CELLULAR AUTOMATA MODELS OF HUMAN TRAFFIC

Richard Holden and Angelo Cangelosi
Centre for Neural and Adaptive Systems
University of Plymouth
rholden@plymouth.ac.uk

Submitted to: First conference of the European Social Simulation Association, September 18-21, Groningen, The Netherlands

Overview

In the area of artificial life and multi-agent systems traffic flow has been modelled with cellular automata, i.e. using simple locally defined rules of communication in the medium of cellular space (Nagel and Rasmussen 1994, Blue and Adler 2000a, 2000b). This use of cellular automata for traffic modelling demonstrates how populations of simple units can interact through time and produce different levels of complexity and unpredictability depending on the level of interaction within the system.

We replicate a car traffic model (Nagel and Rasmussen 1994) in order to show how ‘phase transition’ behaviour depends on specific pre-defined parameters. We also replicate and extend some more complex two-dimensional pedestrian movement models (Blue and Adler 2000a, 2000b). These two separate, von Neumann space models have different application areas, but share the two explicit aims; 1) optimisation of forward movement and 2) resolution of conflict between entities. We extend the pedestrian flow model to Moore neighbourhoods, which allows a comparison to be made between Blue and Adler’s rule sets and new rule sets specifying 8–directional movement. In a final section we redefine ‘density’ in Blue and Adler’s model so that each automata has a definition of density and can act locally according to this, rather than by reference to a globally defined density parameter. The definition of density in these models greatly affects their behaviour. Also, by relying on global definitions of density we argue that certain types of relative movement show better performance, but that this is due to some of the default, stochastic behaviour rather than any rule based conflict resolution.

‘Car traffic’ Nagel and Rasmussen model

The car traffic flow model can be considered as a precursor to the pedestrian flow models as it is very similar to the more simplistic, uni-directional types of pedestrian flow models in terms of the update rule sets. It makes sense therefore to show some of the behaviours of these more simple systems before going on to consider more complex scenarios of human traffic.

The first model is a replication of Nagel and Rasmussen (1994). It consists of a 1-dimensional wrap-around cellular ‘road’ and a number of ‘cars’, which occupy cells along the road. Each car has 3 fundamental rules corresponding to ‘acceleration’, ‘breaking’ and the ‘stochastic effects of a car driver’. Each car moves forward along the road based on a parallel update of each rule. The qualitative behaviour of the model and some of the statistical properties associated with the density parameter are particularly interesting. For example this can be shown with space-time plots. These are a typical technique used to visualise 1D cellular automata development over time. In the Nagel and Rassmusen model we can observe how the traffic 'jams' travel backwards over time. The technique provides a qualitative snapshot of a given section of time and here we see the effect of increasing density on the amount of jams in the model.

Quantitative results show that the time taken for any given car to reach a predefined destination increases exponentially with increasing density. This makes sense; it would take an infinite amount of time to reach a destination where the density in the system is equal to 1, representing a situation where there is no space for any forward movement along the road for any car.

Another feature of the model’s behaviour is that the variations in the time taken to reach a destination by each car sharply increases and peaks at low densities. Nagel and Rasmussen showed how these peaks are found in a regime where the system flow is optimal. So, paradoxically we have a situation where two desirable features of a transport system, local predictability and global performance, are inversely related. The predictability of interactions between cars is at its minimum when the system performs it a maximum. Nagel and Rasmussen see this specific behaviour as part of a more general conflict between the individual traveller plans (Nash Equilibrium) and the travel plans that give overall maximal throughput (System Optimum).

The rules in the Nagel and Rasmussen model are very simple, but we get complex behaviour out of a population of these rules. This complexity is defined by methods in statistical physics. Due to its simplicity, the model has well defined transitions and critical points. Such models lead us away from the view of multi-agent traffic models as fundamentally linear where units are treated in isolation. Statistical measures like mean speed do not give us a full picture of the relationship between local global behaviour. We need to look at other types of statistics like travel time variations. These statistics give us more information regarding interactions inside the system and how these interactions affect global dynamics. (Sole and Goodwin, 2000).

Our results do not contradict these arguments. However, we run the model with varying parameters. For example, we use varying values for the stochastic effects of car driver. Results indicate that the sharpness of the critical behaviour is very sensitive to these values. This serves to illustrate how we can control the behaviour of a system in a fairly predictable manner by changing simple parameters. This illustration is a precursor to a more in depth analysis of the Blue and Alder model where we also change predefined parameters in the model to show how certain types of behaviour cannot really be considered to ‘emerge’ from rule sets, but is rather something, which depends to a large amount on these built in parameters.

Pedestrian Traffic (Blue and Adler model)

The one dimensional car traffic models provide a basis for developing more complex models of movement. The area of pedestrian movement has been used as possible application field for the use of cellular automata (Blue and Adler 1992). In the following sections we will first analyse some of the results of these models by replicating the rule sets of Blue and Adler for von Neumann spaces. It will be shown that the models do indeed optimise forward movement to some extent, but the movement is probably sub optimal at a population level and the relationship to the density parameter is complex across different model comparisons. We will then extend the model to a Moore space.

The models to be outlined in this section contain cellular entities that have a forward direction of movement and the idea is to optimise the speed of the agent in a given direction, under a maximum walk speed constraint. Each agent will account for the position of other agents and their direction of forward movement. In the simplest case we could have an environment where each agent is moving in the same direction as every other agent; 'uni-directional' flow like in the car traffic model above. The next increase in complexity involves 'bi-directional' flow where two types of agents in the population move in opposing directions. Quadri-directional flow further increases complexity to cover all possible local moves in a Von Neumann neighbourhood and defines flow for four directions.

Blue and Adler (2000) identify behaviour which they call 'Dynamic Multiple Lane' formation (DML) as opposed to 'Interspersed Flow' (ISP). The former 'emerges' due to some rules which essentially model the behaviour of pedestrians in situations where, given the same distance between two other agents, if one is moving in the same direction as the other, then the first agent will move to form a lane and avoid any oncoming pedestrians. This is therefore presented as an important part of the algorithm for the optimisation of forward movement.

Another important feature of the crowd behaviour is Mode Locking. This is related specifically to the resolution of lateral conflict in the model. It is argued by Blue and Adler that such behaviour is 'emergent'. A more important aspect of the work of Blue and Adler is the claim that the behaviour is self-organising. This is not important because the claim is necessarily a bad one based on the results Blue and Adler present, but because if we analyse the model in further detail we come to understand how sometimes what comes to be understood as self organising reflects the detail in which we observe a system.

Results indicate that in the Blue and Adler model the so called self-organisation is in fact a reflection of more or less arbitrary factors concerning definitions of different walk speeds for the population of agents. It is not intended that this observation contribute in any way as to what self-organisation is or is not, but does elucidate how decisions made from the top down can influence behaviour which might seem to be self-organising if the fixed parameters are not analysed. Our results do analyse these parameters simply by varying the distributions used for the rule sets.

The variations in the time taken to reach the destination show how variations increase rapidly at different densities. This behaviour is expected and consistent with the earlier car model. However, if we choose different distributions of walk speed for a population we see how different distributions behave in different ways for low density, but come to behave similarly at certain points along the density axis. We argue that this behaviour reflects the choices made regarding walk speed distribution rather than any unforeseeable interaction within the rule sets.

Extension of Blue and Adler model

We carried out some experimental work on an extension of Blue and Adler's model, which involved increasing the complexity of the number of directions of movement to octo-directional. Rule sets are also adapted to account for this neighbourhood. We record the population speed over a range of density (0 - 1). Uni-directional pedestrians are the most optimal population. Bi-directional populations are more optimal than quadri- and octo-directional populations. However, this second point is only true below certain densities.

Simulation snapshots of the population performance at different densities are presented graphically. At a density of 0.3 (where density = number of automata / number of available cells in space), clear differences between the populations have emerged in terms of conflicting clusters. The bi-directional populations are the most conflict ridden population, followed by the quadri- and then octo-directional populations. Considering the emphasis placed on the 'forward movement' rules by Blue and Adler (2000a), we might conclude before seeing the simulation that populations containing automata with less numbers of directions have a better chance of forward movement and this is true. However, only below a density of 0.3 we see 'forward movement' rules work. The snapshots show how the more simplistic bi-directional population accumulates increasing amounts of clusters (blockages) over time and at 0.3 these blockages appear to be worse than those more complex populations (quadri and octo). It could be said therefore that at a density of 0.3, 'forward movement' becomes 'unforward movement'. This sounds silly, but there is a serious point to this. We argue that this unintuitive result is probably a reflection of some of the defaults to stochastic processes, which appear more readily in models with more complex movement. Also, this randomness can aid conflict resolution.

Global or local density

Blue and Adler define a global density parameter for the model. This determines how many agents will occupy the space. So, if we have a density of 0.5 and the space is 50 * 50 there will be 1250 pedestrians in the model. So, Blue and Adler define a global parameter for the density of the model. We introduce a local definition of density, which can be used to control the 'exchange' parameter; the probability that two pedestrians will exchange places in the face of conflict. We define a simple function to allow a decreasing probability of exchange with increasing local density. This new definition of exchange removes a large part of the randomness in the Blue and Alder model where the probability of exchange was fixed globally at 0.5 for all pedestrians throughout the simulation.

We present results that show clear transitions in the variations of the travel times indicating that at these points it is highly unpredictable when an agent will arrive at it's destination. These results show similar patterns to the earlier and much more simple models of Nagel and Rasmussen.

By using this local definition of exchange probability we can see how the Blue and Adler rules for optimisation of forward movement seem to allow the population to perform better in densities up to around 0.4 where all complex types of flow perform poorly. It seems that the rule sets are significantly effecting how populations organise themselves in order to increase forward flow where this is possible. By removing some of the randomness from the rule sets we have given the actual rules for forward movement more chance to show themselves.

References

Nagel, K., and Rasmussen, S. (1994). Traffic at the edge of chaos. In : Artificial Life IV, R.A. Brooks and P.Maes (eds.), 222-230. MIT Press (Cambridge, MA, 1994).

Sole, R. and Goodwin, B (2000) Sign of Life: how complexity pervades biology. Basic Books. New York.

Blue, V. J. and Adler, J.L, (2000a). Cellular Automata Model Of Emergent Collective Bi-Directional Pedestrian Dynamics. In Artificial Life VII Bedau, M.A, McCaskill, J.S, . Packard, N.H and Rasmussen, S (MIT Press, 2000). Available at http://www.ulster.net/~vjblue/research.htm

Blue, V.J. and Adler, J.L. (2000b) Modeling Four-Directional Pedestrian Movements. The Transportation Research Record, Journal of the Transportation Research Board. Available at http://www.ulster.net/~vjblue/research.htm