

Research Article

Loose-coupling a cellular automaton model and GIS: long-term urban growth prediction for San Francisco and Washington/Baltimore

KEITH C. CLARKE

University of California, Santa Barbara, Santa Barbara, CA 93106 USA

and LEONARD J. GAYDOS

United States Geological Survey, EROS Data Center, 242-4 NASA-Ames
Research Center, Moffett Field, CA 94035 USA

Abstract. Prior research developed a cellular automaton model, that was calibrated by using historical digital maps of urban areas and can be used to predict the future extent of an urban area. The model has now been applied to two rapidly growing, but remarkably different urban areas: the San Francisco Bay region in California and the Washington/Baltimore corridor in the Eastern United States. This paper presents the calibration and prediction results for both regions, reviews their data requirements, compares the differences in the initial configurations and control parameters for the model in the two settings, and discusses the role of GIS in the applications. The model has generated some long term predictions that appear useful for urban planning and are consistent with results from other models and observations of growth. Although the GIS was only loosely coupled with the model, the model's provision of future urban patterns as data layers for GIS description and analysis is an important outcome of this type of calculation.

1. Introduction

The human geographical processes of urbanization and urban spread will apparently continue unabated into the twenty-first century. By the last decade of this century (1995), 21 world cities had total metropolitan area populations of over 6 million, led by Tokyo(30 300 000), New York(18 087 251), Seoul(15 850 000), Mexico City(14 100 000) and Moscow(13 150 000). Also by 1995, the number of people worldwide living in settlements of five thousand or more reached 51%, a majority of humankind and a dramatic increase from 29% in 1950. Gottmann (1961) coined the term *megalopolis* to describe the coalescence of metropolitan areas in the northeastern United States. In the era of GIS, remote sensing and digital map products have recorded the birth and growth of similar megalopolises in California and in Mexico, South America, Europe, and Asia. New estimates of the world population in 2100AD indicate an increase from the present population of 5.5 billion to 10 to 20 billion. We have termed the resultant super cities *gigalopolis*, the twenty-first century system of world cities containing billions of people centered on the world's major urban areas. Gibson's fictional view of the future urban United States, for example, talks only of the 'East Sprawl' and the 'West Sprawl' (Gibson 1984).

The magnitude of *gigalopolis* in population terms, however, understates the most critical permanent impact of increased urban-space consumption, which is usually at the expense of prime agricultural land essential for food production. Looked at spatially, each expanding metropolitan area will become both physically and virtually

connected to other growing concentrations of people in the coming century through raw gain of territory as well as broader communication, transportation, and economic ties. Not since the dominance of agriculture within human affairs millennia ago has humankind's habitat changed so quickly and irreversibly. Simultaneously, recent trends show first that people are consuming more space per person within their urban environment, but also that the average household size is decreasing, at least in western cities. Urbanization and urban growth go hand-in-hand, and generate many other land transitions, with several varied land use types eventually converting to urban use. The spatial consequences of the urban transition deserve serious study by scientists and policy makers concerned with global change because they will impact humankind directly and profoundly. Vitousek (1994) called land use/land cover changes, including the urban transition, one of the few certainties of global change, because we are 'certain that they are going on, and certain that they are human-caused'.

This paper summarizes research conducted to describe, model, and predict future urban transitions. We constructed a model using a cellular automaton that simulates the urban growth process. We calibrated the model with historical data for two major metropolitan regions in the United States, and used it to produce one-hundred-year projections of their urban growth. GIS has been an indispensable tool in the model construction and calibration, and will play far more critical a role when the predictions are distributed and reproduced for other areas. The broader purpose of this work is to model and predict the spatial consequences of future urbanization, so that the impact of human-induced land transformations can be better understood.

Cellular modelling grew out of earlier environmental simulation work on wildfire behaviour (Clarke *et al.* 1995), that in itself was based on pioneering work by Michael Batty (Batty and Xie 1994). The role of cellular automata as potential powerful contributions to urban process modelling was demonstrated by Couclelis (1997) and Takeyama and Couclelis (1997). Influencing model choice was the need for a model that was scale independent, so that local, regional and continental scale processes could be described in a single context. Cellular automata are simple models for the simulation of complex systems (Waldrop 1992, Wolfram 1986). A cellular model assumes only an action space (usually a grid), a set of initial conditions, and a set of behaviour rules. Characteristic of such models of complexity is that behaviour is *emergent*, that is, it is generated by repetitive application of the rules beyond the initial conditions. Complex systems are also termed *self-organizing* and are remarkably suitable to computational simulation (Wolfram 1984). Simple cellular automata are characterized by phase transitions between behaviour types, so that a single model can result in stability, stochastic instability or chaos. As such they seem ideally suited to modelling the complexity of urban systems, which typically have many more unknowns than measurable variables. Cellular models are known in ecology as individual based models and this concept, that complex aggregate behaviour results from many interacting self-motivated agents, has great value for both urban modelling and for the data rich environment of GIS. This is especially appropriate in urban modelling, where the process of urban spread is entirely local in nature and aggregate effects, such as growth booms, are emergent.

Modelling cities with cellular automata is a new approach, and one that was virtually impossible without the data management capabilities of GIS and powerful workstation technology. The approach has distant roots in geography in the work of Hagerstrand (1965) and Tobler (1979). Links to prior urban modelling are less

clear (Clarke, Gaydos and Hoppen, 1997), though Batty and Longley (1994), Makse *et al.* (1995) and White and Engelen (1992) have used essentially similar approaches.

2. The role of GIS in urban cellular modelling

Much has been written on the integration of GIS and modelling, especially for environmental issues (Wilson 1995, Goodchild *et al.* 1996, Wagner 1997). In a recent paper, Park and Wagner (1997) implemented a tight coupling of several cellular automaton models (including ours) within Idrisi using the Cellang CA language. In the context of the our cellular model, GIS served at least three important roles, none of which could be called tight coupling. The first of these was as data integrator. In each of the initial applications, data were either already available as ARC/INFO coverages, or were captured by scanning and digitizing (Crawford-Tilly *et al.* 1996). Although the coverages existed, new map extents, projections, and grid resolutions were required, and the GIS was invaluable in ensuring co-registered input data layers for the model. All further modelling and analysis depended on this essential first step, what Chrisman has called a 'universal requirement' for GIS (Chrisman 1997, p. 108). The input data layers were each raster grids, exported as image files and converted by specialized software. Thus the relation of our modelling to GIS is one of loose coupling, as classified and described by Anselin *et al.* (1993).

Secondly, GIS allowed the results to be visualized. This was the weakest component of the loose coupling. Stand-alone versions of software for the model were written to generate map displays, multiple display windows, and results (Clarke *et al.* 1996, Acevedo and Masuoka 1997). The role for the GIS in this instance was one of taking intermediate results, and facilitating their comparison with the original and other layers of information.

Thirdly, the predictions generated were reintroduced into the GIS data sets available for application, allowing planning decisions to be made with the data. An example is the ArcView window interface developed to display the three development scenarios for the Sterling Forest area in New York State (Kramer 1996). It is this third application of GIS that is by far the most powerful from the modelling point of view. Having 'what if' model projections available to perform the more traditional GIS operations and analyses such as buffering and overlay is a very powerful GIS capability. This function alone favours the use of loose coupling. In spite of efforts to build cellular modelling functions into GIS directly (Takeyama and Couclelis 1997; Park and Wagner 1997) and the suitability both of specific GIS packages (e.g. GRASS) and of control languages (e.g. MapObjects and Avenue in ArcView), it is likely that most numerical modelling, especially that requiring exhaustive or rigorous calibration, will need to parallel the GIS rather than work within the software. This approach is similar to the way that statistical functions and spatial analytical capabilities have been integrated with GIS. By broadening the definition of GIS, such as that of GIScience, the model coding can even be viewed as part of the science, if not part of the system (Goodchild 1992, Clarke 1997).

3. Data for the model

Data for the model's calibration and initial use in prediction came from a variety of sources. For the San Francisco Bay area, we assembled data from historical maps and air photos, analog and digital maps for different time periods, from data supplied by such agencies as local governments, and from satellite images (Kirtland *et al.* 1994). After success of the initial model application equivalent data, but to more

specific and documented standards, was compiled for the Washington/Baltimore region (Crawford-Tilley *et al.* 1996). These data are available on-line as part of the Washington/Baltimore Regional Collaboratory (Clark *et al.* 1996).

Data assembly problems included inconsistent feature definitions over time, especially for urban areas and major roads, extensive manual generalization present in historical maps, and the need to integrate multiple image and map sources from different projections, datums, and coordinate systems. Although much of the delineation of urban versus non-urban use had to be interpreted, skilled image interpreters were available to make the interpretation to acceptable standards.

The model's input layers are fourfold: (1) A digital elevation model, for which the GIS was used to create a grid of slopes in percentage, (2) A layer showing the initial or seed configuration of urban areas, plus as many additional historical layers as possible, to calibrate the model, (3) As many historical transportation layers as possible, which the model reads and uses sequentially as their year of construction is reached, (4) A layer of excluded areas unlikely or impossible to urbanize, such as national parks, water bodies and protected wetlands. The latest version for the *gigalopolis* project, in which several additional land cover types such as agricultural and forest are included in the model's behaviour, requires a land use layer as well.

Since the growth rules in this model are defined primarily by physical factors, the San Francisco Bay area was an ideal first test site. Elevations range from sea level to 2500 m, with some clear topographic control over growth by the dichotomy between slopes and flats. The region also is diverse in its patterns of urbanization, reflecting its beginnings in small enclaves clustered around the inland waterway network, the emergence of San Francisco as a transportation hub, and the more recent urbanization of the surrounding valleys closely reflecting the extended highway system. The model's temporal input comprised six raster image maps of urban extent for the years 1900, 1940, 1954, 1962, 1974, and 1990, with 1900 as the initial conditions or 'seed' year. An 1850 data set was too sparse for modelling purposes. Roads were available for 1900, 1920 and 1970. The San Francisco data were included in Clarke *et al.* (1997) and can be viewed on the World Wide Web at <http://edcwww2.cr.usgs.gov/umap/umap.html>.

Similarly, for Washington/Baltimore, urban and road layers were constructed for 1792, 1850, 1900, 1925, 1938, 1953, 1966, 1972, 1982, and 1992 (figures 1 and 2). Again due to the sparse and least reliable early data, 1900 was used as the initial conditions or seed layer. While the Washington/Baltimore Collaboratory digitized layers for railroads also, these were excluded to ensure reliability and compatibility between the two applications (Clark *et al.* 1996).

4. Model description

The urban model is a scale-independent cellular automaton (CA), model with some variations from the traditional CA, and multiple behaviour types. The growth rules are uniform throughout a gridded representation of geographical space and are applied on a cell-by-cell basis. A single time span is one iteration of the CA, and all changes are applied synchronously at the end of each time period. The set of growth rules, and the initial condition (map of urban extent at a period in time) are integral to the data set being used because they are defined in terms of the physical nature of the location under study. In-site calibration then adapts the model to its local environment. One interpretation is that the urban area corresponds to an organism. The initial condition is the start 'seed' layer, growth occurs one cell at a

Two Hundred Years of Urbanization in the Chesapeake Bay Region

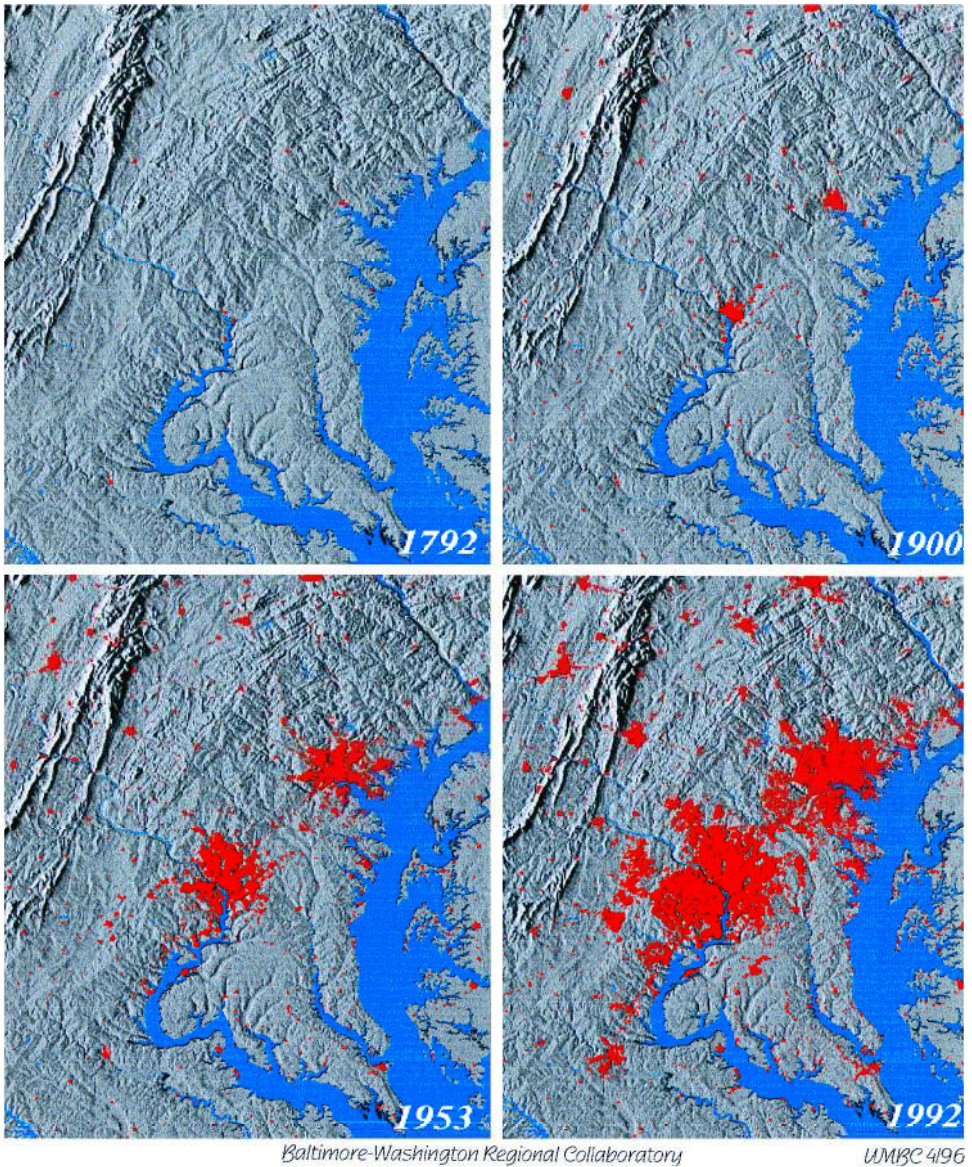



Figure 1. Historical urban development in the Washington/Baltimore area.

time with each cell acting independently of all others, and patterns emerge during growth as the organism ‘learns’ more about its environment.

The model was implemented as a computer program written in the C language. The program operates as a set of nested loops: the outer control loop repeatedly executes each growth ‘history’, retaining cumulative statistical data, while the inner loop executes the growth rules for a single ‘year’. The starting point for urban growth



Transportation System Development in the Washington - Baltimore Region

Figure 2. Historical growth of transportation in the Washington/Baltimore area.

is an input layer of 'seed' cells, the urban extent for a particular year identified from historical maps, atlases, and other sources. The rules apply a cell at a time and the whole grid is updated as the 'annual' iterations complete. The modified array forms the basis for urban expansion in each succeeding year. Potential cells for urbanization are selected at random and the growth rules evaluate the properties of the cell and its neighbours (e.g., whether or not they are already urban, what their topographic slope is, how close they are to a road). The decision to urbanize is based on mechanistic growth rules as well as a set of weighted probabilities that encourage or inhibit growth. The model is described in detail in Clarke *et al.* (1997).

Five factors control the behaviour of the system. These are: a diffusion factor, which determines the overall outward dispersive nature of the distribution; a breed coefficient, which specifies how likely a newly generated detached settlement is to begin its own growth cycle; a spread coefficient, which controls how much diffusion expansion occurs from existing settlements, a slope resistance factor, which influences the likelihood of settlement extending up steeper slopes; and a road-gravity factor which attracts new settlements toward and along roads. These values, which affect the acceptance level of randomly drawn numbers, are set by the user at the outset of every model run.

Four types of growth are possible in the model: spontaneous, diffusive, organic, and road influenced growth. Spontaneous growth occurs when a randomly chosen cell falls in a suitable location for urbanization at the boundary of an existing settlement, simulating the fragmenting influence urban areas have on their surroundings. Diffusive growth permits the urbanization of cells which are flat enough to be desirable locations for development, even if not near an established urban area. Organic growth spreads outward from existing urban cores, representing the tendency of all urban areas to expand. Road influenced growth encourages urbanized cells to develop along the transportation network, reflecting increased accessibility.

A second hierarchy of growth rules, termed self-modification, is prompted by an unusually high or low growth rate above or below a threshold. The growth rate is computed by comparing the number of new pixels urbanized in any time period to the total existing urban area. The limits of 'critical high' and 'critical low' begin an increase or decrease in three of the growth-control parameters. The increase in the parameters is by a multiplier greater than one, 'boom', imitating the tendency of an expanding system to grow ever more rapidly, while the decrease is by a multiplier less than one, 'bust', causing growth to taper off as it does in a depressed or saturated system. However, to prevent uncontrolled exponential growth as the system increases in overall size, the multiplier applied to the factors is slightly decreased or lagged in every subsequent growth year.

Other effects of self-modification are an increase in the road-gravity factor as the road network enlarges, prompting a wider band of urbanization around the roads, and a decrease in the slope resistance factor as the percentage of land available for development decreases, permitting expansion onto steeper slopes. Under self-modification, the parameter values increase most rapidly in the beginning of the growth cycle when many cells are available for urbanization, and decrease as urban density increases in the region and expansion levels off. Without self-modification the model produces linear or exponential growth as long as new land remains available; self-modification generates the typical S-curve growth rate of urban expansion observed within a region.

5. Model calibration

Calibration of the original model was described in Clarke *et al.* (1996). Since then, a new procedure has been developed and tested for the two study areas. After assembly of the various data sets and their conversion to the input format for the model, the calibration problem may be stated: given a starting image of urban extent, which set of initial control parameters leads to a model run which best fits the observed historical data? Pursuit of this question implies that 'best' can be quantified, and used to test statistically the observed against the expected. These parameters are then used for prediction.

We chose four ways to test statistically the degree of historical fit, and are investigating twelve measures in our latest work. The four initial tests are correlation coefficients of predicted outcomes from the model with values computed from the historical map layers. The specific measures were: (1) the *r*-squared fit between the actual and predicted number of urban pixels; (2) the *r*-squared fit between the actual and the predicted number of edges in the images (i.e. those pixels that have contact between urban and non-urban on any side, so that a single isolated urban pixel counts as 4 edges); (3) the *r*-squared fit between the actual and the predicted number of separate clusters in the urban distribution, computed by eroding cluster edges with an image processing routine until all separate blobs collapse onto just one pixel, then counting these pixels; and (4) A modified Lee-Sallee shape index, computed by combining the actual and the predicted distributions as binary urban/non-urban layers, and computing the ratio of the intersection over the union. For perfect correspondence the value is 1.0. Practically, values of about 0.3 were obtained. The four measures were computed as averages of multiple runs. We used 4, 10, and 100 runs in the tests, although more rigorous multi-start tests were conducted in the initial calibration (Clarke *et al.* 1996). We devised a single composite measure, the first three summed and multiplied by the fourth, to rank overall outcomes.

Calibration then proceeded with a UNIX script written in the PERL shell language. The script generated a control file for the program, executed the program so that it iterated over four Monte Carlo runs, wrote log files of its outputs, then executed a stand-alone program that computed the four measures, and wrote these into a master log file. The first of two PERL scripts performed coarse iterations over the control parameters, with increments of 20 units at a time and 25 for Washington/Baltimore. Even so, this procedure involved iterating over the 0–100 range of the diffusion coefficient, the breed coefficient, the spread coefficient, the slope-resistance and the road-gravity control parameters. For San Francisco, six, six, six, five and seven combinations were tested for the parameters respectively, starting in 1900 and terminating at the 1990 data set. This gave *r*-squared values computed for 1940, 1954, 1962, 1974, and 1990, for only five 'observations' although each *r*-squared was a composite of four measures and an average over four model runs. The 7560 combinations executed in about 252 hours of CPU time on a Silicon Graphics Indigo 2 Impact 10000 workstation. The longer runs for the much larger Washington/Baltimore data set mandated fewer combinations, on the order of 3000. This broad coverage phase of the calibration gave those combinations of parameter setting that tended to produce the best 'projections' of the present day.

The initial test was then repeated for unit increments of the control parameters above and below the 'winning' values for the coarse calibration. The additional step or fine calibration involved another 6 x 5 x 5 x 5 combinations of four iterations each for 3000 combinations and another 100 hours of CPU time. Again, for the

larger data set only about 1800 combinations were computed. A final set of parameters was then used for a 100-run iteration, saving the terminating values of the control parameters (changed by the self-modification) as input for a 100-iteration prediction run into the future, stopping at the year 2100.

One obvious problem with this approach is the processing time. Therefore, the Washington/Baltimore data set was used in an experiment to test sensitivity of the calibration process to scale dependence. In the original calibration, at the full resolution and extent of the data set, the mapped area was 486 by 720 pixels with a resolution of 210 m. The entire two step calibration process described above was then repeated at resolutions of 410 m, 820 m, and 1640 m. New data sets were created by direct sub-sampling, i.e. taking every other pixel in each case. To prevent the drop out problem or loss of connectivity in the roads map, the roads were thickened prior to re-sampling. this did indeed cut the processing time. The 1640 m data set was fully calibrated in only 6 hours of CPU time, a 42 fold speed-up.

6. Results

The self-modification internal to the cellular model enhances phase changes. In geographical terms, this means that different periods of time should be dominated by different growth behaviour, and by increasing spatial adaptation to the local environmental conditions. As a result, it was expected that the impact of slope would be more pronounced in San Francisco at first, followed by a 'spilling out' of growth into the Central Valley. On the other hand, topography is clearly less of a constraint to growth in the Washington/Baltimore corridor, so that little difference in behaviour would be expected to result from to the slope factor. There were indeed significant differences in slope behaviour. For example, for the population (number of urban pixels) correlation, the San Francisco fit was high but varied, and increased as the slope resistance increased. On the contrary, in Washington/Baltimore, there was very little variance in response, and fit declined with the slope resistance (figure 3). This implied that slope explains far more of the urbanization in San Francisco than in Washington and Baltimore. While a rather obvious point, it confirms the basic soundness of the model.

Phase changes are obvious from the coarse calibration data. In each case, the five control parameter initial settings were compared to the scores. Some remarkable behaviour was evident within the full range of model applications. For example, for the population r -squared (actual vs. predicted urban extent for each observation period), the mean score decreased as the diffusion coefficient increased, but the




Figure 3. Calibration results: initial values for the model control parameters.

lowest setting had four iterations that scored higher than all others. A phase transition was clearly precipitated by changing the spread parameter. Extremely high correlation (in excess of 0.99 *r*-squared) dropped markedly between spread settings of 1 and 20, but then increased linearly for the remainder of the values. In other words, at a setting between 1 and 20, some interaction was triggered that precipitates a behaviour change. Spatially, for example, this may be the breaking out of the initial constrained urban areas and valleys. The effect is evident, though less clearly, in the Washington/Baltimore data also. Phase transitions were evident in all four of the test measures.

If the success measures can be characterized, then the four are increasingly spatial in nature. The population value is, as expected, the easiest to simulate with the model. Most fits were in the 0.99 range, partly due to the low number of data years. The number of clusters was an indication of how many separate spreading centres or independent communities were present. This value proved a far better discriminator in San Francisco than in Washington/Baltimore. The number of edges was the total number of pixel contact edges between urban and non-urban within the grids. As such, it has already been shown to be influenced strongly by the map source. Remotely sensed data are more speckled than map sources, and thus have more edges. Finally, the Lee Sallee measure is a modified shape index, defined as a ratio of the AND to the OR of the actual and predicted urban images as binary layers. This measure penalizes spatial mismatch twice; for example if a good spatial shape match was displaced slightly, the intersection would be smaller but the union would be larger by twice the displacement area. Values rarely approached even a 30% match, and the ability to discriminate spatial match with this measure was quite good. Again, the spread-parameter phase change was evident, both in San Francisco and Washington/Baltimore. In a few cases, evidence pointed to a local maximum, for example in the breed coefficient for both data sets.

The second round of calibration selected the best outcome from the coarse calibration and iterated single parameter-value increments above and below. Not all values were permuted, since some were already clearly at maxim. Most striking from the fine tuning results was the scaling nature of the model. Phase transitions, trends, maxim and even oscillations (e.g. breed in Washington/Baltimore for the Lee-Sallee measure) were evident. While in each case a single value was chosen as maximum, in fact many individual runs can attain far higher success, in a specific measure and overall. Nevertheless, stability across many runs was considered important for Monte Carlo modelling. We made 100 permutations of the final predictive runs. This was done by starting the model at the present time, but using the terminating values (because of the self-modification) of the control parameters.

Predictions made for over 100 Monte Carlo runs allow us to map probabilities. We decided to divide the probabilities into three categories. Stopping the predictions at different future dates created multiple frame displays (figures 4 and 5). The images of the future urban extent are most convincing when animated with the historical data. In a first attempt to animate the San Francisco data, it was difficult to tell where the observed data ended and the simulations began.

Finally, the results of the multiscale calibration were very encouraging. Table 1 shows the results of the 'best' combinations for the Washington/Baltimore data set arrived at by successive halving of the resolution. As can be seen, the only major differences in the degree of success came after the change from 820 to 1640 m. We attribute this change to the severe distortion of the roads layer when roads must be

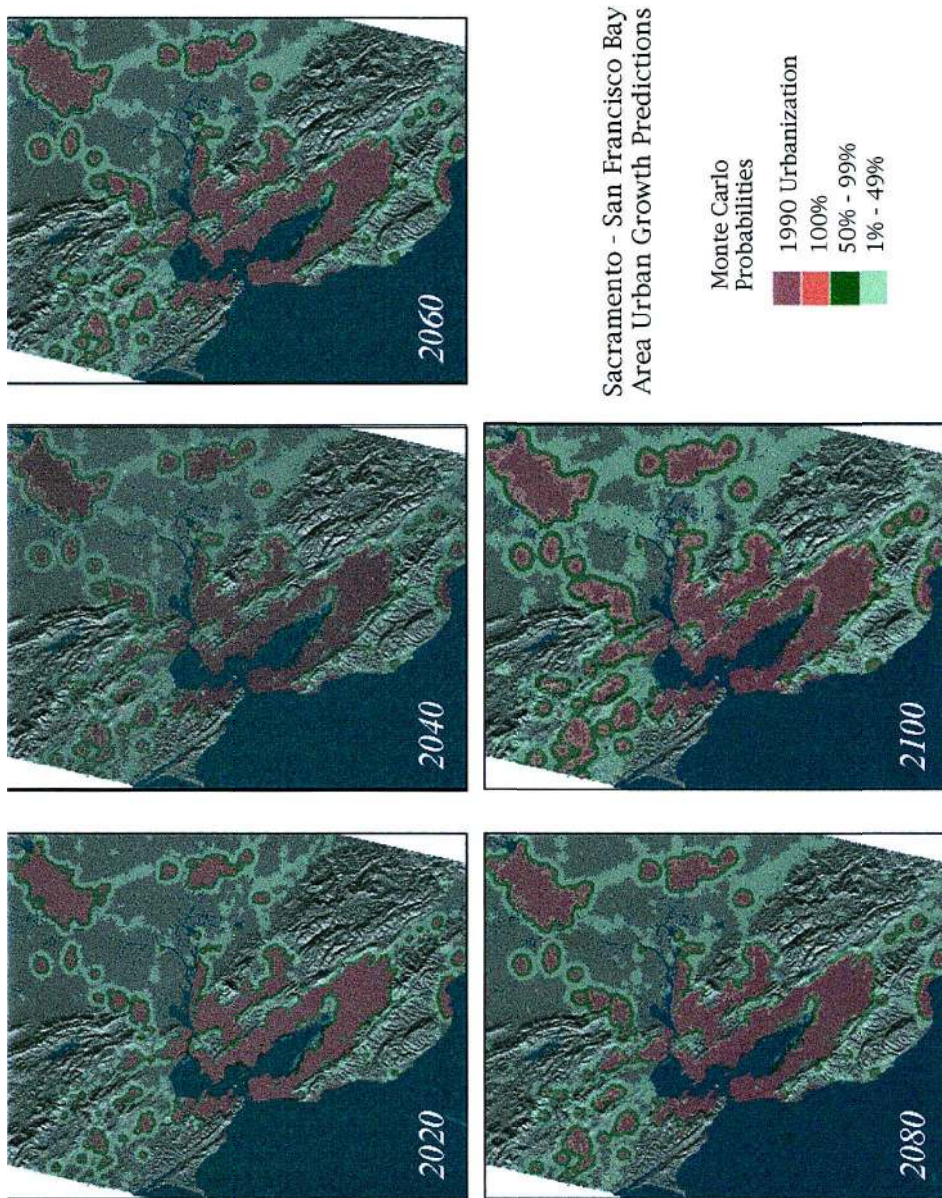



Figure 4. CA model predictions for the San Francisco Bay area.



Washington - Baltimore Urban Growth Predictions




Figure 5. CA model predictions for the Washington/Baltimore area.

over one and a half kilometers wide, and occasionally disconnect. Slope also suffers from the resolution change, with maximum and average slopes diminishing quickly as the scale is made smaller. From table 1, it was clear that the diffusion factor seems most scale sensitive, and that the breed factor seems almost a self-similar fractal, with complete scale invariance. The obvious conclusion from the multi-scale calibration is that a nested or hierarchical approach to calibration would be optimal, first using coarse data to investigate the scaling nature of each parameter in a different city setting, then scaling up once the best data ranges are found. As a result, far fewer combinations need be run at the finest scales, and CPU times can be significantly reduced, with no loss of calibration rigour.

Table 1. Best overall calibrations by resolution: Washington/Baltimore.

Resolution	210 m (486 x 720)	420 m (243 x 360)	840 m (121 x 180)	1680 m (60 x 100)
Composite score 1	0.900014	0.950593	0.993430	0.841441
Composite score 2	3.223740	3.237434	3.226773	3.213986
<i>r</i> -squared urban	0.971070	0.969425	0.960926	0.967818
<i>r</i> -squared edges	0.965636	0.972149	0.972061	0.985269
<i>r</i> -squared clusters	0.978280	0.969293	0.949094	0.973370
Shape match	0.308754	0.326567	0.344692	0.287529
Diffusion	75	10	50	75
Breed	25	2	2	3
Spread	22	26	20	22
Slope	9	1	5	5
Roads	4	25	3	9

7. Predictions

Figures 4 and 5 show model predictions of future urbanization patterns for the two applications areas. Obviously, the short-term predictions are more reliable than the long term. The Monte Carlo probabilities have been reduced to three classes. Shaded pink are those areas that the model predicts as 100% certain of growth. Some of these areas are the existing urbanized areas, but as can be seen from the figures, this area expands over time. Most conservatively, these would be judged as potential growth zones. Because ordinary outward expansion and infilling of existing settlements are the most predictable types of growth, the organic spread of the settlements has the highest degree of confidence. Shaded green are areas with less than 100%, but greater than 50% or even chance of growth. The greater-than-even-odds criterion may be a good compromise for spatial forecasting. Taking in this class of prediction much of the highway-based growth is captured. Finally, a high-risk projection category of greater than zero but less-than-even-odds is plotted in light green. This zone combines more extended and fragmented growth of transportation, but also includes many entirely new urban centers. When a new spreading centre forms in repeated model runs, it could be identified as a 'city waiting to happen', a site so potentially ripe for growth that it is merely a matter of time before urbanization arrives. Examples of such settlements are: Vacaville and Morgan Hill in California, and Chantilly and Manassas in Virginia.

The utility of these predictive layers clearly lies in their being placed in the public realm. The San Francisco Bay area projections have been posted on the World Wide Web for some time. The Washington D.C./Baltimore projections are available at <http://urban.wr.usgs.gov/urban.html>. They will be most useful as data layers in a GIS data format, for integration with existing information.

8. Conclusions

Work on the model so far has established that the approach can produce useful results. While the calibration phase is slow and highly dependent on the size of the input data set and on the quality and quantity of historical data, historical map and image data as the basis for future settlement predictions seems suitable and the multi-scale calibration process could speed the process. Monte Carlo methods, coupled with the variance estimates they generate, allow clear confidence limits to be placed on projected future spatial patterns. As sensitive and potentially controver-

sial as growth predictions are from the political and societal viewpoint, the rigour of the calibration, the repeatability of the experiment and the utility of a prediction-probability layer for further GIS work justify the outcomes of the modelling. Only the real future, as it slowly unfolds, can verify our model. But then, is the purpose of modelling to actually predict, or to help imagine, test, and choose between possible futures? If so, then this model and approach can indeed be useful for scenario planning (Kramer 1996).

While the role of the model as both a consumer and provider of GIS data has been evident, the functional coupling of the model with any particular GIS has been loose, at best. On the other hand, we see such models as valuable enriching sources of GIS data layers, and layers that have real value for planning and GIS application. For example, what location decisions pending today might change in the light of predicted future patterns? Could better decisions on siting waste disposal be made given predictions of urban spread? Would zoning or land-preservation policies change if their future consequences were better known, or even better understood? Could cities test and evaluate alternative growth-control strategies by spatial modeling? Experience with existing models have shown this to be the case.

Future work with the model will involve extending the land transition to a broader set of land use and land-cover changes. We have set the Anderson Level I classes as a first goal for the new work (Anderson *et al.* 1976). In addition, the model needs to be ported to and repeatedly applied to new study areas and at different map scales. We intend to further examine the impact of scale and locale of application of the model on the initial conditions, on the best start parameters, and on the process of self-modification. Our plans call for testing the model at about 1 km resolution for the entire lower 48 United States for the full Anderson Level I classification, and for applications in New York, Chicago, Philadelphia, and Portland. Finally, can the model be used to predict the extent and impact of *gigalopolis*? Can the model be coupled with other models of the global environment to provide multiple assessments of the impact of global change? If these questions are answerable in the affirmative, the cellular model could become a valuable tool for anticipating, and effectively addressing some aspects of our children's, and our childrens children's, urban future.

Acknowledgments

The Human-Induced Land Transformations (HILT) project was supported by the United States Geological Survey (USGS) under a NASA Joint Research Initiative (JRI NCC2-5091). The Gigalopolis project is supported by the USGS National Mapping Division's EROS Data Center under grant 1434-CR-96-SA-01235 (Department of the Interior). Support is gratefully acknowledged.

William Acevedo, Michael Figueroa and Stacy Hoppen provided programming and data analysis support. Data were compiled and provided by USGS under the HILT and Temporal Urban Mapping projects (see, <http://edcwww2.cr.usgs.gov/umap/umap.html>). Jeanette Candau and John Lin prepared the graphics. Additional support for data for Washington/Baltimore came from the Washington/Baltimore Regional Collaboratory courtesy of Dr Tim Foresman of the University of Maryland-Baltimore County (see, <http://research.umbc.edu/~tbenja.bwhp/main.html>). Dr Richard J. Pike and four anonymous reviewers gave valuable comments on the manuscript.

References

- ACEVEDO, W. and MASUOKA, P., 1997, Time-series animation techniques for visualizing urban growth. *Computers & Geosciences*, **23**, 423–435.
- ANDERSON, J. R., HARDY, E. E., ROACH, J. T. and WITMER, R. E., 1976, *A Land Use and Land Cover Classification System for use with Remote Sensor Data*, U.S. Geological Survey Professional Paper 964 (Washington, DC: US Government Printing Office).
- ANSELIN, L., DODSON, R. F. and HUDAK, S., 1993, Linking GIS and spatial data analysis in practice. *Geographical Systems*, **1**, 3–23.
- BATTY, M. and LONGLEY, P., 1994, *Fractal Cities* (London: Academic Press).
- BATTY, M. and XIE, Y., 1994, Modelling inside GIS: Part 2. Selecting and calibrating urban models using Arc/Info. *International Journal of Geographical Information Systems*, **8**, 429–450.
- CHRISMAN, N., 1997, *Exploring Geographic Information Systems* (New York: Wiley).
- CLARK, S., STARR, J., FORESMAN, T. W., PRINCE, W. and ACEVEDO, W., 1996, Development of the temporal transportation database for the analysis of urban development in the Baltimore-Washington Region. In *Proceedings, ASPRS/ACSM Annual Convention and Exhibition, Baltimore, MD, 22–24 April*.
- CLARKE, K. C., 1997, *Getting Started With Geographic Information Systems* (Upper Saddle River, NJ: Prentice Hall).
- CLARKE, K. C., GAYDOS, L. J. and HOPPEN, S., 1997, A self-modifying cellular automaton model of historical urbanization in the San Francisco Bay area. *Environment and Planning B*, **24**, 247–61.
- CLARKE, K. C., HOPPEN, S. and GAYDOS, L. J., 1996, Methods and techniques for rigorous calibration of a cellular automaton model of urban growth. In *Proceedings, Third International Conference/Workshop on Integrating GIS and Environmental Modeling CD-ROM, Santa Fe, NM, 21–26 January, 1996* (Santa Barbara, CA: National Center for Geographic Information and Analysis). http://www.ncgia.ucsb.edu/conf/SANTA_FE_CD-ROM/main.html.
- CLARKE, K. C., RIGGAN, P. and BRASS, J. A., 1995, A cellular automaton model of wildfire propagation and extinction. *Photogrammetric Engineering and Remote Sensing*, **60**, 1355–1367.
- COUCLELIS, H., 1997, From cellular automata to urban models: New principles for model development and implementation. *Environment and Planning B-Planning & Design*, **24**, 165–174.
- CRAWFORD-TILLEY, J. S., ACEVEDO, W., FORESMAN, T. and PRINCE, W., 1996, Developing a temporal database of urban development for the Baltimore/Washington region. In *Proceedings, ACSM/ASPRS Annual Meeting, Washington DC, 3*, 101–110. http://edcwww2.cr.usgs.gov/umap/pubs/asprs_jt.html.
- GIBSON, W., 1984, *Neuromancer* (New York: Ace Books).
- GOODCHILD, M. F., 1992, Geographical information science. *International Journal of Geographical Information Systems*, **6**, 31–46.
- GOODCHILD, M. F., STEYAERT, L. T. and PARKS, B. O., editors, 1996, *GIS and Environmental Modeling* (Fort Collins: GIS World).
- GOTTMANN, J., 1961, *Megalopolis: the urbanized northeastern seaboard of the United States* (New York: Twentieth Century Fund).
- HAGERSTRAND, T., 1965, A Monte-Carlo approach to diffusion. *European Journal of Sociology*, **VI**, 43–67.
- KIRTLAND, D., GAYDOS, L. J., CLARKE, K. C., DECOLA, L., ACEVEDO, W. and BELL, C., 1994, An analysis of human-induced land transformations in the San Francisco Bay/Sacramento Area. *World Resources Review*, **6**, 206–217.
- KRAMER, J., 1996, Integration of a GIS with a Local Scale Self-Modifying Cellular Automaton Urban Growth Model in Southeastern Orange County, New York. MA Thesis. Hunter College CUNY. <http://everest.hunter.cuny.edu/~jkramer/paper.html>.
- MAKSE, H. A., HAVLIN, S. and STANLEY, H. E., 1995, Modelling urban growth patterns. *Nature*, **377**, 608–612.
- PARK, S. and WAGNER, D. F., 1997, Incorporating Cellular Automata simulators as analytical engines in GIS. *Transactions in GIS*, **2**, 213–231.
- TAKEYAMA, M. and COUCLELIS, H., 1997, Map dynamics: Integrating cellular automata and

- GIS through geocalgebra. *International Journal of Geographical Information Science*, **11**, 73–91.
- TOBLER, W. R., 1979, Cellular geography. In *Philosophy in Geography*, edited by S. Gale and G. Olssen, (Dordrecht: Reidel), pp. 379–386.
- VITOUSEK, P. M., 1994, Beyond global warming: Ecology and global change. *Ecology*, **75**, 1861–1876.
- WAGNER, D., 1997, Cellular automata and geographic information systems. *Environment and Planning B*, **24**, 219–234.
- WALDROP, M. M., 1992, *Complexity: The Emerging Science at the Edge of Order and Chaos* (New York: Simon and Schuster).
- WHITE, R. and ENGELN, G., 1992, Cellular automata and fractal urban form: a cellular modelling approach to the evolution of urban land use patterns, Working Paper no. 9264, Research Institute for Knowledge Systems (RIKS), Maastricht, The Netherlands.
- WILSON, J., editor, 1995, Special issue on the integration of GIS and environmental modelling. *International Journal of Geographical Information Systems*, **9**.
- WOLFRAM, S., 1984, Cellular automata as models of complexity. *Nature*, **311**, 419–424.
- WOLFRAM, S., 1986, *Theory and Applications of Cellular Automata* (Singapore: World Scientific).