Environment and Planning B: Planning and Design 1997, volume 24, pages 159–164

Editorial

Urban systems as cellular automata

Cellular automata (CA) date back to the very beginnings of digital computation. Alan Turing and John von Neumann who pioneered the notion that digital computation provided the basis for the universal machine both argued, albeit in somewhat different ways, that digital computation held out a promise for a theory of machines that would be self-reproducible (Macrae, 1992), that computers through their software could embody rules that would enable them to reproduce their structure, thus laying open the possibility that digital computation might form the basis of life itself. This was a bold and daring speculation but it followed quite naturally from the philosophy of computation established in the 1920s and 1930s. Von Neumann perhaps did most to establish the field in that, up until his death in 1956, he was working on the notion that a set of rules or instructions could be found which would provide the software for reproducibility. The idea of automata flowed quite easily from this conception and the notion that the elements of reproducibility might be in terms of 'cells' was an appeal more to the 'possibility' of using computers as analogs to create life than any actual embodiment of such functions through computer hardware.

Von Neumann worked on many projects, CA being only one. His work was published posthumously by his student and colleague Arthur Burks who carried on this work at the University of Michigan in the 1960s and 1970s (Burks, 1970) and this led directly to ideas and methods in complexity based on genetic algorithms, artificial life, and neural nets. But another line was emerging in this development based on a simpler, more visual, approach to automata. Von Neumann drew some of his inspiration from Stanislav Ulam, the mathematician who worked with him on the Manhattan project, who had suggested to him as early as 1950 that simple CA could be found in sets of local rules that generated mathematical patterns in two-dimensional and three-dimensional space where global order could be produced from local action (Ulam, 1976). It was this line of thinking that was drawn out, as much because in 1970 John Conway, a mathematician in Cambridge, England, suggested a parlor game called 'Life' which combined all the notions of CA in a model which simulated the key elements of reproduction in the simplest possible way. Life is described by some of the authors writing in this special issue. It has become the exemplar par excellence of CA but its popularity rests on the fact that a generation of hackers took up Conway's idea and explored in countless ways the kinds of complexity which emerge from such simplicity.

It is probably worth stating the elements of Life for it embodies the key elements of CA. In essence, Life can be played out on any set of cells which exist in any space but usually it is convenient to think of this space as being a regular tessellation of the two-dimensional plane such as a grid where each square represents a cell. Any cell can be alive—'on'—or dead—'off'—and there are two rules for cells becoming alive or giving birth, or dying or not surviving. The rules are simplicity itself. A cell which is not alive becomes alive if there are exactly three live cells immediately adjacent to it. A cell remains alive if there are two or three live cells adjacent to it, otherwise it dies. Fewer than two adjacent cells implies the cell dies from isolation, more than three from overcrowding. The event that set the field humming in 1970 was Conway's challenge reported by Martin Gardner (1970) in his recreational mathematics column in *Scientific American.* Conway would give a prize of 50 dollars to the first person who could unequivocally demonstrate that certain configurations of Life could be self-perpetuating. The challenge was won by Bill Gosper and his group at MIT within the year who showed that a particular configuration of cells and their dynamics called a 'glider gun' would, under these rules, spawn live cells indefinitely (Poundstone, 1985).

Although Life is the best known cellular automaton, it is perhaps the least applicable to real configurations. However, it illustrates all the main principles which strict CA embody. Formally, CA are composed of four elements. First, there are *cells* which are objects in any dimensional space but which must manifest some adjacency or proximity to one another if they are to relate in the local manner prescribed by such a model. Second, each cell can take on only one *state* at any one time from a set of states which define the attributes of the system. In Life, for example, there are two states: dead or alive. Third, the state of any cell depends on the states and configurations of other cells in the *neighborhood* of that cell, the neighborhood being the immediately adjacent set of cells which are 'next' to the cell in question in some precise manner. Finally, there are transition *rules* which drive changes of state in each cell as some function of what exists or is happening in the neighborhood of the cell. The key elements which define a strict cellular automaton are that the rules must be *uniform*, they must apply to every cell, state, and neighborhood; and that every change in state must be *local*, which in turn implies that there is no action-at-a-distance.

It is this property of locality which has become so important to the contemporary applicability of CA. The fact that there is no action-at-a-distance in strict CA means that there is no prescribed or predetermined global order to the patterns and forms that CA might create. But life as we know it is ordered. From the simplest CA, it is easy to show that complex global patterns can emerge directly from the application of such local rules and it is this property of emergence which makes CA so intriguing. For example, recursive local actions which produce well-defined geometrical cell structures in two-dimensional space often generate similar geometries at higher scales as the structure grows and changes. These geometries are fractal in that their basic motif or pattern repeats itself across the scales, demonstrating fractal properties of self-similarity or self-affinity. Most fractals can thus be generated by CA for the recursive logic of form which underpins the basic principles, provides a fundamental generating mechanism for such geometries. In a world where global interactions fuse in subtle and diverse ways with local action, CA looks like a paradigm for the 21st century, resonating with everything from the postmodern mathematics of fractals and chaos to the cry of development theorists "Think globally, Act locally!". Indeed, one of the major developments in complex systems—artificial life—has come from CA—in fact from Arthur Burks' Michigan School which kept CA alight until the glory years began after Conway's popularization (Holland, 1975; Langton, 1989).

The great attraction of CA for systems theorists is not simply this ability to provide a framework for emergence but the fact that many classes of dynamics can be simulated through CA. The simple logic in CA of reducing everything to discrete states, discrete spaces, and discrete time, all of which must be uniform might seem at first to impose excessive restrictions on representation of system properties. But anything that can be translated from differential to difference equations can in principle be represented by CA. Moreover, many physical processes such as diffusion are local in their spatial-temporal impact in that longer range correlations are not central to their formulation. As Toffoli and Margolus (1987) show in their magnificent book on the subject, countless spatial-physical processes can be modeled by using CA, from spin glasses to diffusion-limited aggregation, from predator–prey relations to spatial politics. In fact, the really great attraction of CA is that it gives equal weight to the importance of space, time, and system attributes, thus imposing a frame which forces researchers to think very hard about representing any system where the importance of one of these elements becomes emphasized relative to the others.

The four CA principles involving cells, states, neighborhoods, and transition rules impose a strictness on the representation of any problem which often requires some relaxation so that CA might be suitably adapted to the problem at hand. For example, CA usually involve deterministic processes as in Life although, by relaxation of the rules governing change of state and by choosing such changes according to a uniform set of probabilities, structures which have a degree of randomness built into their regular order can be easily generated. For example, consider a cellular automaton in which cells are filled if, say, just one neighboring cell is already filled. Starting from a single seed, it is easy to show that a fractal pattern reminiscent of a highly structured dendrite emerges where the branches are composed of 'H' patterns which just touch one another. This kind of structure becomes less ordered but has the same generic form if the rule is made that, if 80% of the time a cell has just one active neighboring cell, that cell is made active. As this probability is reduced, the form becomes more random. In problems where a degree of ignorance concerning a change in state is characteristic, such adaptations are essential.

More problematic is the definition of neighborhood. Usually in two-dimensional space, the von Neumann neighborhood—cells which are N, S, E, and W of the central cell in question—or the Moore neighborhood—these cells plus those which are NW, NE, SE, and SW—are used. Neighbors which are more distant from the central cell may influence the change in state but it is assumed in strict CA that the temporal dynamics takes care of more distant effects in that growth and decline always imply spatial diffusion. In some systems, such as urban systems, the strictness of this assumption might be questionable. In static spatial modeling where there are clear distance effects such as gravitation, then the adoption of the CA approach almost always requires neighborhoods that do include action-at-a-distance and the issue then is to judge just how relevant CA are in such circumstances in contrast to other, less strict, approaches. A halfway house does exist, however, which embodies some principles of CA but relaxes the neighborhood definition. These are cell–space (CS) models after Albin (1975) and many of the papers in this special issue invoke such relaxations.

The application of CA to urban systems like CA themselves can be traced back to the beginnings of its field, to the first attempts to build mathematical models of urban systems which began in the 1950s. In the postwar years, social physics was in full swing and models of spatial diffusion an important branch of these developments. Hägerstrand as early as 1950 was building diffusion models, specifically of human migration based on the excellent historical population record in Sweden, and these models, although embodying action-at-a-distance through gravitation effects, were close in spirit to the notion of cells being affected by changes in their immediate vicinity (Hägerstrand, 1967). In the early 1960s CA were implicit in the wave of computer models designed for land use–transportation planning. Chapin and his colleagues at North Carolina in their modeling of the land development process articulated cell–space models where changes in state were predicted as a function of a variety of factors affecting each cell, some of which embodied neighborhood effects (Chapin and Weiss, 1968); whereas Lathrop and Hamburg (1965) proposed similar cell-based simulations for the development of western New York State.

However, strict CA models came into our field from another source—from theoretical quantitative geography. These were largely due to Waldo Tobler (1970; 1975; 1979) who during the 1970s worked at the University of Michigan where Arthur Burks and his Logic of Computers group were keeping the field alive. Tobler himself first proposed cell–space models for the development of Detroit but in 1974 formally began to explore the way in which strict CA might be applied to geographical systems, culminating in his famous paper "Cellular geography" published in 1979. At Santa Barbara, in the 1980s, Couclelis influenced by Tobler continued these speculations, until the late 1980s when applications really began to take off as computer graphics, fractals, chaos, complexity all generated the conditions which have led to the kinds of applications reported here.

There are four main characteristic adaptations of strict CA which find their expression in all the papers published here and it is tempting to suggest that these embody inevitable changes to the paradigm when it comes to be applied to urban systems. First, as we have already implied, most applications to urban systems relax the neighborhood effect to incorporate action-at-a-distance. This is an obvious extension but it is controversial. Only Clarke and his colleagues in their simulation of urban growth in the San Francisco Bay area and Benati in his use of CA to model two-dimensional Hotelling-like retail location problems use neighborhoods of the Moore – von Neumann kind. Yet it is an open question as to whether CA might be devised for urban systems which do simulate action-at-a-distance as a consequence of the dynamic application of strict neighborhood rules. In short, it does seem possible to think of action-at-a-distance effects as the global consequence of the kinds of spatial diffusion that take place locally, thus reinforcing the application of strict CA.

Second, somewhat surprisingly perhaps, the notion of any cell having only one state at one time is adopted quite uniformly in these applications. Empirically, it is hard to identify a suitable scale for urban systems where everything is reducible to one activity in one cell. Further, if we impose a grid onto the cell – space, then the states which constitute real systems can only be approximated by regular cells and activities are bound to be mixed in some way within a typical cell. In fact, it is hard to see how this assumption of CA might be relaxed without moving to more conventional spatial representations such as those used in zoning for social physics models, for example.

Third, there is a major problem of structuring CA to meet actual rates of change. For example, consider a strict cellular automaton in which any cell is developed if at least one cell is already developed in the Moore neighborhood. If each cell is developed or redeveloped in each time period of simulation, then starting from a single seed the number of cells developed increases as 1, 9, 25, 49, 81, ..., which is clearly the square of the diameter of the occupied space. In terms of the distance from the original seed—the radius R—then this growth is approximately R^2. Of course, to simulate real systems, we need much greater control over such growth rates and thus in several of the models developed here, external mechanisms are incorporated to control the totals associated with such change. In all three of these adaptations, the strict CA formalism is relaxed and it is an open question as to how far one can go and still appeal to the logic of CA which, as Couclelis (1985) has argued many times, is best considered as a metaphor for spatial development.

The fourth specific adaptation of the CA approach to urban systems modeling involves its relationship to geographic information systems (GIS). At one level, there is a very clear and unequivocal parallel and that is to raster-based GIS in which data are represented as a raster or grid-based or pixel-based two-dimensional array. Many types of physical data such as those which are remotely sensed are represented in this way and rather specific GIS have emerged to deal with such data such as IDRISI and GRASS. Moreover, especially but not exclusively with raster systems, the idea of treating data as layers, of overlaying data, and of performing map-based operations on the same pixels or cells of different layers has led to the idea of map algebra which ties nicely into CA. In fact, the rules of any cellular automaton might be considered a kind of map algebra. In these papers, this link is exploited most effectively by Couclelis who describes her work with Takeyama in developing a generalization of map algebra called geo-algebra of which CA is a special case. Wagner too develops the link with raster-based GIS and map algebra whereas White and Engelen show how a purpose-built raster-based GIS incorporating CA can be used to simulate regional development.

The first paper is by Couclelis who undertakes both a review of where CA stands in urban systems modeling as well as a synthesis. But her main contribution is to identify and illustrate how the CA metaphor can be moved from an illustrative pedagogy to an explicit forecasting device through generalizing the idea of space within CA to proximal space and the operations of CA to a more general geo-algebra. The link to traditional operational urban systems models is also explored in some detail by Batty and Xie who show how CA models can be designed by using a general linear dynamics framework into which spatial interaction, Lowry, input – output, and other antecedents can be fitted. They show how such a framework is useful in considering how various kinds of dynamics—births, deaths, and mutations—can be simulated and they illustrate the approach with a variety of automata.

Phipps and Langlois then develop the relationship between CA and parallel processing, thus focusing on the way in which simultaneous dynamics is characteristic of urban systems. As an example, they show how the von Thünen model of land use in a monocentric city can be simulated. The last hypothetical simulations which relate to the long tradition of urban modeling and regional science are provided by Benati who develops a CA approach to Hotelling's (1929) classic location problem in which two traders compete for the share of a single spatial market. Benati generalizes the Hotelling model to two-dimensional space and many competing actors and he shows, using a CA approach, how patterns of market domination emerge which are consistent with ideas concerning self-organization.

The papers then change direction a little in that, in the rest of the issue, GIS and empirical applications rather than model principles and hypothetical simulations are emphasized. Wagner then argues that the commonality between raster GIS and CA provides some real potential for extending both, by showing how the temporal logic of CA might inform static nontemporal GIS and how the idea of raster map layers can be used to extend CA. White and Engelen almost implement Wagner's suggestions in their various island models of which the example of St Lucia is presented. Here they show how CA can be used to forecast future environmental change where the focus is upon map layers which are largely physical and raster-based. A slightly more conventional approach is taken by Clarke, Hoppen, and Gaydos in their model of the historical urbanization of the San Francisco Bay area. Although several authors writing in this issue show how deterministic CA must be adapted to deal with probabilities of location, this becomes essential for simulating large-scale urban development as Clarke and his colleagues show. Their model is a well-balanced illustration of CA in that the effect of physical factors such as slope, and hard constraints which exclude development, are central to where development takes place.

The last two papers emphasize how CA can be used to model self-organization. Portugali, Benenson, and Omer continue their work on their City series of models which represent a successive elaboration of the basic probabilistic CA model. In City-2, which is the model developed here, they show how migration can be mapped onto the basic urban landscape which is simulated by City. They emphasize two properties: how activities cluster and segregate and how cognitive dissonance between actors complicates the dynamics of location. In the last paper, the French group based at the University of Paris I illustrate how CA logic can be extended to deal with multi-agents. Their model generates patterns of settlement within an hexagonal tessellation of cells, consistent with power laws such as rank size, which in turn indicate some form of self-organized criticality (Bak, 1996). The model also focuses upon transitions between different urban regimes as a necessary characteristic of urban evolution.

This issue was first suggested when Marc Eichen organized a special session on CA in geographical modeling at the San Francisco Annual Meeting of the Association of American Geographers in April 1994. Another meeting held in the SRATEMA Laboratory of the University of Venice Institute of Architecture in November 1994 indicated how live the field was becoming and the editors decided to take some papers from each of these events as the basis of the special issue. We added papers from the Paris (Sanders et al), Tel Aviv (Portugali et al), and Ottawa (Phipps and Langlois) groups as well as Benati's paper which was submitted to the journal independently. Since then, several other papers on CA have been submitted and these will be published over the next year. As editors of both the journal and this special issue, we would like to encourage papers on these new approaches to spatial simulation, to self-organization, and to urban morphology. Considered critiques, constructive advances, and empirical applications are all necessary to progress these ideas.

M Batty, University College London
H Couclelis, University of California at Santa Barbara
M Eichen, City University of New York

References

Albin P, 1975 *The Analysis of Complex Socio-Economic Systems* (Lexington Books, Lexington, MA)
Bak P, 1996 *How Nature Works: The Science of Self-Organized Criticality* (Copernicus – Springer, New York)
Burks A W (Ed.), 1970 *Essays on Cellular Automata* (University of Illinois Press, Urbana, IL)
Chapin F S, Weiss S F, 1968, "A probabilistic model for residential growth" *Transportation Research* **2** 375 – 390
Couclelis H, 1985, "Cellular worlds: a framework for modeling micro – macro dynamics" *Environment and Planning A* **17** 585 – 596
Gardner M, 1970, "The fantastic combinations of John Conway's new solitaire game 'Life'" *Scientific American* **223** 120 – 123
Hägerstrand T, 1967 *Innovation Diffusion as a Spatial Process* (University of Chicago Press, Chicago, IL)
Holland J, 1975 *Adaptation in Natural and Artificial Systems* (The University of Michigan Press, Ann Arbor, MI); reprinted and published in 1992 by MIT Press, Cambridge, MA
Hotelling H, 1929, "Stability in competition" *Economic Journal* **39** 41 – 57
Langton C, 1989, "Artificial life", in *Artificial Life* Ed. C Langton (Addison-Wesley, Reading, MA) pp 1 – 47
Lathrop G T, Hamburg J R, 1965, "An opportunity – accessibility model allocating regional growth" *Journal of the American Institute of Planners* **31** 95 – 103
Macrae N, 1992 *John von Neumann* (Pantheon, New York)
Poundstone W, 1985 *The Recursive Universe* (William Morrow, New York)
Tobler W, 1970, "A computer movie simulating urban growth in the Detroit region" *Economic Geography* **26** 234 – 240
Tobler W, 1975, "Linear operators applied to areal data", in *Display and Analysis of Spatial Data* Eds J Davis, M McCullaugh (John Wiley, New York) pp 14 – 37
Tobler W, 1979, "Cellular geography", in *Philosophy in Geography* Eds S Gale, G Olsson (D Reidel, Dordrecht) pp 379 – 386
Toffoli T, Margolus N, 1987 *Cellular Automata Machines: A New Environment for Modeling* (MIT Press, Cambrige, MA)
Ulam S, 1976 *Adventures of a Mathematician* (Charles Scribner's Sons, New York)