virus contagion can spread across cells using its natural spreading mechanisms. Such spread is known as *vectored* infection. For our model, the velocity of inter-cell infection is controlled through a p_{input} parameter. The actual probability of spread, $p_{vectored}$, is a function of the vectored infection parameter provided by the user, p_{input} , and the density of susceptible hosts in the local interaction neighbourhood. This is illustrated in Equation 2. A high value for p_{input} would represent a highly transmissible virus such as one that was airborne or carried in bird droppings.

$$p_{vectored} = \frac{\text{susceptible population}}{\text{neighbourhood capacity}} \times p_{input} \quad (2)$$

The other seven parameters, although affecting epidemic spread, are not strictly spatial and their descriptions have been omitted in this discussion.

3.4 Cell update algorithm

A CA cell update function defines each cell’s evolution to its new state. Our cell update function takes as arguments all the parameters outlined in the previous section along with the state information from the cells in the interaction neighbourhood of the cell in question.

The lattice update in our model is performed in two phases: the movement phase and the infection phase. It is important to note that the effects of the movement phase are instantaneous, that is, cell populations are immediately updated to reflect host traversals. During the infection phase, although implemented sequentially over the lattice, each cell evolves synchronously with all the others.

Movement phase algorithm

1. Select a random cell from the world.
2. For each of the individuals in the cell, randomly select a neighbouring cell and move the individual into it. This movement is dependent on the motion probability parameter described earlier and the destination cell not being full.
3. Repeat from step one until all the cells in the world have been accounted for.

Infection and recovery phase algorithm The ordering of events during the infection and recovery phase are such that similar operations are performed consecutively, that is, all operations affecting host numbers are dealt with before infections are calculated. Other inter-leavings should be tried in order to examine possible biases.

1. Select the first (left upper-most) cell.

2. Deduct from the cell population deaths by ‘natural causes’.
3. Deduct deaths resultant from infection.
4. Add to the population any newborns.
5. Add to the population any immigrants.
6. Compute inter-cell (vectored) infections.
7. Compute intra-cell (contact) infections.
8. Compute spontaneous (randomly dispersed) infections.
9. Compute host recoveries.
10. Compute new susceptibles hosts (due to loss of immunity).
11. Repeat from step two for the next cell until all cells have been accounted for.

This model does not take into account other parameters such as the latency or incubation times illustrated in Figure 2. Latency is the lag between being infected and becoming infective and incubation is the delay between being infected and becoming symptomatic. These times could be implemented by introducing another state, for example ‘E’ for exposed hosts, who are infected but not infective [17].

4 Simulation Results

Here we present the results of two experimental scenarios using the model described above. Emphasis is placed on the emergent patterns that are produced by the model, particularly in response to the underlying spatial heterogeneity. That is, we have tried to capture geographical differences over the landscape and observe the effects they have on the macroscopic system behaviour. These simulations provide some evidence of CA’s usefulness for epidemic spread visualization.

4.1 Corridors of Spread

This experiment tries to emulate a real life landscape with imaginary town centres and transport links. Towns and roads are human constructed features that attract high population densities. Conversely, this scenario could be depicting non-cultural features such as rivers that promote development along their shores or mountain ranges which might limit population settlement and movement. Most important is the directed and linear shape of these settlements rather than wide uniform population densities. Figure 3 shows the virtual landscape used in this experiment.

Each cell has a carrying capacity of 1000, with three ‘towns’ set to this maximum. Two of these towns, the north-west one and the south-east one, start with a 1:9 infective to susceptible ratio. The ‘transport links’ comprise a three cell wide bar of susceptibles. Cells on the central axis of this bar have an initial susceptible population of 100, whilst the cells on either side have an initial susceptible population of 75. The rest of the landscape comprises cells with 10 susceptibles in them. The town in the south-west corner contains 1000 susceptibles and no infectives.




Fig. 3. The synthetic landscape in this figure captures the tendency of settlement to form around prominent cultural and geographical features. Here there are three points of high population density – two of which are connected by a transport link. The link itself has settlement developed along it. The north-west and south-east ‘towns’ both contain 100 infectives and 900 susceptibles.

Parameter settings Table 1 lists the parameter values used in this scenario. Essentially all of the parameters are zero except for the infection radius, contact infection probability, and vectored infection probability which are all equal to one. These values capture a simplified environment where the population is completely static geographically and demographically. The purpose of using high probabilities is to accelerate infection spread and isolate the effects of population density on epidemic dynamics.

Results The results of executing our model with the above disease-specific and landscape-specific features are presented in the time lag map shown in Figure 4. The lag map is basically a series of snapshots taken at $t = 0, 20, 40, 60, 80, 100, 200, 300$. Each cell is represented by a square: black squares contain at least one infective host and white squares contain only susceptible and recovered hosts. Figure 4 shows the tendency of the epidemic to follow the lines of population density to produce the ‘fuzzy cross’ pattern.

After 300 epochs, the top left outbreak has reached all four edges of the map but the bottom right outbreak is yet to reach any. Notice that the epidemic spreads outward along the arms or ‘roads’ before filling up the space between the roads. This illustrates how a disease spreads over an area of higher population density faster than over under or unpopulated areas, as we would expect.




Fig. 4. A lag map showing the state of the epidemic at $t = 0, 20, 40, 60, 80, 100, 200, 300$. Notice that the outbreak to the north-west is able to cover a greater distance than the outbreak in the south-east because it has access to the road link and the population associated with that link. Notice that the spread from $t = 20$ in the top left of the map appears asymmetric. This is probably an artifact of the stochastic nature of this model.

Parameter	Value
Infection radius	1
Movement radius	0
Immigration rate	0.0
Birth rate	0.0
Natural death rate	0.0
Virus morbidity	0.0
Spontaneous infection rate	0.0
Vectored infection rate	1.0
Contact infection rate	1.0
Recovery rate	0.0
Resusceptible rate	0.0
Movement probability	0.0

Table 1. The parameter values for the ‘Corridors of spread’ scenario. All of the parameters are set to zero except for the ones that relate directly to population density. That is, all are zero except for the infection radius, vectored infection rate, and contact infection rate which are set to unity.

4.2 Barriers to spread

This experimental scenario depicts how a CA model can simulate the effects of erecting barriers to slow or stop virus spread. As seen in the foot and mouth disease epidemic in Great Britain during 2001, a key to slowing down disease spread is restricting movement [19]. Other measures included the culling of livestock or the inoculation of livestock, hence creating barrier areas through which the virus cannot spread. These measures are simulated by incorporating ‘no spread zones’ in the initial state of the CA lattice.

The starting distribution contains two ‘hot spots’ which have been segregated from the rest of the landscape. One hot spot has a four square wide barrier surrounding it, whilst the other has a one square wide barrier surrounding it. Barriers are implemented as cells with zero carrying capacity. In Figure 5, barriers are represented by black squares and all other (blank) squares contain an equal number of hosts. The grey squares depict the two sources of infectives, both of which confined by ‘buffers’. The barriers restrict host movement and provide no hosts for any viruses to infect and escape.

Parameter settings Whilst it is quite simple to define a typical host population, it is much more difficult to define a typical virus strain and keep generality. Consequently, although the values in Table 2 do not need to quantitatively represent a particular virus, the ratios between the parameters need to be reasonable. Essentially these numbers are educated guesses at the values for a ‘typical’ virus. An interaction radius of two means that the the enclosed infectives in the bottom right corner of Figure 5 should still be able to disperse into the surrounding environment.

Parameter	Value
Infection radius	2
Movement radius	1
Immigration rate	0.01
Birth rate	0.02
Natural death rate	0.01
Virus morbidity	0.05
Spontaneous infection rate	0.0001
Vectored infection rate	0.2
Contact infection rate	0.4
Recovery rate	0.1
Resusceptible rate	0.001
Movement probability	0.001

Table 2. Epidemic spread parameter values for the containment scenario.

Results The resultant epidemic spread is shown in the lag map of Figure 5. Notice that the virus is able to elude containment from the bottom right enclosure. Despite the enclosure itself becoming free of infectives, the neighbouring country side has become infective. The upper left enclosure is no longer contaminated and neither is its surrounding hinterland.

5 Discussion

These scenarios provide evidence of the power of CA models as an underlying theory for novel epidemic simulation and visualization; differential equation models provide no such capability. In Figure 4 we can clearly see where the infective towns are, where the densest populations are situated, as well as identify where the infection velocity is greatest. Features such as *leap frogging* or *infilling* [20] where the epidemic appears to skip regions of land before coming back to fill in the space can be seen in the lag map.

This behaviour was documented in the foot and mouth disease epidemic in Great Britain in 2002 [21] and is due to the stochastic nature of epidemic spread.

Prediction techniques such as these may be used in future to help public health officials efficiently direct containment measures and medical services to deal with the dynamics of infection spread.

6 Conclusion

These epidemic scenarios presented above provide an opportunity to demonstrate the visualization capabilities of a graphical CA model. There are no new epidemiological conclusions to be drawn from either of these two experiments – they show exactly what our intuition would suggest. What is important is that other statistics-based models, such as Markov chain models, do not illustrate




Fig. 5. This lag map shows that buffer zones that are too narrow provide no resistance to the spread of a virus.

such behaviours well. As far as sample data is concerned, very little is available about the spatial behaviour of epidemics; much of infectious disease epidemiology is focussed on statistical data and deriving models that match that data [7]. This is particularly evident with lag maps whose data is gathered during and after the epidemic, rather than being generated preemptively.

This research has shown the applicability of using CA models as a predictive tool for epidemiologists. The next phase of this research is to utilize comprehensive historical data to calibrate and validate models for specific epidemics.

7 Acknowledgements

Thanks are due to Professor Roy Anderson, Department of Infectious Disease Epidemiology, Imperial College, for his help and guidance during the course of this research.

References

1. Australian Associated Press. Flu-hit workers urged: Stay home. *The West Australian Saturday April 20*, page 5, 2002.
2. Department for Environment Food and Rural Affairs. Foot and mouth disease information page. Web page. Available: <http://footandmouth.csl.gov.uk>.
3. James Randerson and Debora MacKenzie. Netherlands' chicken cull tries to stop bird flu virus in its tracks. *New Scientist*, 2391, 19 April 2003.
4. World Health Organisation. Press release: Cases of severe respiratory illness may spread to hospital staff. Web page, 12 March 2003. Available: http://www.who.int/csr/sars/archive/2003_03_12/en/.
5. Denis Mollison, editor. *Epidemic Models: Their Structure and Relation to Data*. Cambridge University Press, 1995.
6. R. Bagni, R. Berchi, and P. Cariello. A comparison of simulation models applied to epidemics. *Journal of Artificial Societies and Social Simulation*, 5(3), June 2002. Available <http://jasss.soc.surrey.ac.uk/5/3/5.html>.
7. Roy M. Anderson and Robert M. May. *Infectious Diseases of Humans: Dynamics and Control*. Oxford University Press, 1991.
8. Adam Kleczkowski and Bryan T. Grenfell. Mean-field-type equations for spread of epidemics: The 'small world' model. *Physica A*, 274(1-2):355-360, 1999.
9. E. F. Codd. *Cellular Automata*. Academic Press, 1967.
10. C. H. Andrewes. *The Natural History of Viruses*. W. W. Norton & Company, 1967.
11. Frank Ryan. *Virus X*. Harper Collins, 1996.
12. G. Rousseau, B. Giorgini, R. Livi, and H. Chate. Dynamical phases in a cellular automaton model for epidemic propagation. *Physica D*, 103:554-563, 1997.
13. Nino Boccara and Kyeong Cheong. Critical behaviour of a probabilistic automata network SIS model for the spread of an infectious disease in a population of moving individuals. *Journal of Physics A: Mathematical and General*, 26(5):3707-3717, 1993.
14. Nino Boccara, Kyeong Cheong, and Mark Oram. A probabilistic automata network epidemic model with births and deaths exhibiting cyclic behaviour. *Journal of Physics A: Mathematical and General*, 27:1585-1597, 1994.

15. Maria Duryea, Thomas Caraco, Geoffrey Gardner, William Maniatty, and Boleslaw K. Szymanski. Population dispersion and equilibrium infection frequency in a spatial epidemic. *Physica D*, 132(4):511–519, 1999.
16. A. W. Burks, editor. *The Theory of Self-Reproducing Automata*. University of Illinois Press, 1966. As cited in [9].
17. E. Ahmed and H. N. Agiza. On modeling epidemics. Including latency, incubation and variable susceptibility. *Physica A*, 253:347–352, 1998.
18. E. Ahmed and A. S. Elgazzar. On some applications of cellular automata. *Physica A*, 296:529–538, 2002.
19. Neil M. Ferguson, Christl A. Donnelly, and Roy M. Anderson. Transmission intensity and impact of control policies on the foot and mouth epidemic in Great Britain. *Nature*, 413:542–548, October 2001.
20. Susan Scott and Christopher J. Duncan. *Biology of Plagues: Evidence from Historical Populations*. Cambridge University Press, 2001.
21. Neil M. Ferguson, Christl A. Donnelly, and Roy M. Anderson. The foot-and-mouth epidemic in great britain: Pattern of spread and impact of interventions. *Science*, 292:1155–1160, May 2001.