



Contents lists available at ScienceDirect

Computers, Environment and Urban Systems

journal homepage: www.elsevier.com/locate/compenvurbsys

Designing and implementing a regional urban modeling system using the SLEUTH cellular urban model

Claire A. Jantz ^{a,*}, Scott J. Goetz ^{b,1}, David Donato ^{c,2}, Peter Claggett ^{d,3}

^a Geography–Earth Science Department, Shippensburg University, 1871 Old Main Dr., Shippensburg, PA 17257, USA

^b The Woods Hole Research Center, 149 Woods Hole Road, Falmouth, MA 02540, USA

^c US Geological Survey, 12201 Sunrise Valley Drive [MS 521], Reston, VA 20192, USA

^d US Geological Survey, Chesapeake Bay Program Office, 410 Severn Ave., Suite 109, Annapolis, MD 21403, USA

ARTICLE INFO

Article history:

Received 21 February 2009

Received in revised form 19 August 2009

Accepted 19 August 2009

Available online xxx

Keywords:

Urban modeling
Urban simulation
Cellular automata
SLEUTH
Chesapeake Bay

ABSTRACT

This paper presents a fine-scale (30 meter resolution) regional land cover modeling system, based on the SLEUTH cellular automata model, that was developed for a 257000 km² area comprising the Chesapeake Bay drainage basin in the eastern United States. As part of this effort, we developed a new version of the SLEUTH model (SLEUTH-3r), which introduces new functionality and fit metrics that substantially increase the performance and applicability of the model. In addition, we developed methods that expand the capability of SLEUTH to incorporate economic, cultural and policy information, opening up new avenues for the integration of SLEUTH with other land-change models. SLEUTH-3r is also more computationally efficient (by a factor of 5) and uses less memory (reduced 65%) than the original software. With the new version of SLEUTH, we were able to achieve high accuracies at both the aggregate level of 15 sub-regional modeling units and at finer scales. We present forecasts to 2030 of urban development under a current trends scenario across the entire Chesapeake Bay drainage basin, and three alternative scenarios for a sub-region within the Chesapeake Bay watershed to illustrate the new ability of SLEUTH-3r to generate forecasts across a broad range of conditions.

© 2009 Elsevier Ltd. All rights reserved.

1. Introduction

The objective of this paper is to describe a regional urban land cover modeling system that was developed for the Chesapeake Bay watershed, which is located in the eastern United States (Fig. 1). We developed a fine-scale (30 meter × 30 meter or 0.09 hectare cell size) regional modeling system, based on the SLEUTH urban land-cover change model (Clarke, Hoppen, & Gaydos, 1997; US Geological Survey, 2007) and applied it to forecast growth up to the year 2030 for the Chesapeake Bay watershed (CBW) and adjacent counties, an area covering 257,000 km².

SLEUTH is one of a class of models known as cellular automata (CA), where the land surface is conceptually divided into cells using a regular grid. SLEUTH then associates with each cell an automaton, an entity that independently executes its own state-transition rules, taking into account the states of the automata associated with nearby cells. Given its success with regional scale urban sim-

ulation, its ability to incorporate different levels of protection for different areas, the relative ease of computation and implementation, and the fact that it is public domain software, we adopted the SLEUTH model (Clarke et al., 1997; Clarke & Gaydos, 1998) to form the basis for this work. SLEUTH incorporates spatial data through a link with geographic information systems (GIS) and, like many recently developed CAs (e.g. Van Vliet, White, & Dragicevic, 2009), relaxes many of the assumptions of classic CA theory, such as homogeneity of space, uniformity of neighborhood interactions, and universal transition functions, to more realistically simulate real urban systems. Because they are interactive, modified CA models like SLEUTH are attractive in applied settings as planning tools (Batty, 1997). Potential outcomes can be visualized and quantified, the models can be closely linked with GIS, and raster based spatial data derived from remote sensing platforms can be easily incorporated into the model.

The utility of CA models for simulating complex systems, including urban systems, has been well documented (Couclelis, 1997; O'Sullivan & Torrens, 2000; Silva & Clarke, 2005; Torrens, 2006; Torrens & O'Sullivan, 2001; Van Vliet et al., 2009). For regional scale modeling, CA models have proven to be effective platforms for simulating dynamic spatial interactions among biophysical and socio-economic variables associated with land-cover change (White & Engelen, 1997). For example, Li and Liu

* Corresponding author. Tel.: +1 717 477 1399; fax: +1 717 477 4029.

E-mail addresses: cjant@ship.edu (C.A. Jantz), sgoetz@whrc.org (S.J. Goetz), didonato@usgs.gov (D. Donato), Pclagget@chesapeakebay.net (P. Claggett).

¹ Tel.: +1 508 540 9900.

² Tel.: +1 703 648 5772.

³ Tel.: +1 410 267 5771.




Fig. 1. Study area. The Chesapeake Bay drainage basin is outlined in black. Our study area, shown in gray, includes all the counties that are contained within or that intersect the watershed boundary. The small watersheds colored white in southeast Pennsylvania represent a case study area that will be presented in this paper.

(2006) develop a modeling approach that relaxes traditional CA transition rules with case-based reasoning and explicitly accounts for the influence of proximity and distance between urban clusters. They have used this CA modeling system to accurately simulate fine-scale (30 m × 30 m) urbanization patterns and interactions between hierarchically organized urban centers in the Pearl River Delta in southeastern China, an area of over 41,000 km². Most notably, Soares-Filho et al. (2006) use the SimAmazonia CA modeling system to integrate factors driving deforestation in the Amazon basin, including market forces, road construction, and government regulations. SimAmazonia was applied over a very large region, more than 8 million km² at a resolution of 1 km × 1 km cells.

We addressed two main challenges in this work. First, because of the size of the watershed and the fine grain of the analysis,

application of the model posed significant logistical and computational challenges; this application required more memory than any previously published analysis of which we are aware. Second, urbanization patterns and patterns of urban land-cover change are extremely heterogeneous across the watershed (Jantz, Goetz, & Jantz, 2005). In overcoming these challenges, this application represents a significant contribution to the software infrastructure for simulation of urban growth and the development of decision support tools for regional ecosystem management. This work was undertaken in close partnership with the Chesapeake Bay Program, providing a direct link between the science of land-cover change modeling and applications for ecosystem management.

The US Environmental Protection Agency has listed the Chesapeake Bay as impaired due mainly to non-point source loads of

nutrients and sediment. The Chesapeake Bay Program (CBP), a federal and state agency partnership established to restore the health of the Chesapeake Bay, has agreed to over one hundred different restoration objectives in the areas of living resources, habitat, water quality, stewardship and sound land use (Chesapeake Bay Program, 2000). Future land use forecasts will inform many of these objectives. Knowing the probability of land conversion from agriculture, wetland or forest (resource lands) to residential, commercial, or industrial use (built) will guide the development of practical alternatives and contingency plans related to Bay trends and indicators (Jantz & Goetz, 2007).

The objective was to create an adaptive modeling system capable of producing dynamic and fine-scale forecasts of urban land-change through the year 2030 within the CBW. In order to achieve this objective, these four problems had to be solved:

1. How to modify and adapt the SLEUTH model. Several modifications were made to the SLEUTH model and calibration methodology to address scale sensitivity and the inability of SLEUTH to consider factors that attract and resist development. Also, SLEUTH's performance was enhanced by a set of code modifications that substantially reduced the model's memory requirements and increased processing speed.
2. How to subdivide the Chesapeake Bay watershed. Because of the heterogeneity of urban patterns and urban land-cover change patterns we divided the Chesapeake Bay into 15 sub-regions, ranging from roughly 7,000 km² to 23,000 km² in size, within which urbanization patterns are relatively homogeneous. These subdivisions comprised the spatial framework of our regional modeling system.
3. How best to calibrate our revised version of the SLEUTH model. We calibrated SLEUTH separately for each of the 15 sub-regions.
4. How to select alternative futures for the CBW. A Bay-wide forecast of future urbanization in 2030 under a "current trends" scenario was completed. We also developed two additional policy and growth scenarios to assess the utility of this modeling approach. We illustrate the use of these alternative scenarios using results for a generally representative sub-region, southeast Pennsylvania (Fig. 1).

2. Methods

2.1. Overview of the SLEUTH model

SLEUTH simulates urban dynamics through the application of four growth rules: spontaneous new growth, which simulates the random urbanization of land; new spreading center growth, or the establishment of new urban centers; edge growth; and road influenced growth. Each type of growth is controlled by an area-wide coefficient (diffusion, breed, spread, road growth) that can range in value from 0 to 100, reflecting the relative contribution of a particular growth type to urban dynamics within a study area. The resistance of development to slope is also controlled by a calibrated parameter, the slope coefficient, which ranges from 0 to 100 (0 indicating low slope resistance, 100 indicating high slope resistance). The user can specify additional resistance rules in an excluded layer, which indicates areas that are partially or completely excluded from development.

Implementation of the model occurs in two general phases: calibration, where historic growth patterns are simulated; and prediction, where historic patterns of growth are projected into the future. For calibration, the original SLEUTH model requires inputs of historic urban extent for at least four time periods, a historic

transportation network for at least two time periods, slope, and an excluded layer.

2.2. Modifications to the SLEUTH model

In our previous work with the SLEUTH model, we identified several limitations. First, when fine resolution data are used, SLEUTH is not always able to generate an appropriate level of dispersed growth because of SLEUTH's bias towards edge growth (Jantz & Goetz, 2005).

Second, most of the fit statistics that have commonly been used to calibrate the model are least squares regression scores (r^2) measuring the relationship between a particular simulation of urbanization and actual (historic) observed urbanization. Thus the historic input data sets used in calibration must cover at least four points in time: one to initialize the model and three additional control points to calculate the regression equation. In addition, use of the r^2 statistic alone can result in an under- or over-fitting of the model. Without additional information, such as the y-intercept of the linear regression equation, a user may identify a simulation that appears to perform well but is actually over- or under-estimating growth rates or patterns.

Third, SLEUTH utilizes computer memory inefficiently. For larger data sets, Unix or Linux based parallel computing is typically used to calibrate the model, but the Chesapeake Bay data set exceeded the memory capacity of our available computing resources (memory requirements ranged from 1.4 GB to more than 5 GB), even when divided into sub-regions.

Finally, SLEUTH usually only incorporates factors that constrain development (Jantz, Goetz, & Shelley, 2004). Providing an ability to identify areas where growth is *more likely* to occur will increase the utility of the model, both in terms of improving SLEUTH's ability to simulate historic patterns and in developing scenarios of future development.

The first three points presented above were addressed through direct modification of SLEUTH's source code (written in the C programming language), resulting in a new version of SLEUTH, SLEUTH 3.0beta_p01 Version R (referred to here as SLEUTH-3r). The fourth point (attracting growth) was addressed methodologically during calibration, as discussed below in Section 2.3.

The source code changes discussed in this section thus consist of three primary modifications to: (i) address scale sensitivity, (ii) calculate new fit metrics, and (iii) decrease SLEUTH's memory requirements and optimize processing speed. A general discussion of these changes is presented here. Technical descriptions are available with the model's source code, which can be downloaded from the USGS Eastern Geographic Science Center's (EGSC) High-Performance Computing Cluster (HPCC) website (<http://egscbeowulf.er.usgs.gov/geninfo/downloads/>). We emphasize that SLEUTH-3r represents added functionality to SLEUTH; the original functions of SLEUTH are completely retained, as are the original theoretical underpinnings.

2.2.1. Modifications to address scale sensitivity

SLEUTH's inability to capture dispersed settlements patterns and its tendency to allow edge growth to dominate the system are both related to the number of pixels that the model selects for potential new spontaneous development in any time step. In the original code, the number of spontaneous urbanization attempts (the dispersion value) depends on the calibrated value for the diffusion coefficient, a constant multiplier, and the number of pixels in the image diagonal, a convention embedded in the original source code (US Geological Survey, 2007):

$$D_V = D_C \times D_M \times \sqrt{R^2 + C^2}$$

where D_V is the dispersion value, D_C is the diffusion coefficient, D_M is the diffusion coefficient multiplier (a constant equal to 0.005 in the original version of SLEUTH), R is the number of rows and C is the number of columns.

In SLEUTH-3r, D_M is no longer a constant, allowing the user to change this multiplier value interactively. When the multiplier is increased or decreased, the number of urbanization attempts for diffusion growth changes accordingly. D_M must be set prior to beginning calibration. To discover an appropriate multiplier value, SLEUTH-3r's growth coefficients were set to produce the maximum level of spontaneous new growth (i.e. diffusion was set to 100 and all other growth coefficients set to 0) (Jantz & Goetz, 2005). Then, several simulations are performed with SLEUTH-3r in calibration mode to test different values for D_M , simulating growth over the length of the historic urban time series. When D_M is set such that SLEUTH-3r is able to capture, or even over-estimate, the number of urban clusters (as measured by the *cluster fractional difference* metric, a new pattern metric discussed below), normal calibration procedures can be initiated (see Section 2.3) to identify the best values for SLEUTH-3r's growth coefficients.

2.2.2. New calibration statistics

In addition to the ability to interactively set the diffusion coefficient multiplier, SLEUTH-3r now also creates new tabular files that include difference and ratio metrics that directly compare the modeled variable (e.g. number of urban clusters) with the observed variable for all control dates. Specifically, SLEUTH-3r calculates (i) the algebraic difference between the observed value and modeled value, (ii) the ratio of the modeled value to the observed value, and (iii) the fractional change in the modeled value relative to the observed value. It does this for most of the original fit statistics, for each run, and for each control year.

When at least four control points are available, these new fit metrics can be used in conjunction with the r^2 values to enhance the calibration procedure. When fewer than three control points are available, the new metrics can be the principal means for calibrating SLEUTH-3r. Table 1 presents a list of the new fit metrics available in SLEUTH-3r.

2.2.3. Decreasing memory requirements and improving processing speed

The final set of modifications made to SLEUTH's source code addressed the model's memory requirements and computational speed. SLEUTH requires space in RAM for numerous internal cell arrays, each with the same dimensions as the modeling unit; our applications required about 18 of these internal cell arrays, so the largest modeling unit in our study area would require space in RAM for more than 1.4 billion cells. Because available versions of SLEUTH required 4 bytes of RAM for each cell, our largest sub-

region would require more than 5.6 Gigabytes of RAM, which exceeded the 2.0 Gigabyte maximum program size under the 32-bit computer operating systems used by the vast majority of users.

Calibrating with the existing versions of SLEUTH was a computationally intensive process for which the required computer processing time was roughly proportional to the size of the modeling unit. Our relatively large sub-regional modeling units would thus require extensive and lengthy calibration computations.

A review of the SLEUTH source code revealed that only one byte of RAM per cell was actually required in any of the internal cell arrays (grids) because the largest number required to be stored for any one cell was 255 or less. Since all integers between 0 and 255 can be represented by a single 8-bit byte of computer storage (using the C-language "unsigned char" data type), in SLEUTH-3r we could use a single byte per cell in the SLEUTH internal arrays instead of the four-byte value which had been allocated in standard SLEUTH. We incorporated this change into SLEUTH-3r and successfully tested it to insure that the change did not introduce any spurious artifacts. With this change in place, we were able to use SLEUTH-3r with our relatively large modeling units.

Additional improvements in SLEUTH's processing speed were also desirable, and processing statistics produced by SLEUTH showed that the single most time-consuming activity in our growth simulations was the road growth algorithm. The original road-search algorithm proceeds stepwise from the location of a new-growth cell within the internal roads array, starting with the square of eight cells immediately surrounding the new-growth cell. The algorithm begins with the northwest cell (topmost and leftmost) and proceeds counter-clockwise around the square, checking each cell to see if it is a road cell (Fig. 2). If no road cell is found, the algorithm steps out to the next square of cells and repeats this process. The road-search ends when the first road cell is found. This is inefficient because it requires that each cell within a potentially large square area be checked every time a road-search is conducted. Furthermore, since the algorithm does not remember from one search to another where the roads are located it performs an inordinate amount of repetitive processing. The algorithm is also biased because it systematically selects road cells to the northwest even if there are equally close road points to the south, east, or northeast.

The key to speeding up the road-search was to prevent the algorithm from "forgetting" where the roads are. We created a new, compressed data structure that contains only the coordinates of the points in the road grid that are road cells. This structure is much smaller than the source grid because only a small proportion of any area will be covered by roads. Based on this data structure, we created a new road-search algorithm which sequentially checks the rows of cells above and below the new-growth cell until the closest road cell is found or it is determined that there is no

Table 1
New fit metrics available in SLEUTH-3r. For each of the metrics described below, SLEUTH-3r writes the following three quantities to a ratio. Log file: (i) the algebraic difference between the observed value and modeled value (diff), (ii) the ratio of the modeled value to the observed value (ratio), and (iii) the fractional change in the modeled value relative to the observed value (fract). Measurements derived from the modeled data are averaged over the set of Monte Carlo trials. It does this for each run, and for each control year.

| Fit statistic | Definition |
|-------------------------|---|
| Pixels (pix) | Modeled urban pixels compared to actual urban pixels for each control year. Referred to as "population" and as "area" in SLEUTH's output files |
| Edges (edges) | Modeled urban edge pixels compared to actual urban edge pixels for each control year |
| Clusters (clusters) | Modeled number of urban clusters compared to actual urban clusters for each control year. Urban clusters are areas of contiguous urban land. In cell space, clusters can consist of a single pixel or multiple, contiguous urban pixels. Contiguity is determined using the eight-neighbor rule |
| Cluster size (mn_cl_sz) | Modeled average cluster size compared to actual average urban cluster size for each control year. This is not an area-weighted mean |
| Slope (avg_slope) | The average slope for modeled urban pixels compared to actual average slope for urban pixels for each control year |
| % Urban (pct_urba) | The percent of available pixels urbanized during simulation compared to the actual urbanized pixels for each control year |
| X-mean (xmean) | Average x-axis values for modeled urban pixels compared to actual average x-axis values for each control year |
| Y-mean (ymean) | Average y-axis values for modeled urban pixels compared to actual average y-axis values for each control year |
| Radius (radius) | Average radius of the circle that encloses the simulated urban pixels compared to the actual urban pixels for each control year |

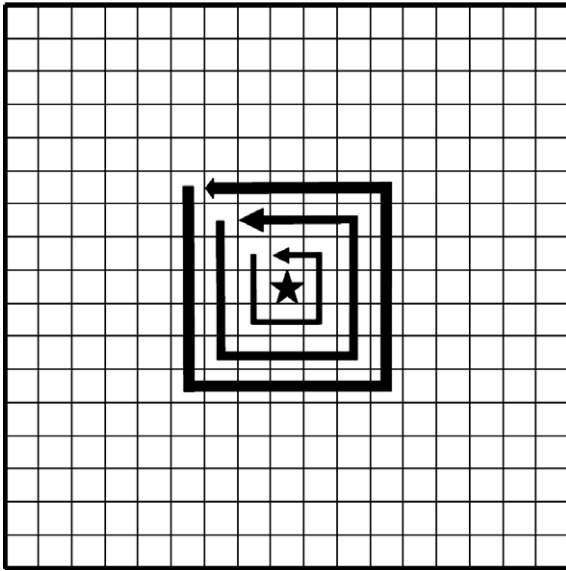


Fig. 2. Illustration of the search algorithm in the original SLEUTH Program.

close-by road. Because this algorithm uses the new, compact data structure for roads, it does not perform any repetitive, cell-by-cell checking and so it is much faster than the old algorithm. The new algorithm also corrects the biases of the old one. When the new algorithm finds a closest road cell, there is no closer cell in the sense of a Pythagorean distance metric, although on occasion there may be another equally close cell.

Having found that SLEUTH processes many of its internal grids in a cell-by-cell sequence, looking for non-zero cells, in our final code modifications we created new data structures for several grids which list coordinates for just the non-zero cells and we modified the procedures which process these grids so they would skip the now unnecessary checking of zero-valued cells and process only the non-zero cells.

2.3. Sub-dividing the Chesapeake Bay watershed

As noted in Section 1, one of the main objectives of this project was to develop an urban modeling system that could be applied

across the Chesapeake Bay watershed while maintaining the high spatial resolution of the available urban land cover maps. The urban land cover data consist of maps of impervious surface cover, derived from Landsat TM and ETM+ imagery, which captured urbanization patterns between 1990 and 2000 (Goetz et al., 2004; Jantz et al., 2005). Urbanization in the region, as characterized in the Landsat maps, was primarily associated with existing urban centers, such as Washington, DC, Baltimore, MD, and Norfolk, VA. In many exurban counties, however, rates of change exceeded those in urban areas (Jantz et al., 2005). In addition, urbanization patterns in urban and suburban counties tend to be characterized by clustered, high-density development. In exurban counties development patterns tend to be more dispersed.

Because SLEUTH's growth coefficients are applied globally within a study area, the heterogeneous urbanization patterns observed across the Chesapeake Bay watershed, as well its size, required that the study area be subdivided; an analogous problem was faced by Soares-Filho et al. (2006) in the Amazon region. We used *k*-means cluster analysis, a robust method for identifying groupings within a data set where within-group variability is minimized and between-group variability is maximized (Aldenderfer & Blashfield, 1984), to characterize rural, suburban and urban landscapes at the county scale. These broad groups were then subdivided further into smaller intermediate sub-regions, based on a combination of political boundaries, rural–urban commuting patterns, and physiographic provinces.

The cluster analysis was performed using the 208 counties that intersect the Chesapeake Bay watershed and each county was categorized as being rural, suburban or urban based on several variables, which are briefly discussed here and summarized in Table 2. Fragstats (version 3.3, build 4) (McGarigal & Marks, 1995), a pattern analysis software package, was used to calculate county-level pattern metrics from the 2000 urban land cover map: percent area developed, area-weighted mean urban cluster (patch) size, urban edge pixel density and urban cluster density. Population density was derived from US Census data (US Bureau of the Census, 2000). The following measures of change between 1990 and 2000 were also included: change in percent area developed, change in urban cluster density, change in area-weighted mean urban cluster size, change in urban edge pixel density and the percent change in population.

We also used the level-III EPA ecoregions (US Environmental Protection Agency, 2003) that comprise the Chesapeake Bay

Table 2

Input variables for multivariate *k*-means clustering. Variables were calculated for each county.

| Variable | Data source |
|---|---|
| Percent developed area in 2000 | Derived from 2000 impervious surface map (Goetz et al., 2004; Jantz et al., 2005) using FRAGSTATS (McGarigal & Marks, 1995) |
| Area-weighted mean urban cluster size in 2000 | Derived from 2000 impervious surface map (Goetz et al., 2004; Jantz et al., 2005) using FRAGSTATS (McGarigal & Marks, 1995) |
| Urban cluster density in 2000 | Derived from 2000 impervious surface map (Goetz et al., 2004; Jantz et al., 2005) using FRAGSTATS (McGarigal & Marks, 1995) |
| Urban edge pixel density in 2000 | Derived from 2000 impervious surface map (Goetz et al., 2004; Jantz et al., 2005) using FRAGSTATS (McGarigal & Marks, 1995) |
| Population density in 2000 | US Bureau of the Census (2000) |
| Dominant rural–urban commuting classification in 2000 | USDA Economic Research Service (2000) |
| 1990–2000 Change in percent developed area | Derived from 1990 and 2000 impervious surface maps (Goetz et al., 2004; Jantz et al., 2005) using FRAGSTATS (McGarigal & Marks, 1995) |
| 1990–2000 Change in area-weighted mean cluster size | Derived from 1990 and 2000 impervious surface maps (Goetz et al., 2004; Jantz et al., 2005) using FRAGSTATS (McGarigal & Marks, 1995) |
| 1990–2000 Change in urban cluster density | Derived from 1990 and 2000 impervious surface maps (Goetz et al., 2004; Jantz et al., 2005) using FRAGSTATS (McGarigal & Marks, 1995) |
| 1990–2000 Change in urban edge pixel density | Derived from 1990 and 2000 impervious surface maps (Goetz et al., 2004; Jantz et al., 2005) using FRAGSTATS (McGarigal & Marks, 1995) |
| 1990–2000 Change in population density | US Bureau of the Census (1990), US Bureau of the Census (2000) |
| Dominant ecoregion | US Environmental Protection Agency (2003) |

watershed and gave each one of the seven ecoregions a numerical identifier. The counties were then labeled with the ecoregion code that comprised the majority of the county area.

Finally, we incorporated a simplified rural–urban commuting classification based on the United States Department of Agriculture (USDA) Economic Research Service (ERS) rural–urban commuting area (RUCA) codes (USDA Economic Research Service, 2000). Defined at the Census tract-level, RUCA codes are derived from Census measurements of population density, and daily commuting patterns to identify urban areas and the adjacent economically and functionally integrated areas. These tract-level codes were aggregated to the county scale to identify urban, suburban, and rural counties, based on the dominant commuting patterns.

Using these input variables, the *k*-means analysis was then used to identify rural, suburban and urban counties within the CBW. However, these initial subdivisions resulted in regions that were still too large to be modeled as individual units. Fifteen sub-regions, ranging from roughly 7100 km² to 23,000 km², were therefore identified using the initial *k*-means groupings as a basis for splitting the initial subdivisions. Finally, a 10-km buffer was applied to each of the 15 sub-regions, creating an overlap area between adjacent sub-regions to minimize edge effects that might otherwise result from different growth parameters being applied to each sub-region.

2.4. Calibration of the SLEUTH-3r model

The goal of SLEUTH calibration is to find a set of values for the five parameters (discussed in Section 2.1) that can accurately reproduce actual past land-cover change within the study area. Calibration is typically undertaken using what is referred to as a “brute force” methodology. That is, a large number of combinations of parameter values are tested automatically and the user evaluates the results, locating a “best fit” set of parameter values through the use of fit statistics (Table 1). We performed what is referred to as a coarse calibration, where the values for each parameter ranged from 1–100, but only increments of 25 were tested (i.e. 1, 25, 50, 75, and 100). This resulted in 3125 unique parameter combinations. In our previous work, we found that any gains in performance achieved by testing additional parameter values are minimal, particularly given the substantial increase in computing time (Jantz & Goetz, 2005).

The choice of appropriate goodness of fit measures is important, since it determines how SLEUTH will simulate urban patterns and how forecasts of urban growth will be created (Silva & Clarke, 2002). However, there is no consensus regarding which goodness of fit measure or set of measures to use. Clarke et al. (1997) relied primarily on four metrics: *population*, *edges*, *clusters*, and the *Lee and Sallee statistic*. Recent examples show that others have relied on a weighted sum of all the statistical measures (Yang & Lo, 2003), or an unweighted product score of several metrics (Candau, 2002; Silva & Clarke, 2002). Dietzel and Clarke (2007) suggest an optimum fit statistic, a product of seven of SLEUTH's fit statistics that were found to produce robust and unique results. We emphasize, however, the potential difficulty in evaluating the fit of the model using a composite score. For example, Jantz and Goetz (2005) found that the parameter sets producing a high fit score for one statistic were opposed to those producing a high fit for another, making interpretation of the model's behavior problematic when using composite metrics. For the calibration procedure in this work, we therefore focused on two metrics we considered most relevant to the application: the *pixel fractional difference* (PFD) and the *clusters fractional difference* (CFD).

The PFD and CFD metrics are direct comparisons between the number of urban pixels and the number of urban clusters, respectively, in the control maps and the corresponding simulated maps.

Achieving an accurate fit for the PFD metric ensured that the overall amount of development would be matched. The CFD metric is a simple pattern metric that focuses on the frequency of clusters in the urban system. Achieving an accurate fit for this metric indicates that the model is capturing an important aspect of urban form (i.e. clustered vs. dispersed settlement patterns). We selected parameter sets that were able to match both of these fit statistics within $\pm 10\%$.

SLEUTH is stochastic and thus utilizes the Monte Carlo method to generate multiple simulations of growth for each unique parameter set, so the fit statistics that SLEUTH-3r calculates are averaged over the Monte Carlo trials. For calibration, we initially used only seven Monte Carlo trials to economize computational processing time. Based on these initial results, we selected a subset of parameter sets that performed well. Then, each parameter set was tested by running the model in calibrate mode for 25 Monte Carlo trials. Twenty-five Monte Carlo trials were, we found by experimentation, sufficient for quantifying the spatial variability resulting from random processes. We were therefore able to achieve acceptable computational efficiency while maintaining a rigorous calibration procedure.

For calibration, the original version of SLEUTH requires inputs of historic urban extent for at least four time periods, a historic transportation network for at least two time periods, slope, and an excluded layer. Because of the new fit statistics, SLEUTH-3r requires only two inputs of historic urban extent. We were therefore able to take advantage of our existing data set for the Chesapeake Bay watershed for 1990 and 2000, as noted in Section 2.3 and as documented in Goetz et al. (2004) and Jantz et al. (2005).

A USGS 7.5 min digital elevation model was used to create an input layer for slope. The road network used in this study is based on limited access and other major highways, derived from the US Streets data set (Environmental Systems Research Institute, 2003), which reflects the ca. 2000 primary road network. The primary road network was used because it is the transportation network that likely has had the largest influence on regional growth patterns. We assumed no change in the primary road network between 1990 and 2000 due to the lack of data for 1990 of comparable quality to the 2000 road data.

The final input is the excluded layer, which designates lands that are resistant to urban development. For the excluded layer used in calibration, federal, state, and local parks, easements, and water bodies were entirely excluded from development. The excluded layer is typically scaled from 0 (no exclusion) to 100 (completely excluded). However in our calibration of SLEUTH-3r, instead of using zero as the default value to indicate areas theoretically open for development, we used a base value of 50. This allows the user to indicate areas that are more likely to be developed by applying values less than 50 in the excluded layer, effectively creating an exclusion/attraction layer. This exclusion/attraction layer provides added functionality for both calibration and forecasting and, we contend, enables improved overall model performance by allowing the inclusion of growth attractors (e.g. areas of anticipated population growth) as well as constraints.

SLEUTH also has a ‘self-modification’ function (Clarke et al., 1997), which is intended to more realistically simulate different rates of growth over time. When the rate of growth exceeds a specified critical threshold, the growth coefficients are multiplied by a factor greater than one, simulating a development ‘boom’ cycle. Likewise, when the rate of development falls below a specified critical threshold, the growth coefficients are multiplied by a factor less than one, simulating a development ‘bust’ cycle. Without self-modification, SLEUTH will simulate a linear growth rate until the availability of developable land diminishes. Because we used only two actual historic data sets, we did not invoke the self-modification function for calibration. As discussed in the next section, however, we did utilize self-modification when creating forecasts.

In order to provide additional assessments of the accuracy and utility of the model simulations, beyond those calculated by the model during calibration, we performed an extensive accuracy assessment. After the best-fit parameters were identified for each sub-region, the model was initialized in 1990 and run in predict mode to 2000, with 25 Monte Carlo trials. This resulted in a predicted development probability surface for 2000, which was then compared to the observed patterns for 2000. We assessed the performance of the model across multiple extents and scales: counties, Hydrologic Unit Code 11 (HUC 11) watersheds, and an array

of $7290\text{ m} \times 7290\text{ m}$ grid cells (Fig. 3). The cell size for this array was selected to achieve a cell resolution between 5 km^2 and 10 km^2 , and so that the number of rows and columns would match the extent of our study area.

2.5. Forecasts to 2030

When forecasts are created with SLEUTH, the model is initialized with the latest urban extent map, in our case the year 2000, and the growth coefficient values that were derived during

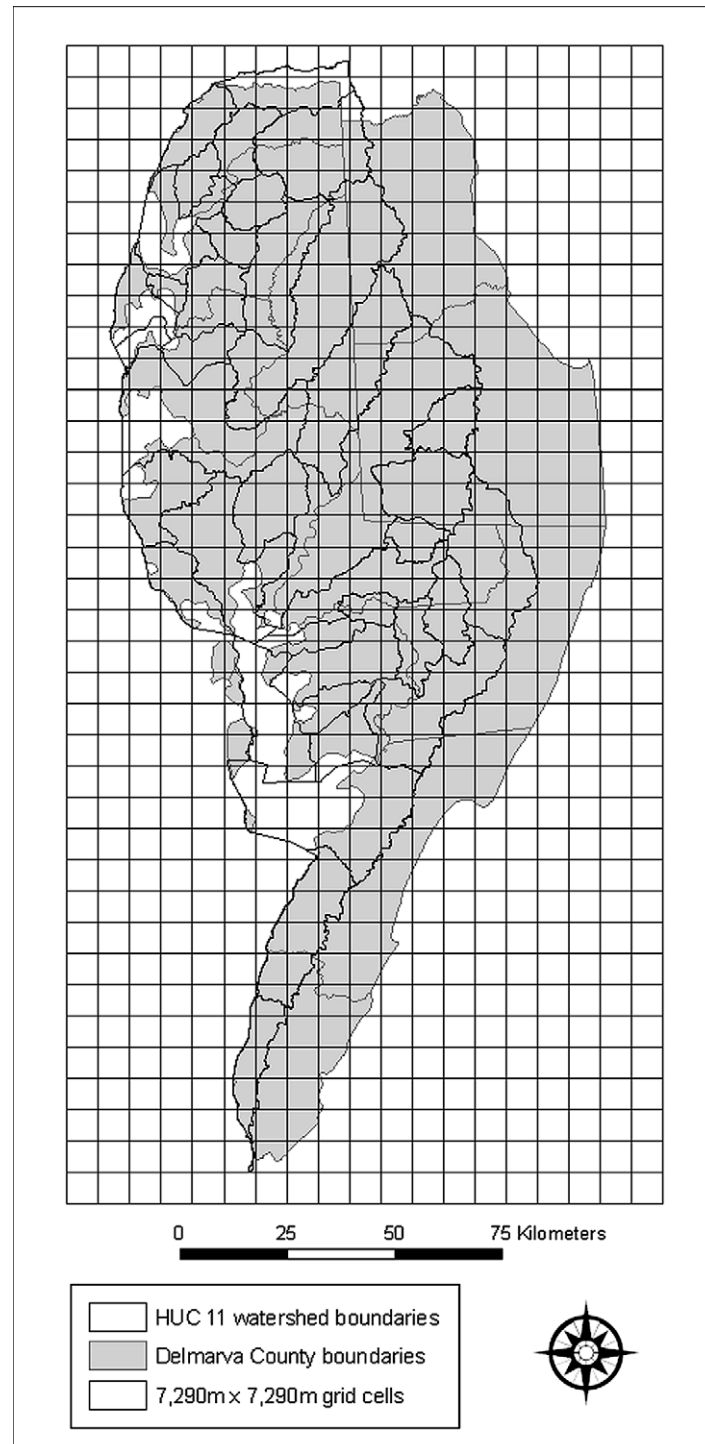


Fig. 3. Areal units for multi-scale calibration accuracy assessment, shown for the Delmarva peninsula.

calibration. The user sets the target year in which to stop the forecast; in our case we chose the year 2030. As in calibration, 25 Monte Carlo trials were performed and each sub-region was modeled separately. For the forecasts presented here, we utilized the same exclusion/attraction layer that was used for calibration,

assuming no change in spatial factors that would influence urban patterns in the future. For most sub-regions, we assumed a linear growth trend and thus did not invoke the model's self-modification functionality. However, for urban sub-regions, such as the Washington, DC–Baltimore, MD region, or urbanizing sub-regions,

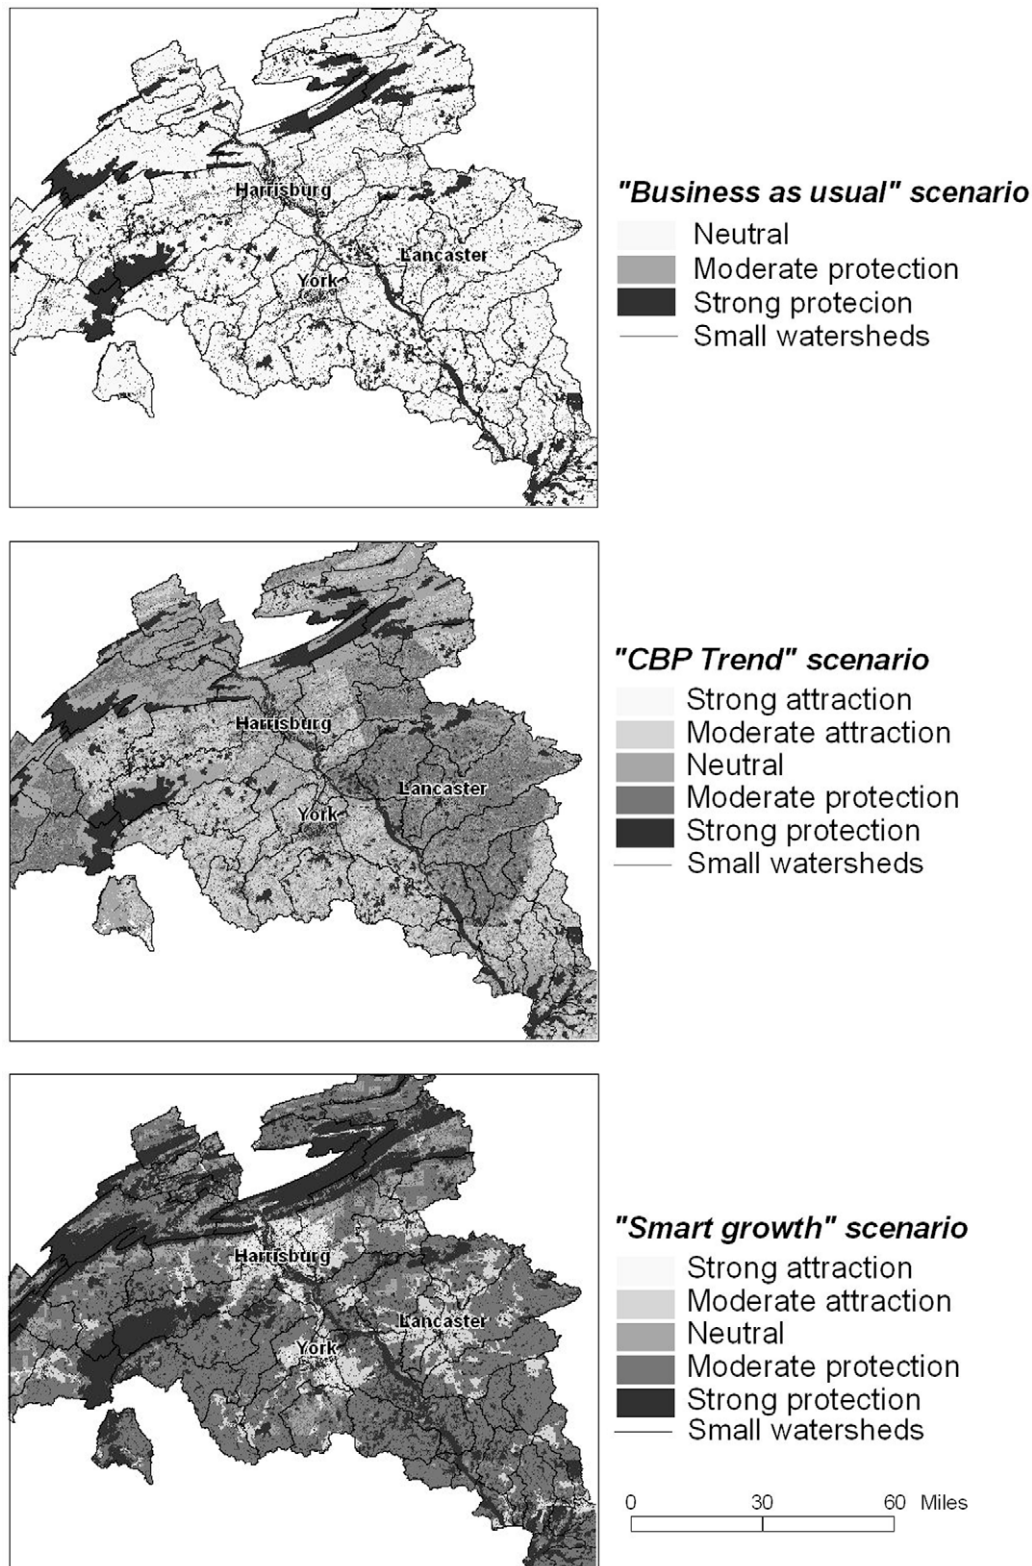


Fig. 4. Forecast scenario maps (exclusion/attraction layers) for southeast Pennsylvania.

such as the Delmarva Peninsula, the maintenance of linear growth rates was thought to be an implausible assumption. As land available for urbanization decreases, actual growth rates will slow as there is a greater economic incentive for growth to occur in denser clusters. Thus for these sub-regions, we allowed the model's "bust" self-modification to operate, using a multiplier of 0.95. When the model enters into the "bust" cycle, the self-modification multiplier is applied to the diffusion, breed, and spread coefficient values before each annual growth cycle begins, effectively lowering those values and slowing growth.

To illustrate the capability of SLEUTH to simulate alternative scenarios, we developed three different scenarios to forecast future development for the southeast Pennsylvania sub-region, an area of expanding population and exurban growth (Fig. 1). This sub-region is a good case study area because it has a large urban center, Harrisburg, PA, several small urban centers, such as Lancaster and York, PA, and a heterogeneous exurban landscape that includes agricultural valleys and forested ridges. In addition, it is a region that has been experiencing rapid growth in recent years due to its proximity to the Washington, DC–Baltimore, MD and Philadelphia, PA metropolitan regions.

For the test sub-region, alternative scenarios were implemented by developing exclusion/attraction layers that reflect different land use policy scenarios (e.g. Jantz et al., 2004). In our case, we developed three scenarios (Fig. 4):

1. A "business as usual" (BAU) scenario that assumed no change in the excluded layer.
2. A trend scenario developed by the Chesapeake Bay Program (CBP trend) that incorporates an agricultural vulnerability model. In addition to using the existing protected lands to designate areas that are off-limits for new development, this scenario identifies county agricultural lands that are either more or less likely to be developed based on the relative difference between the extent of modeled agricultural lands in 2030 and mapped agricultural lands in 2002 at the county scale (R. Burgholzer, pers. comm., based on participation in the CBP Agricultural and Nutrient Reduction Workgroup). Because we calibrated SLEUTH using an exclusion/attraction layer with a base value of 50, county agricultural lands that are less likely to be developed were given a value

between 51 and 80 (expressing the degree of exclusion of urbanization); county agricultural lands more likely to be developed were given a value between 20 and less than 50 (expressing the degree of attraction of urbanization). Agricultural resistance values were not stretched to the minimum or maximum range of possible values because no farmlands could be considered to be powerfully attractive or repulsive to urban growth based solely on farm animal production trends.

3. A "smart growth" (SG) scenario that uses the exclusion/attraction layer to code existing protected lands as completely off-limits for new development, and to apply stronger protection for lands with associated cultural or natural value. Also, areas around existing urban centers were assumed to be more likely to become developed. Resource lands were identified using the Chesapeake Bay Program's Resource Lands Assessment (RLA) data sets (Chesapeake Bay Program, 2005). The RLA consists of five GIS-based models that represent ecological, environmental or cultural elements: ecologically valued forests, economically valued forests, forests valued for water quality protection, prime agricultural soils, and cultural assets. We applied a resistance to development using a GIS overlay model that combined the five RLA maps using a weighting scheme as outlined in Fig. 5. Resistance within "smart growth" areas was eliminated to simulate a higher likelihood of development in these areas. Smart growth areas were centered on existing developed areas and were modeled using US Census designated urbanized area boundaries (Fig. 5).

Three versions of each scenario were run: one where the growth coefficients stayed static over the forecast time period (linear growth), and two where self-modification "bust" multipliers would be applied to cause growth rates to decline over the forecast time period. The first "bust" scenario used a multiplier of 0.90 and the second used a multiplier of 0.80. The critical growth rate was set so that the system would go into a bust cycle beginning with the first forecast year. Because growth rates for southeast Pennsylvania were high throughout the calibration time period, we assumed that a linear growth trend (with a bust multiplier of 1.0) would represent a high growth scenario, the 0.90 multiplier would simulate a moderate growth scenario, and the 0.80 multiplier would

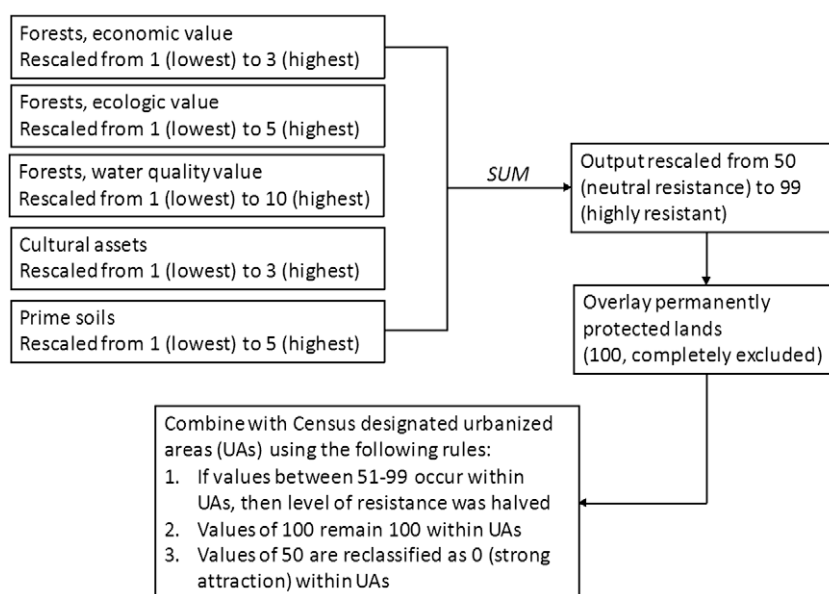


Fig. 5. GIS overlay model for smart growth scenario.

simulate a low growth scenario. Thus, a total of nine forecast scenarios were run, a low, medium and high growth forecast for each of the three policy scenarios.

3. Results

3.1. Modifications to the SLEUTH model

The ability to interactively set a diffusion coefficient multiplier that reflects the unique characteristics of the study area is a key advancement in the SLEUTH-3r model. For all sub-regions we were able to identify a value for the diffusion coefficient multiplier that would over-estimate the number of urban clusters (by roughly 30%) when diffusion growth was maximized (Table 3). This ensured that SLEUTH-3r would be able to simulate an appropriate level of diffusion growth (see Section 3.3 for calibration results).

Our modifications of the SLEUTH code to speed-up processing and to reduce memory requirements proved effective. The net

reduction in required RAM was approximately 65%, though the exact savings depend on the size and content of the input layers. With these reduced memory constraints we were able to make runs for our relatively large modeling units on computers equipped with just 1.5 Gigabytes of RAM. The new road-search algorithm proved to be about 800 times faster than the old algorithm and, because of the prominent role of this algorithm, the overall speed of processing was at least doubled. The processing speed-up from the new road-search algorithm in conjunction with increased speed resulting from more efficient processing of other internal arrays resulted in an overall speed increase of a factor of five. This allowed us to run calibrations for all 15 of our sub-regions on 10 nodes of a Beowulf cluster over the course of about a month; without the increased speed our calibration runs would have taken nearly five months of CPU time.

3.2. Subdivision of the Chesapeake Bay watershed

The initial results of the multivariate *k*-means clustering is shown in Fig. 6A. Urban centers, such as Washington, DC, Baltimore, MD and Richmond, VA, are clearly identified, along with their surrounding suburban or suburbanizing counties. As discussed in Section 2.3, the county groupings identified in this initial analysis were, however, too large and required further subdivision. Using the methodology described in Section 2.3, and with input from the Chesapeake Bay Program, we settled on the final regional subdivisions shown in Fig. 6B.

3.3. Calibration of the SLEUTH-3r model

The calibration results for each sub-region provide both the best-fit parameter set (Table 4) and corresponding fit metrics (Table 5). For all sub-regions we were able to match the overall amount of development within 10% and for all but four sub-regions (New York central, Richmond–Norfolk Virginia, Virginia south-cen-

Table 3
Diffusion coefficient multiplier (D_M) values for each sub-region.

| Sub-region | D_M |
|------------------------------|-------|
| Delmarva | 0.040 |
| New York central | 0.040 |
| New York east | 0.001 |
| New York west | 0.025 |
| Pennsylvania north-central | 0.060 |
| Pennsylvania northeast | 0.015 |
| Pennsylvania northwest | 0.040 |
| Pennsylvania south-central | 0.050 |
| Pennsylvania southeast | 0.130 |
| Virginia central | 0.060 |
| Virginia Richmond–Norfolk | 0.080 |
| Virginia south | 0.030 |
| Virginia south-central | 0.045 |
| Virginia west | 0.040 |
| Washington, DC–Baltimore, MD | 0.120 |

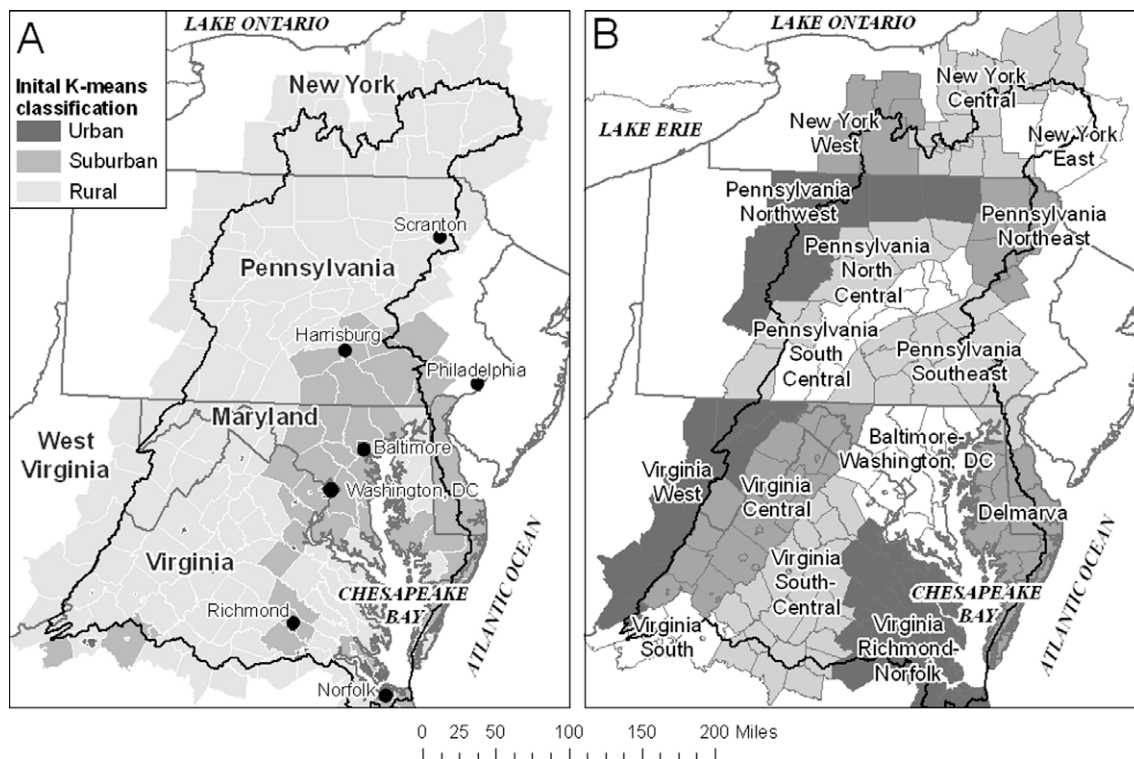


Fig. 6. Initial *k*-means stratification (A) and final stratification (B).

tral and Virginia south) we achieved a match within 5%. Likewise, for matching the number of urban clusters, all sub-regions except for New York east achieved a match within 10% and most were matched within 5%.

SLEUTH calculates the fit metrics globally for each sub-region. We also present results of the model's performance at the county scale, HUC 11 watershed scale, and using the 7290 m × 7290 m lattice. Table 6 shows the results of linear regression analyses that compare the observed and simulated development in 2000 and urban land-cover change between 1990 and 2000 for each areal unit. Fig. 7 illustrates, for the 7290 × 7290 lattice, the spatial patterns of the differences between simulated and observed urban land cover estimates for 2000.

3.4. Forecasts to 2030

The basin-wide forecasts (Fig. 8) indicate a continuation and intensification of development trends that were observed in the 1990–2000 time period (Jantz et al., 2005). We note, for example, the intensification of urbanization in southeast Pennsylvania, between Harrisburg and Philadelphia, and on the Delmarva Penin-

sula. Likewise, exurban development throughout Virginia is also apparent.

The results for the nine 2030 forecasts for southeast Pennsylvania (Fig. 9) indicate that the overall amount of development between the BAU and SG scenarios, regardless of the growth rates, was similar. Fig. 10 focuses on these two scenarios to illustrate the *spatial differences* in growth patterns. The CBP trend scenario consistently resulted in higher levels of growth, likely due to the fact that there are more pixels available for urbanization in this scenario. As an example of an impact assessment, we compared the types of land converted to development by overlaying the forecast maps with the RLA map developed for the smart growth scenario (Fig. 11).

4. Discussion

4.1. Modifications to the SLEUTH model

The added functionality of SLEUTH-3r has greatly enhanced the model's ability to capture urbanization patterns across a wide range of conditions. That we found diffusion coefficient multiplier

Table 4

Parameter sets for each sub-region.

| Sub-region | Diffusion | Breed | Spread | Slope | Road growth |
|------------------------------|-----------|-------|--------|-------|-------------|
| Delmarva | 100 | 75 | 50 | 25 | 50 |
| New York central | 50 | 25 | 25 | 50 | 50 |
| New York east | 50 | 50 | 100 | 50 | 25 |
| New York west | 100 | 50 | 50 | 100 | 25 |
| Pennsylvania north-central | 75 | 25 | 25 | 1 | 75 |
| Pennsylvania northeast | 100 | 50 | 75 | 100 | 50 |
| Pennsylvania northwest | 75 | 50 | 50 | 25 | 1 |
| Pennsylvania south-central | 50 | 50 | 25 | 50 | 50 |
| Pennsylvania southeast | 75 | 75 | 25 | 1 | 50 |
| Virginia central | 100 | 25 | 50 | 25 | 25 |
| Virginia Richmond–Norfolk | 100 | 100 | 25 | 75 | 25 |
| Virginia south | 50 | 50 | 50 | 100 | 50 |
| Virginia south-central | 75 | 25 | 75 | 50 | 50 |
| Virginia west | 75 | 1 | 75 | 50 | 1 |
| Washington, DC–Baltimore, MD | 100 | 50 | 25 | 1 | 50 |

Table 5

Calibration accuracy results for each sub-region. The number of urban pixels and the percent urban, and the number of urban clusters for 1990 and 2000 are given, along with the simulated number of pixels and clusters for 2000. For the pixels and clusters fractional difference metrics, a zero value indicates a perfect match between the simulated and observed data sets. Negative values indicate underestimation; positive values indicate overestimation.

| Sub-region | 1990 Pixels (% urban) | 2000 Pixels (% urban) | 2000 Simulated pixels | Pixels fractional difference | 1990 Clusters | 2000 Clusters | 2000 Simulated clusters | Clusters fractional difference |
|------------------------------|--------------------------|--------------------------|-----------------------------|---------------------------------|------------------|------------------|----------------------------|-----------------------------------|
| Delmarva | 327,589 (0.80) | 798,168 (1.96) | 789,505 | −0.01 | 92,427 | 185,299 | 166,896 | −0.09 |
| New York central | 521,977 (0.85) | 842,267 (1.37) | 766,146 | −0.09 | 75,436 | 126,166 | 119,151 | −0.05 |
| New York east | 35,300 (0.14) | 83,827 (0.34) | 84,244 | 0.00 | 9840 | 25,361 | 9831 | −0.60 |
| New York west | 214,957 (1.03) | 432,758 (2.07) | 425,102 | −0.01 | 51,886 | 89,171 | 88,390 | −0.01 |
| Pennsylvania north-central | 236,622 (0.32) | 394,575 (0.54) | 377,223 | −0.04 | 48,243 | 85,696 | 85,090 | −0.01 |
| Pennsylvania northeast | 229,234 (1.01) | 370,372 (1.63) | 384,873 | 0.04 | 32,121 | 56,535 | 55,482 | −0.02 |
| Pennsylvania northwest | 123,451 (0.18) | 286,435 (0.42) | 280,295 | −0.02 | 29,886 | 70,959 | 71,381 | 0.01 |
| Pennsylvania south-central | 235,264 (0.45) | 418,710 (0.80) | 406,840 | −0.02 | 57,983 | 109,305 | 102,917 | −0.05 |
| Pennsylvania southeast | 1354,671 (2.34) | 2045,556 (3.53) | 2024,113 | −0.01 | 248,499 | 349,571 | 315,144 | −0.09 |
| Virginia central | 197,389 (0.25) | 495,483 (0.63) | 502,965 | 0.02 | 45,043 | 118,785 | 129,152 | 0.08 |
| Virginia Richmond–Norfolk | 1037,356 (1.86) | 1608,125 (2.88) | 1507,658 | −0.06 | 124,248 | 203,814 | 213,845 | 0.04 |
| Virginia south | 102,681 (0.44) | 289,000 (1.24) | 262,528 | −0.09 | 18,828 | 62,161 | 56,182 | −0.09 |
| Virginia south-central | 142,646 (0.24) | 467,646 (0.79) | 494,381 | 0.06 | 35,359 | 122,487 | 121,610 | −0.01 |
| Virginia west | 75,543 (0.10) | 226,422 (0.29) | 219,623 | −0.03 | 23,871 | 65,167 | 68,078 | 0.00 |
| Washington, DC–Baltimore, MD | 2211,517 (4.53) | 3031,176 (6.21) | 2973,476 | −0.01 | 229,441 | 299,615 | 275,798 | −0.07 |

Table 6
Calibration accuracy results for counties, HUC 11 watersheds, and an array of 7290 × 7290 m grid cells. Ordinary least squares (OLS) regression scores are presented for both estimates of total developed area and for estimates of change in developed area between 1990 and 2000.

| | Estimated developed area (r^2) | Estimated change in developed area, 1990–2000 (r^2) |
|---------------------------------|------------------------------------|---|
| Counties, $N = 208$ | 0.97, $p < 0.01$ | 0.74, $p < 0.01$ |
| HUC 11 watersheds, $N = 505$ | 0.98, $p < 0.01$ | 0.65, $p < 0.01$ |
| 7290 × 7290 lattice, $N = 5126$ | 0.97, $p < 0.01$ | 0.74, $p < 0.01$ |


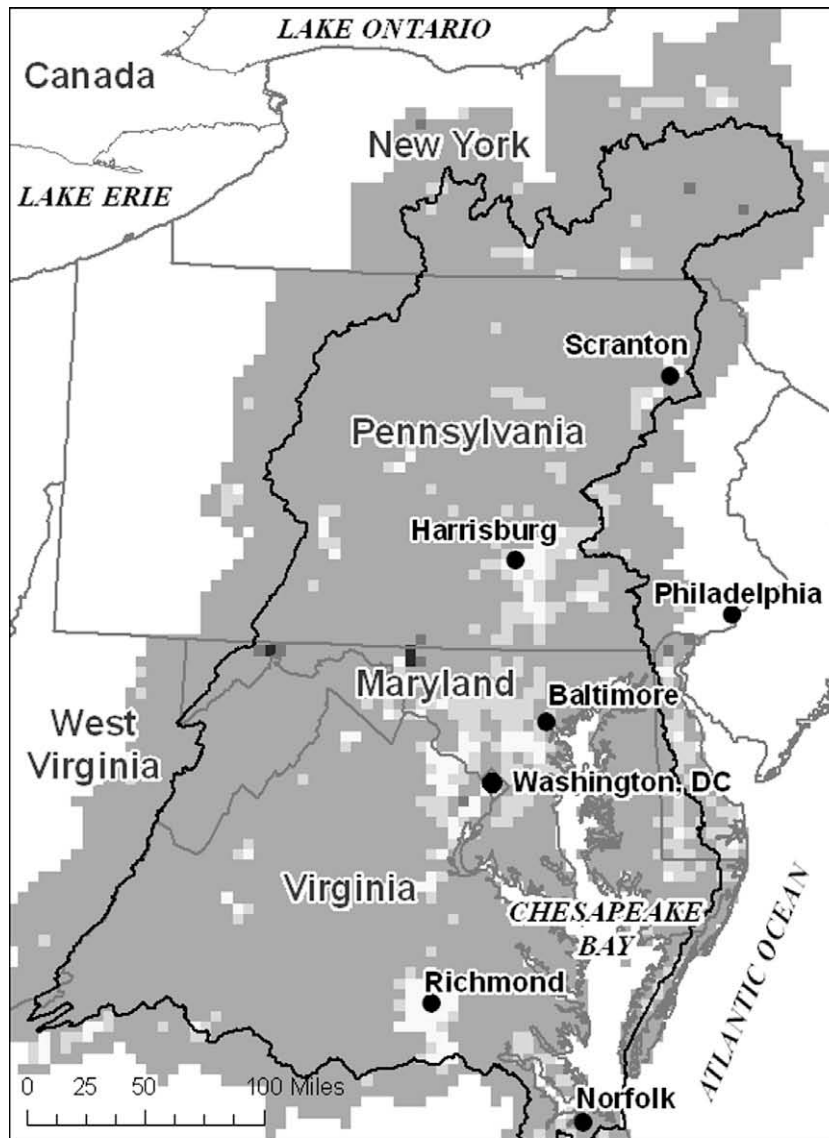


Fig. 7. Difference in estimates of percentage developed area, where the simulated estimates were subtracted from the mapped estimates. Negative values thus represent that SLEUTH-3d is under-estimating development in 2000, while positive values represent an overestimation.

values ranging from 0.001 to 0.130 for the Chesapeake Bay watershed sub-regions (Table 3), illustrates both the heterogeneity of urbanization patterns found in the study area and the model's new ability to adapt to these conditions. While it is outside the scope of this paper to address how these patterns may reflect the

process of urban growth, the diffusion multiplier could offer new insight into these questions. The new fit metrics enable application of SLEUTH-3r in areas that lack more than two data sets representing historic urban land cover. We found the fractional difference metrics (PFD, CFD) particularly useful because they quantified both

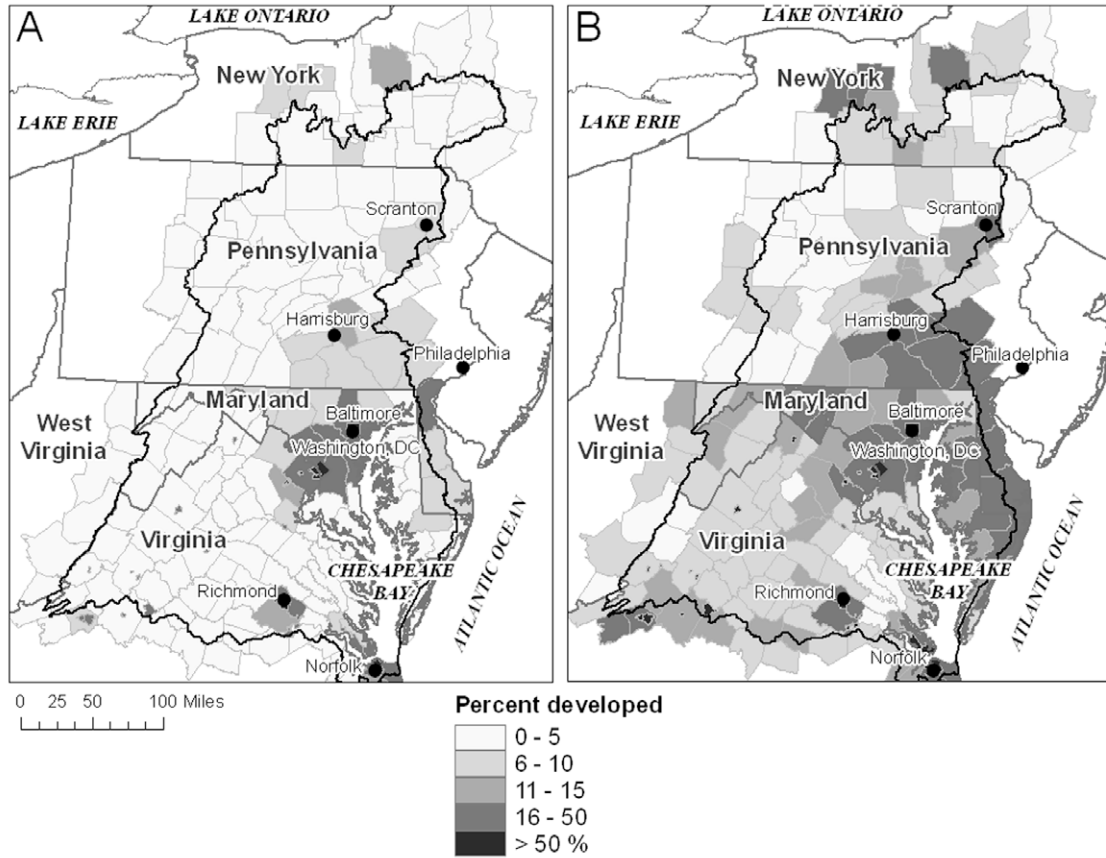


Fig. 8. Basin-wide forecasts to 2030. The percentage of each county's area that is urbanized in 2000 is shown in (A), and the forecast for 2030 is shown in (B).




Fig. 9. Forecast results for southeast Pennsylvania, 2000–2030.

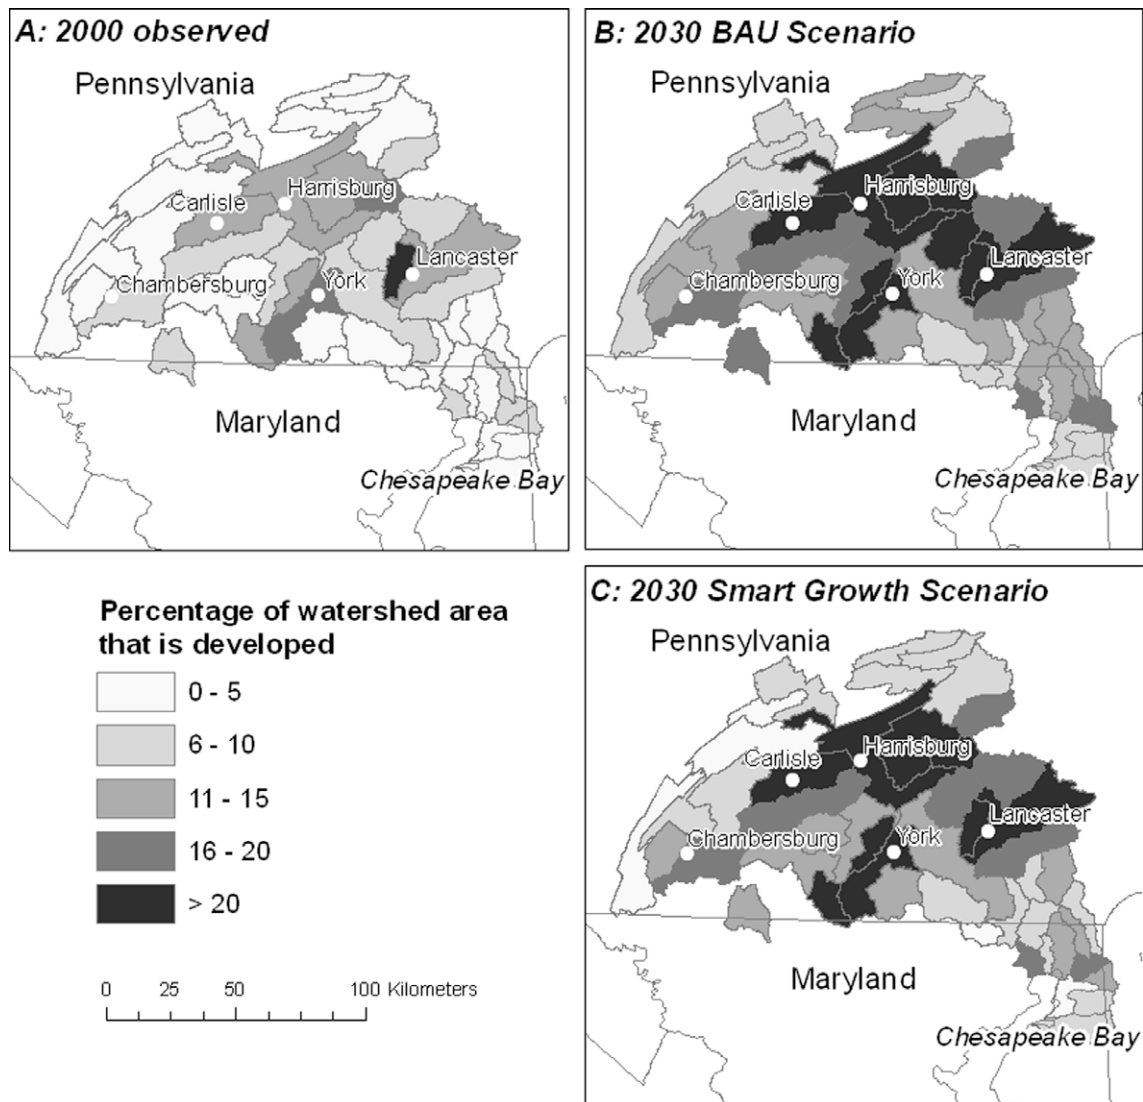


Fig. 10. Example forecast results for southeast Pennsylvania. The percentage of watershed area that is developed in 2000 is shown in (A), and 2030 forecasts for the Business as Usual (BAU) linear trend scenario and the Smart Growth linear trend scenario are shown in (B) and (C).

the model's fit and the direction of error. Finally, the enhanced processing speed and memory allocation allow for processing of larger input layers, expanding SLEUTH-3r's ability to model larger areas with fine-grained (small) cells, thus better capturing sub-unit variability (e.g. within grid cells, watersheds, regions, districts).

4.2. Sub-dividing the Chesapeake Bay watershed

We found the input of and collaboration with the Chesapeake Bay Program invaluable throughout the modeling process, particularly for sub-dividing the study area both in terms of the selection of variables for the *k*-means clustering, and in the production of the final subdivision units. The use of political boundaries, such as counties and states, was based on the recognition that governmental jurisdictions influence land-use change, while variables such as the rural–urban commuting areas provide linkages between metropolitan centers and their suburban and exurban counties. The inclusion of biophysical variables, such as ecoregion classifications, enabled the model to react to environmental influences on regional urbanization patterns. Finally, many of the variables were directly derived from existing urban patterns or from observed urban land-cover change, and these same variables were used to

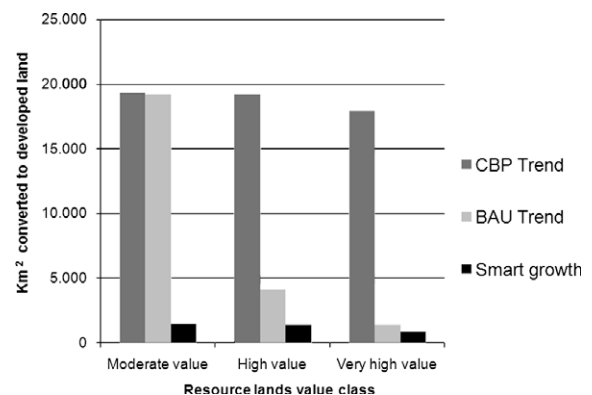


Fig. 11. Forecasted impacts on resource lands for three scenarios for southeast Pennsylvania. Only the higher valued classes are shown here.

drive the calibration of the SLEUTH-3r model. This ensured consistency between the methodology and data used to guide the subdivision of the study area and the modeling approach implemented in SLEUTH-3r. The resulting sub-regions were logical to the

Chesapeake Bay Program partners, a good indication of the pragmatic validity of our choice of 15 sub-regions.

4.3. Calibration of the SLEUTH-3r model

The calibration results show wide variation in the best-fit parameters derived for each sub-region (Table 4), reinforcing the justification for sub-dividing the study area. These results also indicate the sensitivity of the model to the input data. For example, the best score for the clusters fractional difference metric for the New York east sub-region was -0.60 , indicating a 60% underestimation (Table 5). This case reflects the poor quality of the 1990 data for this sub-region (due to cloud cover in the original satellite imagery) rather than the model's inability to capture the observed urbanization rates and patterns. The performance of the model for the other sub-regions was quite good at the aggregate level of the sub-regions (Table 5).

While the model also performs well at finer scales, trade-offs become apparent. In examining the difference between the amount of urbanization estimated by the model and the observed urbanization for the year 2000 (Fig. 8), areas where the model consistently overestimated tended to fall within the exurban or rural landscapes. In contrast, the areas where the most significant underestimation occurred were associated with urban centers. This is particularly evident for Richmond, VA, Washington, DC and Harrisburg, PA. These results indicate the challenge of capturing local scale heterogeneity in densely developed areas. Nonetheless, we note that for nearly the entire CBW watershed we are able to capture the amount of urban development within 5%, which we feel are particularly positive results given the scale and scope of this project.

4.4. Forecasts to 2030

The basin-wide forecasts presented here (Fig. 8) were created based on a set of assumptions that reflect "business as usual," so it is not surprising that the results indicate an intensification of historic development patterns. These results do, however, provide an important baseline from which alternative scenarios can be evaluated, both in terms of the spatial pattern of development, potential impacts on resource lands, and impacts on water quality and hydrology. These are on-going efforts at the time of this writing, as noted below in Section 5.

The forecasts presented for southeast Pennsylvania represent examples of how alternative future scenarios can be developed in SLEUTH-3r. The use of the exclusion/attraction layer and the use of self-modification to simulate high, medium and low growth rates, are innovations that warrant special attention in future applications. While we did not conduct quantitative sensitivity analyses related to these innovations, we have demonstrated that the utility of this approach is both important and promising. We were able to evaluate alternative futures that exhibited similar levels of new development across sub-regions, but differed in terms of the spatial patterns of development and the types of land converted to new development (Figs. 7 and 8). These results can be used as input to other models to quantify impacts on, for example, water quality, flood risk, or wildlife habitat (e.g. Goetz, Jantz, & Jantz, 2009). It is precisely this type of information that is important for ecosystem management and sound land use decision-making (e.g. Jantz & Goetz, 2007).

5. Conclusion

This paper presents the broadest scale application of the SLEUTH model to date – a rare example of a fine-scale land-cover change model applied across a large region – and introduces a

new version of SLEUTH with substantially augmented functionality. We have introduced new fit statistics that significantly enhance the calibration process, and enable application of the model when historic data are available for only two points in time. The use of relative exclusion/attraction values has expanded the capability of SLEUTH to incorporate economic, cultural and policy information with bearing on historic and future urbanization trends (e.g. Jantz & Goetz, 2007). Taken together, all of these changes open up new avenues for the integration of SLEUTH with other land-change models (e.g. Goetz, Jantz, Towe, & Bockstael, 2007). Future research will also be able to use the new information generated by SLEUTH-3r to address questions related to how urban patterns relate to the process of urban land-cover change. In addition, we have made significant advances in the model's computational efficiency.

The Chesapeake Bay Program is currently using the results from SLEUTH-3r to prepare future land use inputs for their Hydrologic Simulation Program-Fortran (HSPF) watershed model. Because we were able to preserve the high-resolution of the land cover data, SLEUTH-3r provided the capability of visualizing alternative future scenarios at a detailed scale, which helped to engage stakeholders in the scenario development process. SLEUTH-3r and HSPF are especially good complements because HSPF, as a lumped parameter model, tends to dampen and correct any absolute spatial errors in SLEUTH's forecasts by aggregating results to broader spatial scales (larger grain sizes).

We believe this project represents an important advancement in computational modeling of urban growth. In terms of simulation modeling, we have presented several new advancements in the SLEUTH model's performance and capabilities. More importantly, however, this project represents a successful broad scale modeling framework that has direct applications to land use management.

Acknowledgements

We acknowledge the support of the NASA Land Cover Land Use Change program (Grant NNG06GC43G to SJG), the EPA Science to Achieve Results program (Grant R82868401 to SJG), the US Geological Survey Geographic Analysis and Monitoring Program and the EPA Chesapeake Bay Program. We thank Rob Burgholzer for his assistance in developing the CBP trend scenario, and William (Skip) Little for some of the initial SLEUTH code modifications. We also thank Jules Opton-Himmel, Dan Steinberg and Greg Fiske at WHRC, and Kyle Shenk and Gary Lasako at Shippensburg University of Pennsylvania, for assistance with GIS data processing. We also acknowledge the helpful and constructive reviews provided by four anonymous reviewers.

References

- Aldenderfer, M. S., & Blashfield, R. K. (1984). *Cluster analysis*. Newberry Park, CA: Sage Publications.
- Batty, M. (1997). Cellular automata and urban form: A primer. *Journal of the American Planning Association*, 62, 266–275.
- Candau, J. (2002). *Temporal Calibration Sensitivity of the SLEUTH Urban Growth Model*. MA thesis. Department of Geography, University of California, Santa Barbara, CA.
- Chesapeake Bay Program (2000). Chesapeake 2000: A watershed partnership. <<http://www.chesapeakebay.net/agreement.htm>>.
- Chesapeake Bay Program (2005). Resource lands assessment. <<http://www.chesapeakebay.net/rla.htm>>.
- Clarke, K. C., Hoppen, S., & Gaydos, L. J. (1997). A self-modifying cellular automaton model of historical urbanization in the San Francisco Bay Area. *Environment and Planning B: Planning and Design*, 24, 247–261.
- Clarke, K. C., & Gaydos, L. J. (1998). Loose-coupling a cellular automaton model and GIS: Long-term urban growth prediction for San Francisco and Washington/Baltimore. *International Journal of Geographical Information Science*, 12, 699–714.
- Couclelis, H. (1997). From cellular automata to urban models: New principles for model development and implementation. *Environment and Planning B: Planning and Design*, 24, 165–174.

- Dietzel, C., & Clarke, K. C. (2007). Toward optimal calibration of the SLEUTH land use change model. *Transactions in GIS*, 11, 29–45.
- Environmental Systems Research Institute. (2003). US streets (data set). Redlands, CA.
- Goetz, S. J., Jantz, C. A., Prince, S. D., Smith, A. J., Varlyguin, D., & Wright, R. (2004). Integrated analysis of ecosystem interactions with land use change: The Chesapeake Bay watershed. In R. S. DeFries, G. P. Asner, & R. A. Houghton (Eds.), *Ecosystems and land use change* (pp. 263–275). Washington, DC: American Geophysical Union.
- Goetz, S. J., Jantz, C. A., Towe, C. A., Bockstael, N. (2007). Modeling the urbanization process across Maryland in the context of Chesapeake Bay restoration. Conference Paper 06 Smart Growth @ 10. College Park, MD. <<http://www.rff.org/irff/Events/SmartGrowthat10.cfm>>.
- Goetz, S. J., Jantz, P. A., & Jantz, C. A. (2009). Connectivity of core habitat in the northeastern United States: Parks and protected areas in a landscape context. *Remote Sensing of Environment*, 113, 1421–1429.
- Jantz, P. A., Goetz, S. J., & Jantz, C. A. (2005). Urbanization and the loss of resource lands within the Chesapeake Bay watershed. *Environmental Management*, 36, 343–360.
- Jantz, C. A., & Goetz, S. J. (2007). Can smart growth save the Chesapeake Bay? *Journal of Green Building*, 2, 41–51.
- Jantz, C. A., & Goetz, S. J. (2005). Analysis of scale dependencies in an urban land-use-change model. *International Journal of Geographical Information Science*, 19, 217–241.
- Jantz, C. A., Goetz, S. J., & Shelley, M. K. (2004). Using the SLEUTH urban growth model to simulate the impacts of future policy scenarios on urban land use in the Baltimore–Washington metropolitan area. *Environment and Planning B: Planning and Design*, 31, 251–271.
- Li, X., & Liu, X. (2006). An extended cellular automaton using case-based reasoning for simulating urban development in a large complex region. *International Journal of Geographical Information Science*, 20, 1109–1136.
- McGarigal, K., Marks, B. J. (1995). FRAGSTATS: Spatial pattern analysis program for quantifying landscape structure. USDA Forest Service General Technical Report PNW-351. <<http://www.umass.edu/landeco/pubs/Fragstats.pdf>>.
- O'Sullivan, D., & Torrens, P. M. (2000). Cellular models of urban systems. In S. Bandini & T. Worsch (Eds.), *Theoretical and practical issues on cellular automata* (pp. 108–116). London: Springer.
- Silva, E. A., & Clarke, K. C. (2002). Calibration of the SLEUTH urban growth model for Lisbon and Porto, Spain. *Computers, Environment and Urban Systems*, 26, 525–552.
- Silva, E. A., & Clarke, K. C. (2005). Complexity, emergence and cellular urban models: Lessons learned from applying SLEUTH to two Portuguese metropolitan areas. *European Planning Studies*, 13, 93–116.
- Soares-Filho, B. S., Nepstad, D. C., Curran, L. M., Cerqueira, G. C., Garcia, R. A., Ramos, C. A., et al. (2006). Modelling conservation in the Amazon Basin. *Nature*, 440, 520–523.
- Torrens, P. M. (2006). Simulating sprawl. *Annals of the Association of American Geographers*, 96, 248–275.
- Torrens, P. M., & O'Sullivan, D. (2001). Cellular automata and urban simulation: Where do we go from here? *Environment and Planning B: Planning and Design*, 28, 163–168.
- US Bureau of the Census (1990). Census 1990 summary file 1 (SF 1) 100-percent data.
- US Bureau of the Census (2000). Census 2000 summary file 1 (SF 1) 100-percent data.
- US Environmental Protection Agency (2003). Level III ecoregions of the continental United States. <<http://www.epa.gov/waterscience/basins/metadata/ecoreg.htm>>.
- US Geological Survey (2007). Project gigalopolis: Urban and land cover modeling. <<http://www.ncgia.ucsb.edu/projects/gig/>>.
- USDA Economic Research Service (2000). 2000 rural–urban commuting area codes. <<http://www.ers.usda.gov/Data/RuralUrbanCommutingAreaCodes/>>.
- Van Vliet, J., White, R., & Dragicevic, S. (2009). Modeling urban growth using a variable grid cellular automaton. *Computers, Environment and Urban Systems*, 33, 35–43.
- White, R., & Engelen, G. (1997). Cellular automata as the basis of integrated dynamic regional modelling. *Environment and Planning B: Planning and Design*, 24, 235–246.
- Yang, X., & Lo, C. P. (2003). Modelling urban growth and landscape changes in the Atlanta metropolitan area. *International Journal of Geographical Information Science*, 17, 463–488.