Cellular Automata and Urban Form: A Primer

Michael Batty

Artificial processes for locating urban activities based on simple rules pertaining to local circumstances give rise to complex global patterns that mirror the spatial organization of cities. These systems are called Cellular Automata (CA). They provide a useful means of articulating the way highly decentralized decision-making can be employed in simulating and designing robust urban forms. CA can be easily programmed in a variety of software, and as such provide a suggestive way of exploring actual as well as optimal patterns and plans. This primer provides a pedagogic guide to these ideas and to potential computer applications.

Batty is Director of the Centre for Advanced Spatial Analysis and a professor in the Bartlett School and Department of Geography, University College, London. He has been a professor and Chair of the Department of City Planning in the University of Wales at Cardiff, where he also acted as Dean (1979–1989), and was Director of the National Center for Geographic Information and Analysis at the State University of New York at Buffalo, New York from 1990 until 1995. His most recent books are *Fractal Cities* (Academic Press, 1994) and *Spatial Analysis: Modelling in a GIS Environment* (Pearson GeoInformation, 1996). He is the editor of the journal *Environment and Planning B.*

Journal of the American Planning Association, Vol. 63, No. 2, Spring 1997.

Sketch Planning as Computational Pedagogy

Computer models of cities either attempt to simulate existing urban form or provide procedures for the design of optimal forms, but rarely both. The mechanisms used to model actual cities usually embody local behavioral descriptions without explicit optimizing,[1] whereas those that produce idealized forms seek to optimize in a more global fashion, often mirroring the viewpoint of the designer. Recently, however, a class of models has emerged that has the potential to represent both. By replacing traditional mathematical functions with rule-based procedures, functions of many kinds can be reduced to rules that mirror how actual systems work and how they might work under idealized conditions. Furthermore, as rule-based systems can be built up from the simplest modules, it is possible to strip real systems down to their fundamentals and concentrate on their essential working.

An excellent example of this discipline is based on a class of models called Cellular Automata, or CA for short. CA are models in which contiguous or adjacent *cells,* such as those that might comprise a rectangular grid, change their *states*—their attributes or characteristics—through the repetitive application of simple rules. CA models can be based on cells that are defined in more than 2 dimensions, but the 2-d form that makes them applicable to cities is the most usual. The rules for *transition* from one cell state to another can be interpreted as the generators of growth or decline, such as the change from an undeveloped to a developed cell or vice versa. This change is a function of what is going on in the *neighborhood* of the cell, the neighborhood usually being defined as immediately adjacent cells, or cells that "in some sense" are nearby. Urban growth and decline in real city neighborhoods provide excellent examples.

CA models were first suggested at the dawn of computer history. The English mathematician, Alan Turing, demonstrated the ideas in some early illustrations of computers that could "reproduce" themselves, but it was the Hungarian-American, John von Neumann, who set the field alight in the 1950s, initiating the scientific study of CA. The first applications, however, were as serious games implemented on computers.[2] In 1970, John Conway presented his Game of Life, in which a cell would be developed if it was adjacent to 3 already developed cells, would remain in the same state, that is survive, if surrounded by 2 or 3 developed cells, but would die otherwise. The game was popularized through Martin Gardner's column in *Scientific American,* but it is only since computers have become truly graphic within the last decade, and since the current obsession in science with neo-Darwinism and complexity theory has reached fever pitch, that the area has really taken off. Recently, the field of "Artificial Life" has emerged around chaos theory and nonlinear dynamics, using the mechanisms of CA, and there are several applications to cities and other spatial ecologies under active development.[3]

The Generic Development Principle

What is attractive about CA is that they embody a principle of generic development that fits very well with the way systems in general and city systems in particular appear to develop, or might be developed. In essence, the state of a cell changes if something does or does not happen in its neighborhood, however that be defined. The principle can be stated in its most general form as

IF *something* happens in the *neighborhood* of a *cell*
THEN *some-other-thing* happens to the *cell.*

Another clause is often added, dealing with the case where the conditional is not met—the **ELSE** clause—but this complicates the logic, and in elementary expositions we can proceed without it.

To give the principle specific meaning, we need to define the *things,* the *cells* and the *neighborhood.* The cells might be sites for development; the things, "states" or types of development; and the neighborhood, regions where development might take place. The rules effecting transition between states would thus imply growth, decline, or simply a change in state. Following Poundstone (1985), Conway's Game of Life can thus be translated into three decision rules:

IF there are 3 cells developed in the 8 cells (the neighborhood) adjacent to the cell in question
THEN the cell is developed.

IF there are 2 or 3 cells developed in the neighborhood
THEN the cell remains in its existing state.

IF there are fewer than 2 or more than 3 cells developed
THEN the cell is emptied of any development.

This implies that if there is no cell or only one cell developed in the neighborhood, development in the cell in question dies or the cell remains empty—through isolation; while if the cell has 4 or more developed neighbors, it also dies or remains empty—through overcrowding. Note that three **IF-THEN** rules are needed, although by using the more complicated **IF-THEN-ELSE** form we could write this in a more parsimonious way.

There are many ways this principle can be elaborated to achieve realism. Many **IF-THEN** rules might be concatenated; the size and configuration of neighborhoods can be varied from the most local to the entire system; different types of state or development, such as different land uses and their attributes, might be characterized; and different configurations or starting points for these automata can be defined. The best way to explore possibilities is through examples.

Idealized Development

The simplest of all developments is that of contiguous growth. Imagine a city growing from one cell, the historic core of development. If there is any development in the 8 cells that form the square neighborhood around the cell (the so-called "Moore" neighborhood), then the cell is developed. The principle is shown in figure 1, and the regular compact growth it generates in figure 2a. However, if growth is allowed only when development in the neighborhood is restricted to one cell, or one or two cells, the resultant patterns are much sparser but nevertheless still regular, as shown in figures 2b and 2c. If the neighborhood is smaller, consisting of only those cells north, south, east, and west (the "von Neumann" neighborhood in figure 1), the resultant pattern, with development restricted in the neighborhood to one cell, is shown in figure 2d. These patterns are reminiscent of ideal cities such as those suggested by Renaissance scholars; they result from the rigorous application of generative rules that organize space exactly and leave no room for any variety in locational decision-making. They mirror the simplicity that idealism, however unwittingly, displays.

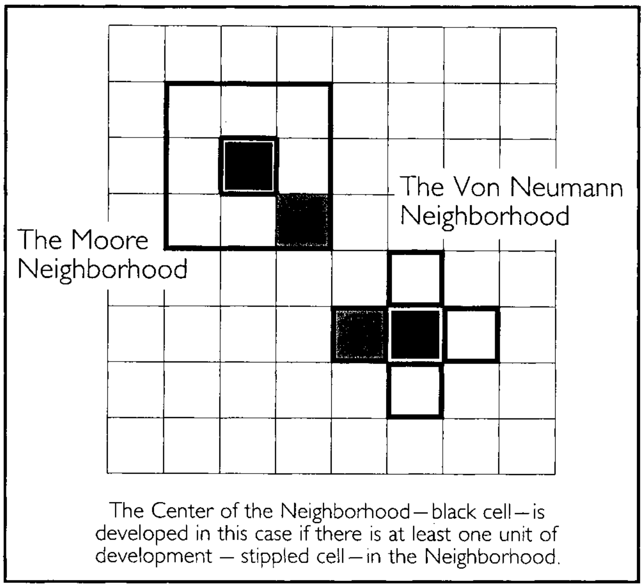

FIGURE 1. Cellular neighborhoods

The compact pattern in figure 2a fills space entirely, and those in 2b to 2d fill lesser amounts. These types of objects are sometimes called "fractals," because their patterns repeat themselves across many scales, due to the fact that their growth reflects the successive application of the local neighborhood principle. They can be measured by a kind of density known as a "fractal dimension" that reflects how much space is filled.[4] In figure 2a, the dimension is 2, the same as the 2-dimensional Euclidean space that is entirely filled; in figures 2b, 2c, and 2d, the dimensions are 1.63, 1.75, and 1.61, respectively. Countless variations in these patterns and their associated fractal dimensions can be generated by choosing and applying different neighborhood principles.

There are many examples of planned cities that can be generated by CA. Figure 3a shows a CA-like picture of Savannah, Georgia in 1734, with the neighborhoods and their subdivisions laid out in the manner mandated by John Oglethorpe, the colony's expedition leader (Reps 1965). Clarence Stein's Radburn layout (Stein 1966) in figure 3b, based on more complex rules, can be generated by two automata working in concert, the first generating the tree-like, rear-access roadways in hierarchical form, the second simply placing development in close proximity to these streets. It is not difficult to suggest rules that generate highly ordered forms such as the idealized utopias of great architects like Vitruvius, Palladio, and Corbusier, and one of the most useful features of CA is to identify exactly what these rules are. But most cities are not so regular, being based on organic growth, and to generate these using CA, some basis for historical accident and "random" decision must be invoked. In short, more realism is required.

Generating Greater Realism

The plans illustrated in figures 2 and 3 all embody the principle that local action generates global pattern: decisions made locally, which are not coordinated centrally in any way, generate patterns that appear to have been manufactured by some central intelligence (Couclelis 1987). This is an age-old conundrum. It is currently one of the most debated topics in science, and reflects the notion that systems are self-organizing. It also reflects the idea that systems organize themselves from the ground up, so to speak from the grassroots, thus generating hierarchical or fractal organization of the kind seen most clearly in figures 2b, 2c, and 2d. Several extensions of the CA model immediately provide the kind of realism these notions imply, and we will examine two here.

Our CA models have so far generated urban forms that result from the application of local growth rules around a single seed, in this case located at the historical origin of that growth. Some cities do in fact grow from a single source, but most grow from more than one, particularly if a city grows to embrace other growing cities or villages within its emerging urban field. Systems of cities of course depend on a system of seeds as reflected in central place theory, for example; and in some cases, CA models can be extended to show how such systems can be generated spontaneously. Here, however, we simply illustrate in figure 4 how two cities whose local growth is based on the neighborhood rule of figure 2b [**IF** there is one unit of development in the neighborhood of an empty cell, **THEN** that cell gets developed] grow and merge together; from this example, it is obvious how quite complex forms can emerge, not only if the number of seeds is increased, but also if the initial shape of those seeds is varied.

Real cities, of course, rarely display such geometric order, and if CA is to provide a useful mechanism for urban development, then some way of introducing more variety is required. This can be achieved using the same kinds of neighborhood rules as those we have presented, but making the choice to develop (or not) dependent upon probabilities. For example, in the simplest model illustrated in figure 2a, if the cell in question has a finite probability of being developed, let us say 80 percent, then over many runs of the model, only 80 percent of the cells that meet the requirements posed by the neighborhood rule will be developed. But because we never know at what particular time a cell might be developed, the rule must be operated randomly. Thus the resultant patterns will con-

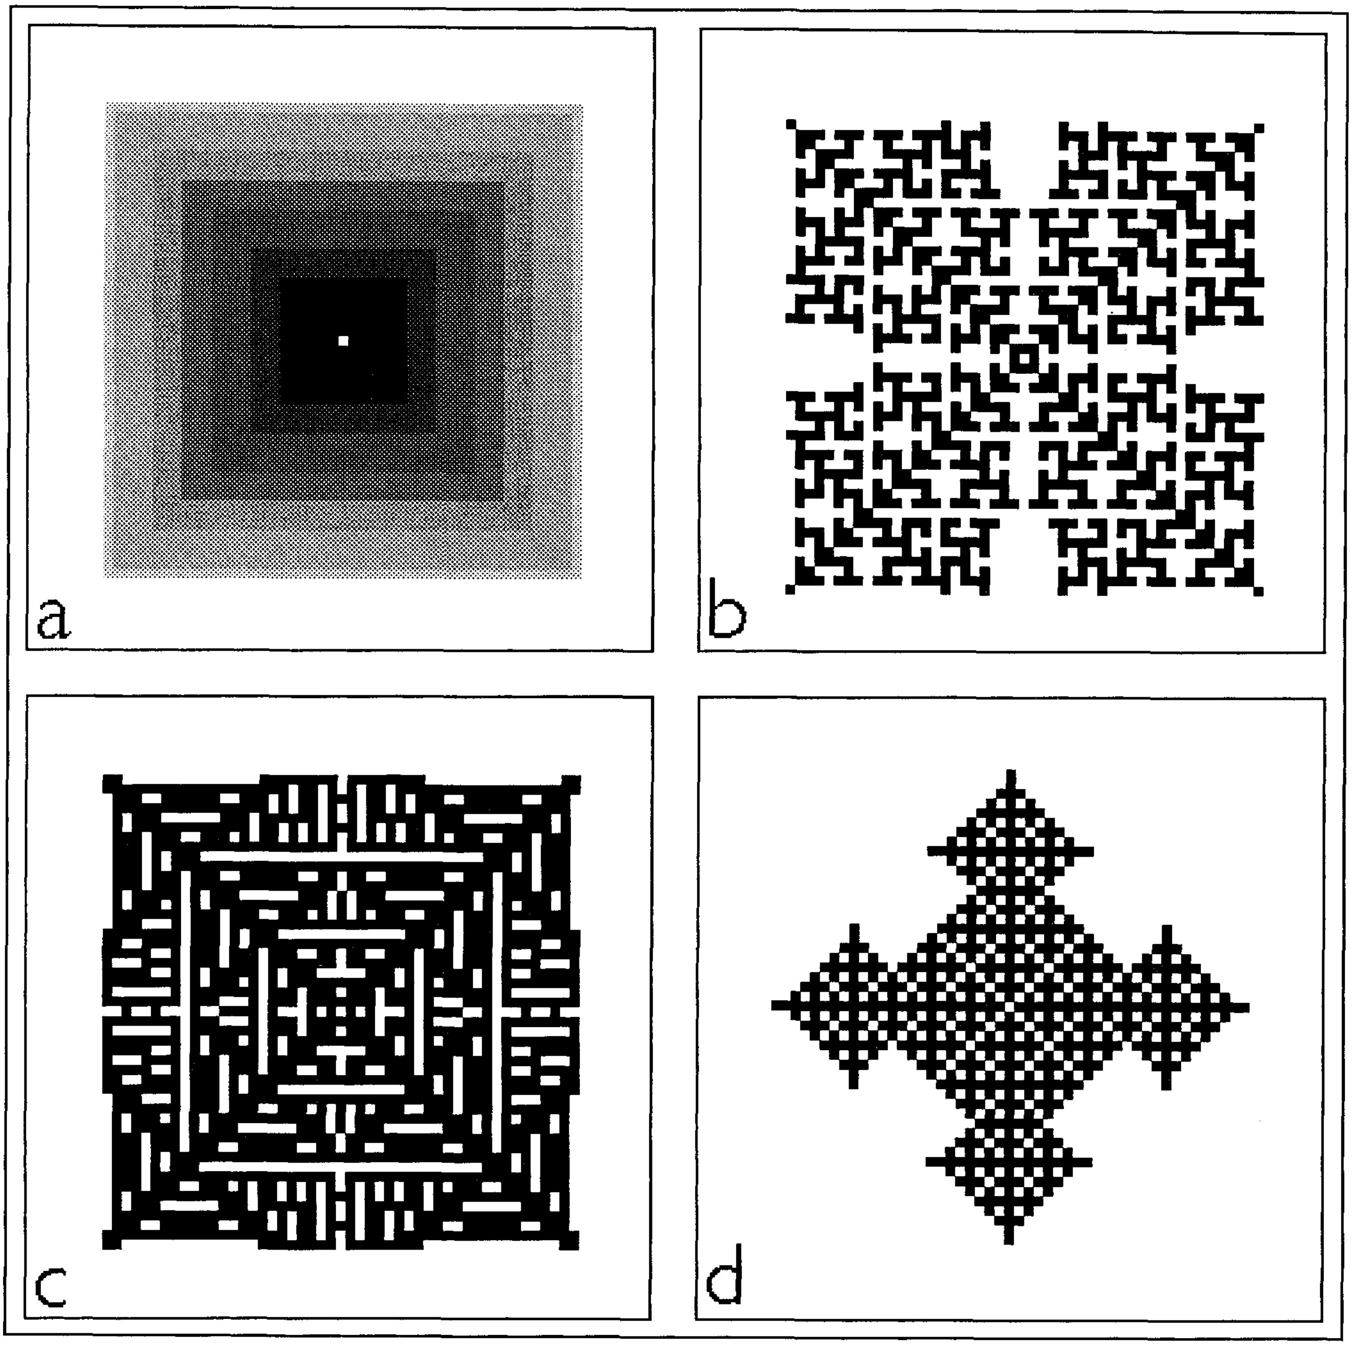

FIGURE 2. Regular cellular automata generating idealized geometries

(a) Development occurs in a cell if there are *one or more* units of development in the Moore neighborhood (grey tones show chronology of development).
(b) Development occurs if there is only *one* unit of development in the Moore neighborhood.
(c) Development occurs if there are *one or two* units in the Moore neighborhood.
(d) Development occurs if there is only *one* unit in the von Neumann neighborhood.

tain randomly located vacant sites. We have modified the rule that generates the entirely compact urban form in figure 2a in the following way:

IF there is at least one developed cell in the Moore neighborhood around the cell in question

THEN the cell is developed with a probability ρ.

Whether it is developed depends on the standard simulation technique of drawing a random number, and if that number is less than ρ, then development takes place. If the cell is not developed, we can also intro-

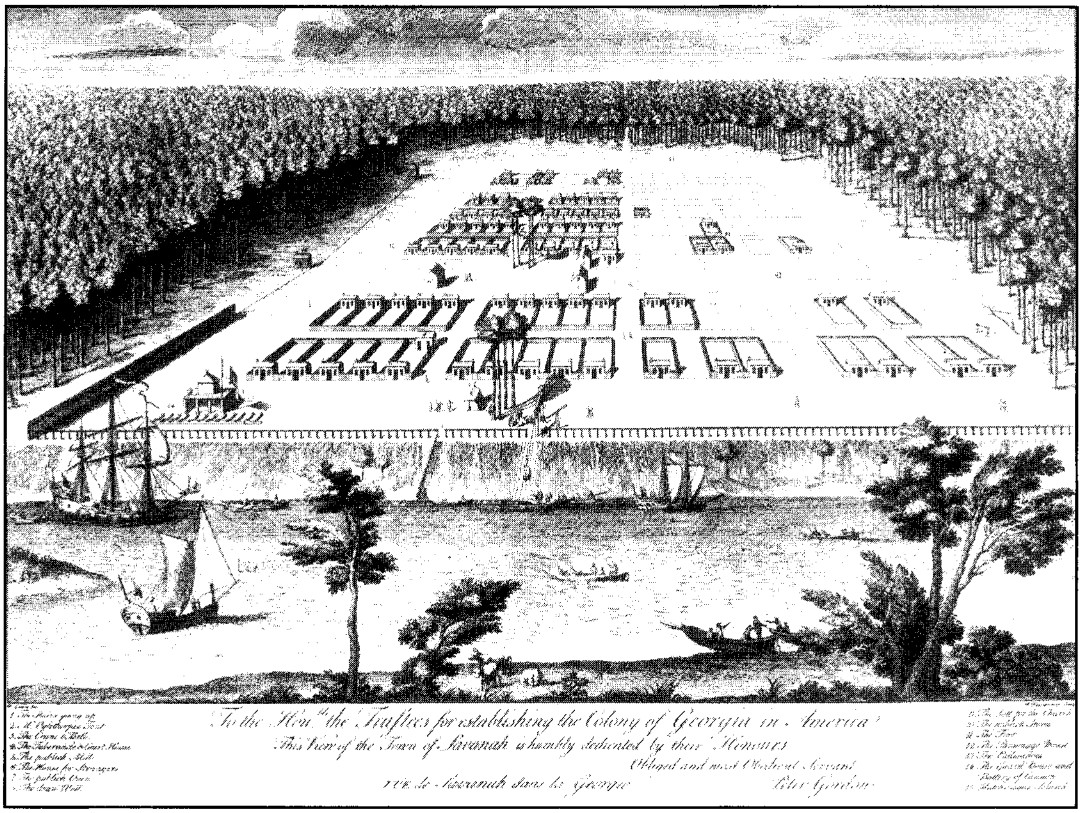

(a) Savannah pictured in 1734 as laid out by the first settlers (courtesy of Historic Urban Plans, Inc., Ithaca, New York)

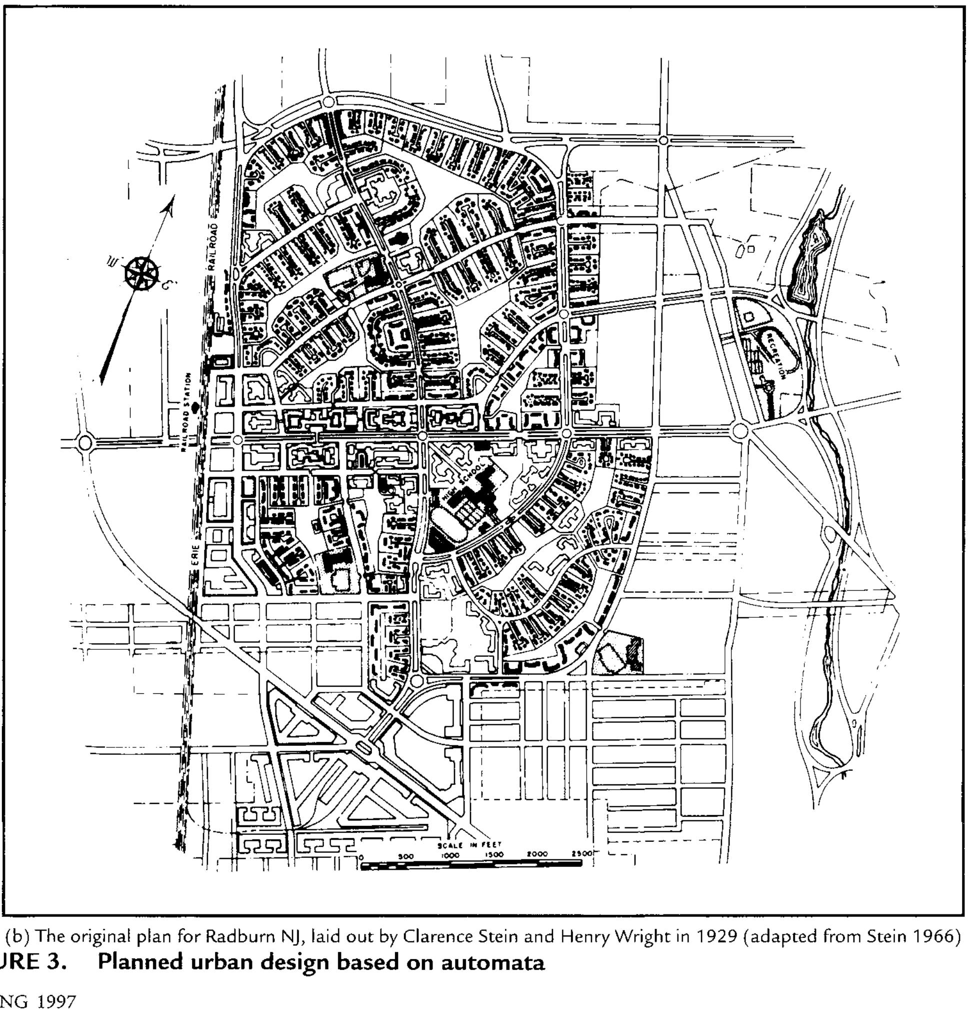

(b) The original plan for Radburn NJ, laid out by Clarence Stein and Henry Wright in 1929 (adapted from Stein 1966)

FIGURE 3. Planned urban design based on automata

duce a modification to the rule that says that the next time the cell is considered for development, the probability of development is lower, given by ρ*ρ, the third time by ρ*ρ*ρ, and, in general, the nth time by ρ^n. This ensures that a level of vacancy exists that reflects the operation of the land market. Simulations are shown in figures 5a and 5b, for two levels of probability.

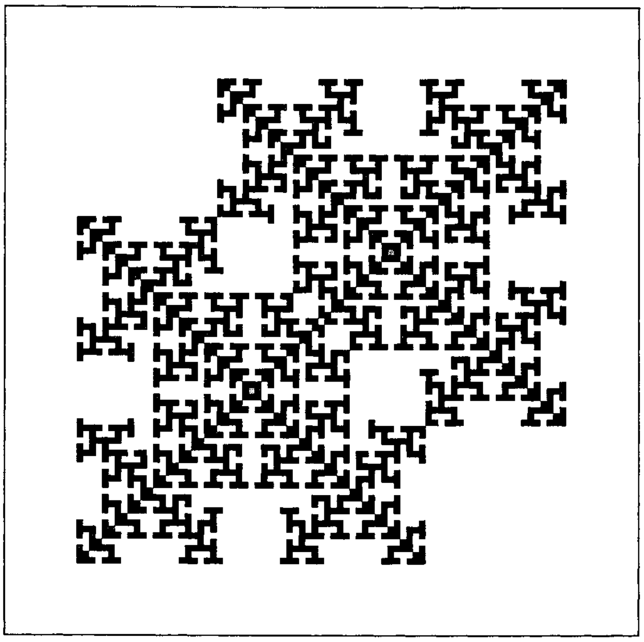

FIGURE 4. Two urban automata merging from different seed sites

We can also show how to introduce the kinds of constraints on land posed by difficult terrain that clearly influence development. In figure 6, we illustrate a simulation of urban growth astride the United States-Canadian border in the Buffalo-Niagara Falls area, where we plant seeds at Buffalo, Niagara Falls, and St. Catherine's. Using the probabilistic CA model just outlined, with ρ = 0.55, we are able to generate the patterns shown in figure 6. Development is constrained to occur on land, an obvious enough requirement, but also the probabilities are lowered if development moves to higher ground.[5]

Many simulations within this framework are possible. In fact, with the software we are using here it is possible to change the rules as the simulations proceed, on the fly, so to speak, thus combining planned and organic forms. Figure 7 shows a simulation where we begin using the sparse but regular form based on the neighborhood rule generating figure 2b, then change to that based on the compact form in figure 2a, finishing with the random rule that invokes vacant land as in figure 5a. CA is a completely general tool, and as such provides a sketch pad for wide experimentation in the manner alluded to in these examples.

Software and Applications

CA models can be implemented within many types of software. It is possible to program them in

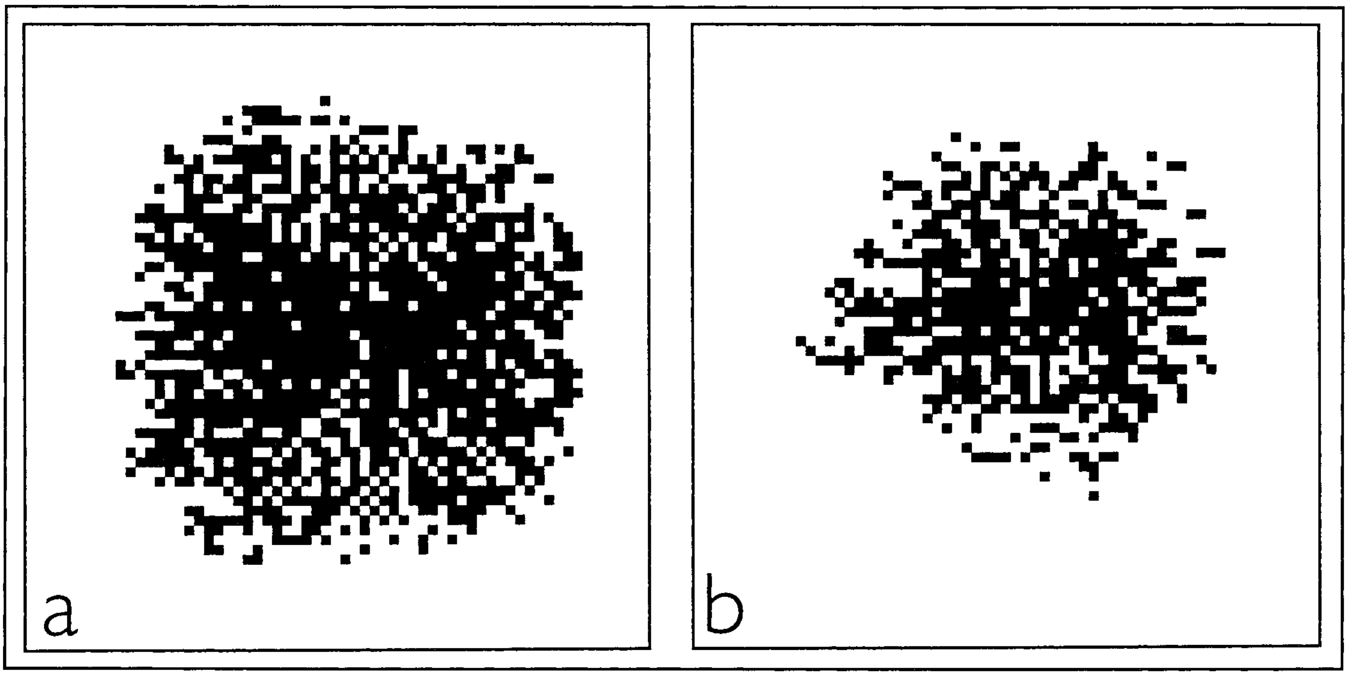

FIGURE 5. Introducing randomness in automata for (a) ρ = 0.8 and (b) ρ = 0.5

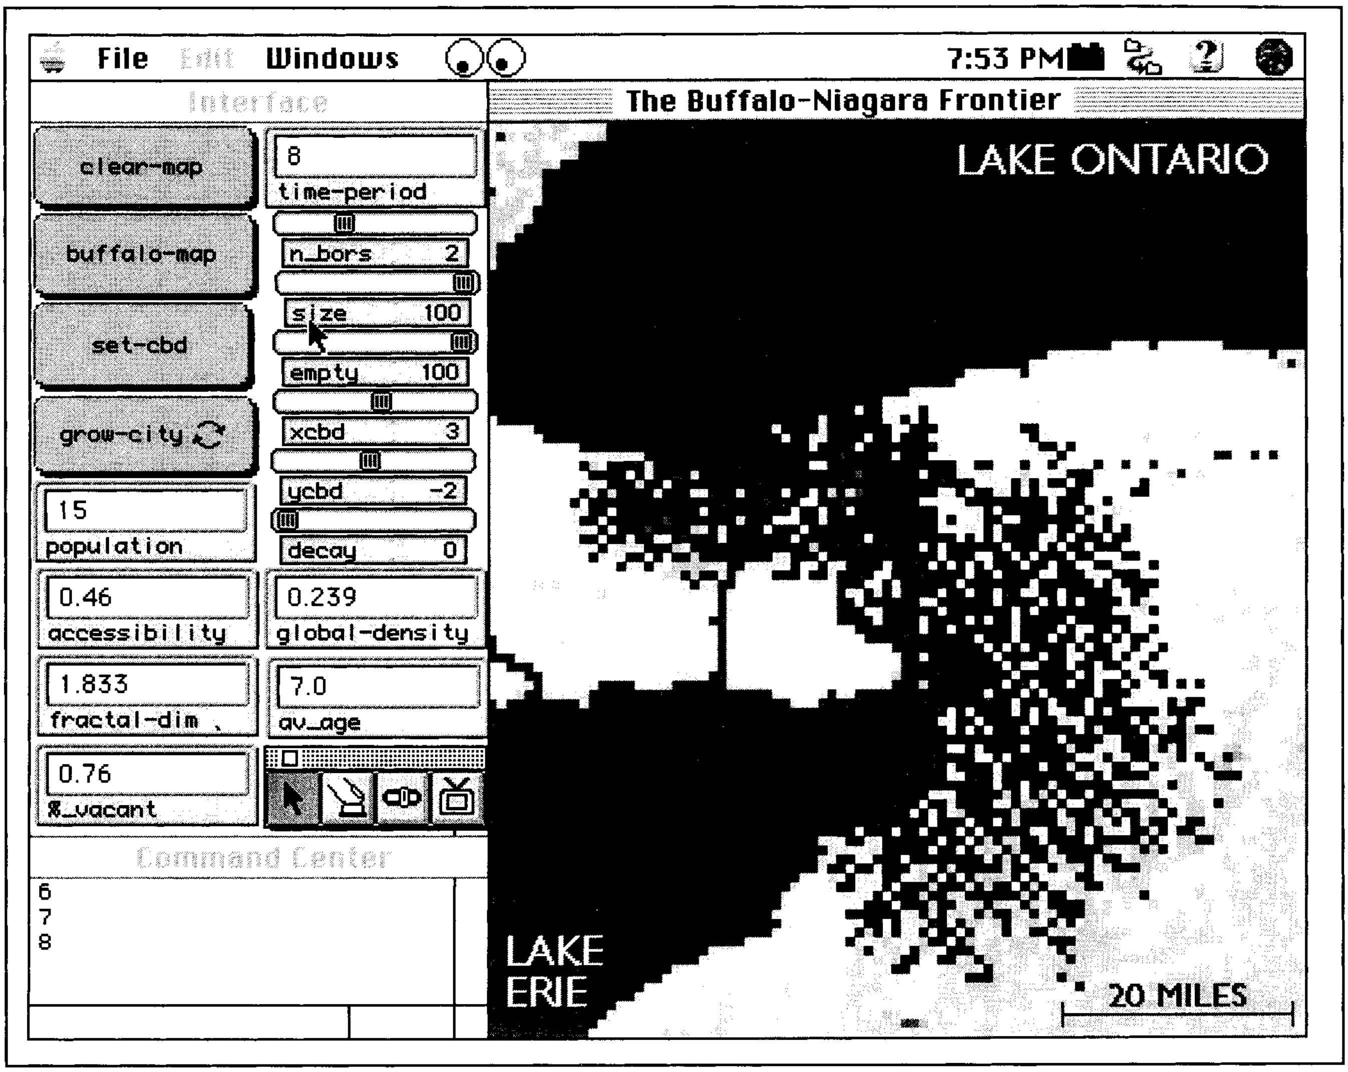

time-period 8
n_bors 2
size 100
empty 100
xcbd 3
ycbd -2
decay 0
population 15
accessibility 0.46
global-density 0.239
fractal-dim 1.833
av_age 7.0
%_vacant 0.76

6
7
8

FIGURE 6. Simulating large-scale urban form, using CA

spreadsheets such as *Excel, LOTUS 1-2-3,* and *Quattro-Pro,* using the chart function as a means of displaying the 2-dimensional grid, and this is aided even further if the modules for plotting maps within the spreadsheet are available. Cartwright (1993) in fact shows how the Game of Life can be programmed in this way. Most desktop GIS and CAD systems now have their own programming languages, and CA models can be easily developed that use the extensive graphics capabilities of these packages. CA models have been developed within the raster GISs *Idrisi* and *MAP II* (Itami 1988; Itami and Clark 1992), although the object-oriented *Avenue* language within *ArcView 3,* or *mini-Pascal* within *MiniCAD* provide even faster programs and more user-friendly interfaces. As packages become more open to one another, and as software continues to incorporate ever more generic programming capabilities and display functions, dedicated software for developing CA becomes less important.

Nevertheless, purpose-built software is still likely to be best. A recent but now somewhat dated review is given by Hielbeler (1990), but the field is developing rapidly, and as CA modeling has very wide applicability, it is impossible to generalize about appropriate software. In the examples developed here, a graphics language based on a massively parallel *Logo,* called *StarLogo,* developed at the MIT Media Lab, has been used, and all the computations illustrated here have been generated using this package (Silverman et al. 1995). *StarLogo* is illustrative of several packages that can be used to program CA. It consists of a series of high-level commands that when issued apply to all the objects—cells in the universe—that invariably are the pixels of the 2-dimensional screen. These rules are in the

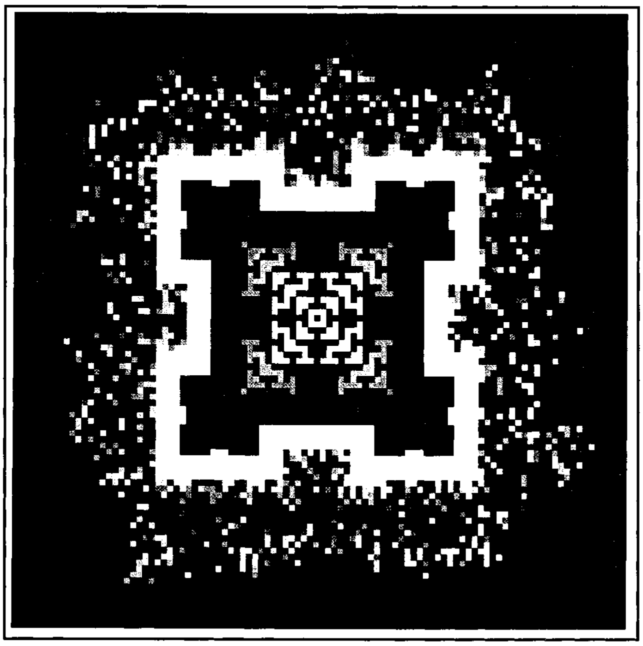

FIGURE 7. Regular to random cellular simulation

form of **IF-THEN-ELSE** statements and provide an accessible way of developing simple CA models, which can be set up to reflect simulations of actual development, or plans of idealized development.

Applications of CA models to cities date from the first time computers were used to model urban land use. Lathrop and Hamburg's (1965) simulation of the Buffalo region, presented in this journal 30 years ago, proposed a CA-like simulation, as did Chapin and Weiss (1968) in their simulation of land development in Greensboro, NC. Since then, the development of raster-based GIS for overlay analysis in landscape planning has been based on a CA-like logic (Itami 1988). Tobler's (1970) simulation of the growth of Detroit is still the seminal geographical application, but much more recently, applications to Cincinnati (White and Engelen 1993), East Amherst, NY (Batty and Xie 1994), the Rhone region in southern France (Mikula et al. 1996), Ottawa-Hull in Ontario (Langlois and Phipps 1995), the San Francisco Bay Area (Kirtland et al. 1994), and Tel-Aviv (Portugali, Benenson, and Omer 1994) are all noteworthy. However, as CA is an exploratory modeling technique, it is more suited to assessing the effect of simple principles for idealized development of the kind shown here and illustrated by Cecchini and Viola (1992) than it is to full-scale simulation. As the software becomes more and more user-friendly, applications to finer and finer scales are likely, with the cutting edge for this kind of logic being more appropriate for urban design than for large-scale strategic planning.

Further Reading

A good general overview of CA can be found in the book by Coveney and Highfield (1995); the best technical, but nevertheless gentle introduction is by Toffoli and Margolus (1987). Applications to artificial life are described in Emmeche (1994). A very powerful but readable account linking CA to decentralization, complexity, and self-organization is in Resnick's (1994) wonderful little volume, which incidentally also contains an excellent nontechnical introduction to the language *StarLogo*. Details of *StarLogo* and how to download the program can be found on the webpage http://lcs.www.media.mit.edu/groups/el/Projects/starlogo/ As good an introductory account as any of the Game of Life is the original article by Gardner (1970), but a much fuller and readable exposition (as well as a computer program in BASIC) is given by Poundstone (1985). Its application to geographical systems is explored in the chapter by Tobler (1979), and several applications to cities are contained in the special issue of *Environment and Planning B: Planning and Design*, 23, issue 2 (1997). Relations between CA and fractals are illustrated in Batty and Longley's (1994) book.

NOTES

1. There is one class of urban models that is the exception—those based on the theories of individual economic optimization that lie at the heart of classical micro-economic theory. Spatial interaction modeling has been much influenced by such theory, and a good exposition can be found in Wilson et al. (1981).
2. Conway first played the game as a many-player game of draughts (checkers) in his Cambridge College common room, but he swiftly moved to computer applications once he realized the complexity of this kind of simplicity. (See Poundstone [1985].)
3. Good summaries of this emergent field are contained in most of the popular expositions of complexity theory, but Emmeche (1994) and Sigmund (1993) provide more detail.
4. Density and fractal dimension are closely related. In strict fractals such as these, the relation between the density of development $V(r)$ and distance r from the seed site is given by a power law of the form $V(r) \sim r^{D-2}$. D is the fractal dimension that always lies between 1 and 2, thus measuring the extent to which the fractal fills the 2-dimensional space available to it.
5. This application shows how *Starlogo,* which is discussed in the main text below, can be used to set up various parameters of the simulations. The slider bars are used to set parameter values such as the x and y coordinates of seeds (xcbd and ycbd), the size of the cluster to be grown, the number of developed cells (n_bors) in the neighborhood, the probability that the development re-

mains empty, and the probability that a developed cell decays to vacancy.

REFERENCES

Batty, Michael, and Paul Longley. 1994. *Fractal Cities: A Geometry of Form and Function.* San Diego, CA: Academic Press.

Batty, Michael, and Yichun Xie. 1994. From Cells to Cities. *Environment and Planning B 21:* s31–s48.

Cartwright, Timothy. 1993. LIFE: The Game. In *Spreadsheet Models for Urban and Regional Analysis,* edited by Richard E. Klosterman, Richard K. Brail, and Earl G. Bossard. New Brunswick, NJ: Center for Urban Policy Research.

Cecchini, Arnaldo, and Filipo Viola. 1992. Ficties—Fictitious Cities: A Simulation for the Creation of Cities. A paper presented to the *International Seminar on Cellular Automata for Regional Analysis.* Venice, Italy: DAEST, Universitario di Architectura.

Chapin, F. Stuart, and Shirley F. Weiss. 1968. A Probabilistic Model for Residential Growth. *Transportation Research* 2: 375–90.

Couclelis, Helen. 1987. Cellular Dynamics: How Individual Decisions Lead to Global Change. *European Journal of Operational Research* 30: 344–6.

Coveney, Peter, and Roger Highfield. 1995. *Frontiers of Complexity: The Search for Order in a Chaotic World.* New York: Ballantine Books.

Emmeche, Claus. 1994. *The Garden in the Machine: The Emerging Science of Artificial Life.* Princeton, NJ: Princeton University Press.

Gardner, Martin. 1970. The Fantastic Combinations of John Conway's New Solitaire Game 'Life.' *Scientific American* 223: 120–3.

Hielbeler, David. 1990. A Brief Review of Cellular Automata Packages. *Physica D* 45: 463–76.

Itami, Robert M. 1988. Cellular Automatons as a Framework for Dynamic Simulations in Geographic Information Systems. *GIS/LIS '88 Proceedings.* San Antonio, TX. 590–7.

Itami, Robert M., and John D. Clark. 1992. Spatial Dynamic Simulations Using Discrete Time and Discrete Event Theory in Cell Based GIS Systems. *Proceedings: 5th International Symposium on Spatial Data Handling.* Charleston, SC. 702–12.

Kirtland, David, Len Gaydos, Keith Clarke, Lee De Cola, William Acevedo, and Cindy Bell. 1994. An Analysis of Transformations in the San Francisco/Sacramento Area. *World Resource Review* 6: 206–17.

Langlois, Alan, and Michel Phipps. 1995. *Cellular Automata, Parallelism and Urban Simulation.* Ottawa, Ontario, Canada: Final Report of the SUAC Project, Department of Geography.

Lathrop, George T., and John R. Hamburg. 1965. An Opportunity-Accessibility Model for Allocating Regional Growth. *Journal of the American Institute of Planners* 31: 95–103.

Mikula, Boguslav, Helene Mathian, Denise Pumain, and Lena Sanders. 1996. Integrating Dynamic Spatial Models with GIS. In *Spatial Analysis: Modelling in a GIS Environment,* edited by Paul Longley and Michael Batty. Cambridge, UK: Geo-Information International. 283–295.

Portugali, Juval, Izak Benenson, and Izak Omer. 1994. Socio-Spatial Residential Dynamics: Stability and Instability within the Self-Organizing City. *Geographical Analysis* 26: 321–40.

Poundstone, William. 1985. *The Recursive Universe.* New York: William Morrow and Company.

Reps, John. 1965. *The Making of Urban America.* Princeton, NJ: Princeton University Press.

Resnick, Mitchell. 1994. *Turtles, Termites and Traffic Jams: Explorations in Massively Parallel MicroWorlds.* Cambridge MA: MIT Press.

Sigmund, Karl. 1993. *Games of Life: Explorations in Ecology, Evolution, and Behaviour.* New York: Oxford University Press.

Silverman, Brian, et al. 1995. *StarLogo Reference Manual.* Cambridge, MA: Epistemology and Learning Group, The Media Lab, MIT.

Stein, Clarence. 1966. *New Towns for America.* Cambridge, MA: MIT Press.

Tobler, Waldo. 1970. A Computer Movie Simulating Urban Growth in the Detroit Region. *Economic Geography* 26: 234–40.

Tobler, Waldo. 1979. Cellular Geography. In *Philosophy in Geography,* edited by S. Gale and G. Olsson. Dordrecht, Holland: D. Reidel Publishing Company. 379–86.

Toffoli, Tommaso, and Norman Margolus. 1987. *Cellular Automata Machines.* Cambridge, MA: MIT Press.

White, Rodney, and Guy Engelen. 1993. Cellular Automata and Fractal Urban Form: A Cellular Modelling Approach to the Evolution of Urban Land Use Patterns. *Environment and Planning A* 25: 1175–93.

Wilson, Alan G., Jose D. Coelho, Sally M. Macgill, and Huw C. W. L. Williams. 1981. *Optimization in Locational and Transport Analysis.* Chichester, UK: John Wiley.