
UrbanSim: Using Simulation to Inform Public Deliberation and Decision-Making

Alan Borning¹, Paul Waddell², and Ruth Förster³

¹ Dept. of Computer Science & Engineering, University of Washington, Box 352350, Seattle, Washington 98195, USA, borning@cs.washington.edu

² Evans School of Public Affairs, University of Washington, Box 353055, Seattle, Washington 98195, USA, pwaddell@u.washington.edu

³ Institut für Raum- und Landschaftsentwicklung, ETH (Swiss Federal Institute of Technology), Zürich, Switzerland, ruth.foerster@irl.baug.ethz.ch

Summary. Decisions regarding urban transportation investments such as building a new light rail system or freeway extension, or changes in land use policies such as zoning to encourage compact development and curb low-density urban sprawl, have significant and long-term economic, social, and environmental consequences. Further, land use and transportation decisions interact. Integrated simulation models can help government agencies and citizens make more informed decisions about such issues. We briefly describe the history of urban models, and present a taxonomy of needed refinements to them. We then present a case study of the UrbanSim model system and a significant application of it to the Puget Sound Area in Washington State. We also situate this effort in the digital government research arena, in which is presents a significant research opportunity: hard technical problems, unmet demand from government users, and important issues around supporting a more democratic planning process.

1 Introduction

Decisions regarding major urban transportation investments, as well as regarding policies to improve air quality and to manage urban development to reduce the adverse effects of low-density urban sprawl, are critical, interdependent choices that shape the long-term quality of life in urban areas. These choices and the problems they attempt to address have important social, economic and environmental impacts that spill over jurisdictional boundaries and that are impacted by decisions made by a wide range of institutions. In the United States, they fall into the scope of metropolitan governance, where the institutional frameworks for forming and implementing policy are less robust than at higher or lower levels of government. These metropolitan governance structures hover between the vise-grip of local governments' control of land use decisions, and the state and federal control of resources for transportation

and environmental regulations. In the gap, Metropolitan Planning Organizations (MPOs) have been created by states under federal requirements to better coordinate the allocation of federal investments in transportation, and air quality planning. These MPOs generally do not have any taxing or direct implementation or operational responsibility, but are charged with creating regional transportation plans and coordinating these with land use and air quality planning. It is a tall order.

Putting institutional difficulties aside for the moment, the task of developing regional transportation plans is complex enough at a technical level. How can an almost infinite list of alternative transportation investments proposed by local governments, states, and other entities be examined systematically and an investment plan adopted that reflects a democratic process based on a robust assessment of the alternatives? Over the past several decades, MPOs and their predecessor institutions have used simulation models to predict the volumes of traffic on the transportation network, given assumptions about the land use patterns that would generate patterns of travel demand on this network. The traditional models are called ‘four-step models’ because they break this task into (1) predictions of the number of trips generated and attracted in each zone of the metropolitan area, (2) the trip distribution patterns from zone to zone, (3) the mode choice of trips (automobile, transit, etc.) between any two zones, and ultimately, (4) how these trips are assigned to the capacity-constrained network, leading to patterns of travel time and congestion. These four-step models were originally developed within the discipline of civil engineering in the late 1950’s and early 1960’s to address a very specific problem: how to estimate the amount and location of additional road capacity needed to satisfy a given demand for transportation. They became ingrained into the planning process for transportation, reinforced by federal investment and regulation. In the 1960’s and 1970’s, the four-step travel models were brought into mainstream use and became the mainstay analysis tool used to support decisions on alternative road investments.

Since the 1980’s, however, the models and the decision-making process have come under increasing scrutiny and criticism, leading to substantial pressure to revise both [3]. One of the central criticisms is that the models, and the way they have been generally used, assume that changes in land use result in different demands on the transport system, but that changes in the transportation system do not cause land use changes to occur. Aside from the mountain of theoretical and empirical evidence to the contrary, this assumption violates common sense. Building a major highway through farmlands cannot be expected to have absolutely no impact on the probability that sites along the new highway, or accessible to it, will develop. And if there is an impact on development, the logical extension is that it will in turn impact travel demand. This idea is what has been referred to as induced demand, and one of the reasons scholars have become increasingly skeptical that it is possible to “build your way out of congestion” (see [10] for example). Since the U.S. Clean Air Act Amendments of 1990 and the Intermodal Surface Trans-

portation Efficiency Act of 1991, federal policy has recognized the need to link transportation and land use, in order to account for this relationship. Since that time, refinement of transportation planning practice has been slow, partly due to the technical difficulties of accounting for the interactions, and partly due to political constraints and the increasing role of public involvement in decision-making processes such as these.

Early use of technology such as transportation models to support transportation investments dates to a conception of planners as technocrats who provide answers that are to be taken at face value and used as an objective basis for public decisions. Public participation in these decisions, and in the technical analyses behind them, was decidedly not on the agenda. Much has changed since then, especially at the local government level. An increasingly sophisticated and skeptical set of stakeholders demands public participation, as well as transparency and access to information about the decision-making process and the assumptions and analyses behind it. Conflicting interests are played out in public meeting after public meeting and in committee after committee that is deliberating land use policies or transportation investments. Environmental advocates have increasingly come to use the courts to prod planning agencies to refine their analyses to address shortcomings such as the omission of land use feedback effects [16].

1.1 Urban Modeling as a Digital Government Research Area

The domain of land use and transportation modeling thus provides an significant opportunity for digital government research: it is of great interest to government agencies, and it includes a set of hard, open problems, both technical and procedural. This chapter is intended for digital government researchers and students who are generally computer- and policy-literate, but who are not necessarily expert in either the domain of urban modeling or of land use and transportation policy. In the chapter, we first present a taxonomy of needed refinements to urban models themselves, and to the process of applying them. We then present a case study of UrbanSim, an urban modeling system that our group has been developing at the University of Washington, including a short history, more recent research initiatives, and some significant applications to planning activities.

Our focus in this chapter is primarily on the U.S. context. However, controversies regarding land use and transportation occur world-wide, and analogous issues arise around using models to inform decision-making in other countries.

1.2 A Taxonomy of Model and Process Refinements

Our research is intended to contribute both to improving the technical modeling capacity to address issues such as the land use consequences of transportation investments, as well as to improving the process of using models in a democratic decision-making context. To help structure this case study,

- Refinement of Models
 - Validity
 - Accuracy
 - Handling of uncertainty
 - Policy sensitivity
 - Comprehensiveness
 - Real estate development and prices
 - Employment location
 - Household location
 - Transportation system
 - Environmental impacts
- Refinement of Process
 - Refinement of the model construction and application process
 - Feasibility of data preparation
 - Performance
 - Usability
 - Support for software evolution
 - Support for a more effective democratic process
 - Responding to stakeholder interests and concerns
 - Transparency
 - Fairness
 - Facilitating stakeholder access to models and their output

Table 1. Model and Process Refinements

as well as providing a framework for evaluating urban models, we offer the following taxonomy of model and process refinements (Table 1). We hope that this taxonomy will be of value beyond this particular case study as well, for other studies of modeling and simulation in the policy arena. In developing this framework, we draw on and extend earlier work that has criticized earlier urban models (for example [18]). We then describe how our project and several research initiatives within it have emerged to address these challenges.

Refinement of Models

At the top level, we distinguish between *refinement of models* and *refinement of process*. Refinement of models focuses on the models themselves. In turn, we can classify the work on refinement of models as work on *validity* and on *comprehensiveness*.

Validity includes improving the accuracy of the models, and also their sensitivity to policies of interest. Accuracy means that the predicted values (for example, of population density in different neighborhoods) are close to the observed values. This raises the obvious problem of how to evaluate the accuracy of predictions of events in the future. One technique is *historical validation*, in which the model is run on historical data, and the results compared with what actually transpired (see [27] for example). This has the clear

merit of comparing with real outcomes. There are difficulties as well, however. First, in many cases the needed historical data is not available. Also, for the relatively small number of regions for which data is available, there may not have been major land use and transportation changes over the period being tested, so that the model in effect isn't being used to simulate major decisions. An alternative technique that is often used is to run the model system with fairly extreme scenarios (e.g. doubling the capacity of selected roadways, or removing zoning restrictions on height limits in a neighborhood). The results are then evaluated by an expert review panel.

Predicting the future is a risky business. There are numerous, complex, and interacting sources of uncertainty in urban simulations of the sort we are developing, including uncertainty regarding exogenous data, the model structure and specification, the parameters of the model, and from the stochastic nature of the simulation. Nevertheless, citizens and governments do have to make decisions, using the best available information. Ideally we should represent the uncertainty in our conclusions as well as possible, both for truthfulness and as important data to assist in selecting among alternatives. However, to date there has been only a small amount of work done on handling uncertainty in urban modeling in a principled fashion [24].

We often also want to improve the sensitivity of the model to policies of interest. For example, if a region is interested in policies that foster walkable neighborhoods, then the model should be able to model walking for transportation as well as for health and recreation. Which policies are of interest is of course a political and societal question; but given such policies, whether the model responds suitably to them becomes a question of validity.

Yet another sort of refinement of the models is increasing their comprehensiveness to include other actors and processes in the urban environment. For example, for households, we might model additional demographic processes, such as household formation and dissolution. Or for environmental impacts, we might model consumption of additional kinds of resources, or the impacts of decisions on biodiversity as well as on particular species of interest (for example, due to Endangered Species Act considerations).

There are important pitfalls and tensions associated with the goal of increasing the comprehensiveness of models: namely what Lee [18] called the problem of hyper-comprehensiveness. One aspect of this is pressure to model more and more aspects of the urban environment because these aspects are important to someone — even though they might have little relation to land use and transportation. For example, there might be demands to model voter turnout rates. These pressures are relatively straightforward to deal with, by reminding stakeholders of the purpose of the modeling work and the need to remain focused. A more difficult issue is that a seemingly endless number of factors influence urban land use and transportation. For example, crime is clearly an important factor in residential location choice, in transportation choice, and others. But we need not just data on current crime rates — and perhaps more importantly, on people's perceptions of crime — but also a pre-

dictive model of crime in the future under different possible scenarios. This is both difficult and controversial. For example, what are the major determinants of the crime rate? Economic conditions? Family stability and moral instruction? The nature of the criminal justice system? How far should the modeler go down this path? Or as another example (relevant to the region around Seattle), suppose we want to model the return rate of wild salmon in rivers and streams that flow through urban regions. There are many factors affecting this: the amount of impervious surface, pollutants from agricultural runoff, the number of fish caught by both commercial and sport fishers, oceanic conditions (including temperature, since the salmon grow to maturity in the ocean before returning to fresh water streams to spawn), and many others. Among the pitfalls of overly ambitious modeling are increasing model complexity, additional data requirements, and in some cases the credibility of the overall modeling effort.

Refinement of Process

Returning to the top level of the taxonomy, refinement of process includes first, improving the process of developing, extending, and applying models; and second, supporting their more effective use in a democratic society. The first of these is concerned with instrumental values such as usability and feasibility: data preparation issues, adequate performance, usability of the software, and accommodating changes in requirements, data, and the like. It must be feasible to prepare the data needed to run the model. Typically, this implies that the data must already be in hand — collecting new data is enormously expensive. But the data in hand may be of varying quality, or in the wrong format, and so forth. Performance should be adequate, and the software must be usable by the technical staff at the planning organization. Also, the model system architecture should allow for the system to evolve as requirements and the questions asked change over time.

Another set of process issues revolve around the desire to use modeling as part of a more participatory, open, and democratic process, rather than in a technical, back-room exercise. One aspect of this is improving the relevance of the modeling and output to the diverse range of stakeholder concerns (in other words, increasing its comprehensiveness in response to stakeholder values). Transparency of the model itself, of the input data preparation process, and of the overall context in which it is used also play an important role as well. Another aspect of this is improving the fairness of the model (for example, in not omitting an important transportation mode, or short-changing the interests of renters as compared with home owners). Again, this can result in additional demands for refinement of the model (either its validity, comprehensiveness, or both). The results of running the model, and ideally even the ability to experiment with alternatives, should be opened up to a wider range of stakeholders, rather than being restricted to the technical modelers. System performance is relevant here also: for example, if the model takes weeks

to run, clearly this would make it difficult to use to support deliberation, in which model results are discussed, and in response new questions are asked of the model or new scenarios are proposed for testing.

There are some obvious tensions among these objectives for refinement of models and process. Pressures to increase policy sensitivity in order to avoid bias from omission of certain policies from consideration, for example pedestrian and bicycling modes, increase the need for a very high level of behavioral and spatial detail. This will certainly come at a cost in performance, and quite possibly also at a cost of some reduction in the accuracy of the results. How can model sensitivity, data requirements, transparency, computing performance, and accuracy be compared against each other? How are the interests of different stakeholders served by alternative compromises among these? How do these choices affect the legitimacy of the model system and the process for using it in the decision-making process? These are difficult problems, and ones that have not received sufficient attention to date. We seek to address these concerns in our project in addition to the more purely technical issues of model refinement.

2 The UrbanSim Project: A Case Study

UrbanSim is an open-source land use modeling system, which integrates with an external travel model to provide a capability for integrated modeling of urban land use, transportation, and environmental impacts. It is reusable, in the sense that it can be applied to a new region by supplying the appropriate data, as opposed to being for a single region only. Its development has been motivated by the desire to refine both the models and processes, as described in Section 1.2.

Our research agenda draws broadly from the academic fields of urban planning and design, public policy, economics, geography, sociology, civil engineering, and computer science, and balances academic research interests with pragmatic relevance in addressing important social problems and public policy choices. Engagement with local and metropolitan governments and planning processes has been essential to this agenda. This engagement informs the research in numerous ways, and motivates it. We therefore devote attention in this chapter to the process of adoption of UrbanSim by public agencies for operational use, and examine two case studies of applying UrbanSim in the metropolitan regions around Salt Lake City and Seattle. In this section we first review the initial development of the UrbanSim project, and then describe some of the significant research initiatives that we have launched in response to lessons learned along the way.

2.1 Related Work

UrbanSim has been developed over the past decade, and has both benefited from prior work and, we think, substantially extended prior modeling ap-

proaches. While space prevents a complete treatment of how UrbanSim compares to other modeling approaches, we offer here a concise treatment, along with pointers to more in-depth reviews of urban models, including UrbanSim itself [9, 20, 25, 31]. We begin by contrasting UrbanSim with two alternative modeling approaches that were available by the mid-1990's, when the design and development of UrbanSim began. The approach that has been most widely used to do land use modeling is based on the Lowry gravity model [19], which adapted the law of gravity from physics and applied it to predict the flow of travel between locations as a function of the sizes of origin and destination and the ease of travel between them. This approach, also widely referred to as a spatial-interaction approach, was used widely in transportation modeling to predict trip distribution patterns, was later refined and adapted to predict land use patterns in the DRAM/EMPAL model system developed around 1970 [22]. This approach has been criticized as over-emphasizing the role of transportation in location choices, and lacking behavioral content such as a representation of real estate markets (housing supply, demand, and prices are not considered, only household and job locations in zones). An alternative approach is the spatial input-output approach that was also initially developed around 1970, adapting the input-output models developed to describe monetary flows in the U.S. economy to predict monetary flows between economic sectors and between zones [8, 11]. This is roughly analogous to an international trade model in which the countries are replaced by some geographic subdivisions of the study area. These spatial input-output models translate from monetary flows into tons of freight, and by dividing the monetary flow from labor to a sector by the average wage rate in the sector, derive the number of workers and the implied commuting flows. While this approach is intriguing for large-scale areas where there is interest in predicting inter-metropolitan freight movement, it seems less suited to developing a behaviorally-transparent model of how households choose where to live or business choose to locate, or of the behavior of real estate developers. Unlike the spatial-interaction approach, the spatial input-output approach does recognize the role of markets, and treats demand, supply, and prices. Both of these modeling approaches share the approach used by virtually previous all urban models: they formulate an equilibrium set of conditions and solve for an equilibrium that is not path-dependent, rather than attempting to predict dynamic changes over time.

UrbanSim differs significantly from these and other earlier modeling approaches in several respects. First, it emphasizes behavioral theory and transparency. This leads to a more explicit treatment of individual agents such as households, jobs, and locations, and to a microsimulation of the choices that these agents make over time. UrbanSim was the first operational modeling system to use very small geographic units such as parcels or small grid cells, and to microsimulate individual agents and their choices in a framework that includes an explicit representation of real estate demand, supply, and prices and their evolution over time. One other approach that has emerged

recently warrants comment. The use of rule-based models that are embedded in highly graphical environments such as Geographic Information Systems has increased in the past several years. These are not behavioral models in the sense that they do not attempt to reflect the pattern of agent behavior observed in the real world. Instead, these tools allow a user to impose on the system a set of assumptions, such as the distribution of population twenty years in the future, and then to visually examine the result of applying those assumptions. While the intent is quite different than the ones we have outlined above, there is some risk that users may misinterpret the output of such tools as containing more behavioral information than it in fact does. (In the above example, simply asserting that a given neighborhood will have a given population density is very different from simulating the results of policies, such as transportation and zoning changes, designed to achieve that result. After all, the policies might or might not achieve the desired result.) There may in fact be some useful complementarity between these kinds of visually-engaging tools and more behaviorally-focused models: the visual tools could provide a useful interface to obtain information from stakeholders about their preferences, and the behavioral models could provide stakeholders feedback about how effective alternative policy actions might be in achieving those objectives, and to help examine what the trade-offs are among the objectives.

2.2 Early Development of UrbanSim

Paul Waddell originally became interested in the problem of land use and transportation modeling in the 1980's, when he worked for a metropolitan planning organization and encountered first-hand the unsatisfactory state of modeling practice. At that time (and it is still widely the case today), the models available for use in metropolitan planning lacked behavioral realism, had insufficient theoretical underpinnings, were largely insensitive to important public policies, and as a cumulative result of these considerations, lacked credibility with both technical planners and with policy-makers. These models were, in short, ineffective.

In the mid-1990's Waddell began designing and developing UrbanSim to address some of the clear shortcomings of existing models to support metropolitan planning. The design choices in this early work focused on improving behavioral realism, validity, and policy relevance. UrbanSim was designed based on a micro-simulation approach to modeling household choices of residential location, business location choices, and real estate development and prices, and the use of a dynamic, annual simulation of the evolution of cities over time. It attempted to make explicit the role of policy inputs, such as comprehensive land use plans and transportation infrastructure, in order to support the evaluation of alternative policy scenarios and their effects. Waddell supervised the development of a prototype version in Java for testing in Eugene, Oregon [26], supported by a contract with the Oregon Department of Transportation.

In 1998, Alan Borning began collaborating with him on this work, based on common interests in using information technology to improve the process of metropolitan planning. In 1999 we obtained the first of a series of grants from the National Science Foundation to support the project, this one from the Urban Research Initiative. This more flexible funding allowed us to put together a team consisting of graduate and undergraduate computer science students, along with students from urban planning, civil engineering, and economics, who rewrote the model and its interface (still in Java).

2.3 Subsequent Research Initiatives

After the initial development work on the UrbanSim prototype and a first effort to improve the software engineering and design of the system, three priorities emerged. First, we wanted to further refine the models, to improve their validity and comprehensiveness (see the taxonomy in Section 1.2). Second, we wanted to explore more systematically the issues around opening the process of using the model to more participation, as well as carefully considering issues of transparency, fairness, and the range of stakeholders affected by the decisions that were to be informed by information coming from the model system. These priorities in turn gave rise to the third priority, namely, addressing the significant computational and software engineering challenges that then arose due to the level of detail the system was attempting to represent, the complexity of the models, and the need to support system evolution and change.

Guided by these priorities, we submitted two additional NSF proposals. In 2001 we were fortunate enough to be awarded a 5-year NSF grant under the Information Technology Research Initiative, administered through the NSF Digital Government program, along with a companion grant from the Digital Government program itself. This increased level of funding and longer funding period enabled us take on a substantially larger and more risky research agenda, while still pursuing our overall strategy of integrating the research with close collaboration with local and metropolitan governments and planning processes. In the remainder of this section we summarize these initiatives. Relating back to the taxonomy, these initiatives all apply to various items in the taxonomy of refinements, but often an initiative will pertain to several — there is usually not a one-to-one mapping between initiatives and items in the taxonomy.

Improving Model Validity

A significant amount of the work in the past several years has concerned improving the validity of the models. Much of this work has been quite technical, and for purposes of this case study, we mention two of these activities, along with providing references to papers that provide additional detail.

The location choice models (for Household Location Choice and Employment Location Choice), for example, have undergone a series of improvements. Despite this, throughout the UrbanSim 3 efforts these two models suffered from a problem of dispersion. For example, highly desirable residential areas — e.g. with a lake or ocean view — typically have higher real estate values and wealthier residents. Absent some significant external cause we would expect this to continue; yet in the simulation results we were seeing a dispersion of wealthy households from these neighborhoods, as well as poorer households moving into these desirable areas. (While for policy reasons we might desire more mixed neighborhoods, achieving this would require policy changes — the effect we were seeing was due to some combination of model specification and software implementation.) This shortcoming was an issue in an assessment of the application of UrbanSim in Salt Lake City (see Section 2.4, and a full description in reference [30]). It was only in our most recent version of the system, UrbanSim 4 (Section 2.3), that we were able to address this. This problem was one of several that precipitated a complete redesign of the software implementation, since diagnosing and addressing it proved to be too challenging in the UrbanSim 3 implementation.

The Developer Model, which simulates the actions of real estate developers, is one of the most complex of UrbanSim’s component models. We have had more versions of this model component than any other. In part this reflects the relatively smaller amount of academic research on real estate supply than on demand. Two approaches are discussed in the real estate literature: the site looking for a use and the use looking for a site. The former approaches the problem from the land owner’s perspective: whether to develop a parcel, and if so, into what use. The latter approaches the problem from the perspective of a specialized developer, say, of office space, who needs to choose a profitable location for the project. Through UrbanSim 3 we adopted the site looking for a use perspective, which was consistent with the literature in land cover change. We learned through several model applications, however, that this approach was problematic from the perspective of policy sensitivity. If we significantly changed the land area available for development, for example, modeling the introduction of an Urban Growth Boundary to constrain low-density development, then the amount of development was significantly reduced, which does not conform to expectations. Instead, we would expect development to continue, but at higher prices, and focused in the areas available for development. So in the development of UrbanSim 4, we re-designed the developer model from the use looking for a site perspective, which has a natural and behaviorally-correct response to a significant change in the area of land available for development: it scales the probabilities among the remaining available sites.

Increasing Model Comprehensiveness — Modeling Environmental Impacts

Another aspect of model refinement is increasing the comprehensiveness of UrbanSim’s component models to encompass additional phenomena. In some cases this is of primary interest as a means of improving model validity. For example, one such effort has concerned improving the Demographic Transition Model, which simulates the creation and removal of households (to simulate the real world processes of births, deaths, children leaving home, and movement into and out of the study area). The primary purpose of this activity is to improve the validity of the location choice models.

Another reason to model additional phenomena is to capture important impacts of land use and transportation decisions, even if these impacts don’t feed back directly into the operation of the other component models. Environmental impacts are the most significant class of such phenomena. In our current implementation, we produce indicators of greenhouse gas emissions from transportation, and have a component model that simulates land cover change. (Land cover is important in its own right, and is also useful as a factor in producing other environmental indicators — for example, impervious surface affects water runoff characteristics.) We plan to add additional indicators in the near future, beginning with air quality and resource consumption indicators.

As noted above, these indicators do not currently feed back into the operation of the other component models — although adding such feedback (e.g. to the residential location choice model) is a research topic for the future. In any case, these indicators of environmental impact become important considerations in the deliberative process in which the model is used.

Value Sensitive Design

The domain of urban planning is both value laden and rife with long-standing disagreements. In addition, as described in the introduction, there is increasing desire to move the modeling and planning process from a technically-focused “back room” operation to a more open and participatory one, which immediately leads to questions of who are the stakeholders and who should participate, along with other issues of transparency, accountability, and openness. In terms of our taxonomy, these issues initially fall into the category of “Support for a more effective democratic process,” but fairly quickly lead to issues regarding the process of constructing and applying the models, as well as to pressures for refinements to the model.

To approach these value issues in a principled fashion, we rely on the Value Sensitive Design theory and methodology [14]. Value Sensitive Design is an approach to the design of information systems that seeks to account for human values in a principled and comprehensive way throughout the design process.

Key features are its interactional perspective, tripartite methodology (consisting of conceptual, empirical, and technical investigations), and emphasis on indirect as well as direct stakeholders.

For UrbanSim in its current form, the direct stakeholders are the urban modelers and technical planners who use UrbanSim and manipulate its results. The indirect stakeholders are those who do not use the system directly, but who are affected by it. They include for example elected officials, members of advocacy and business groups, and more generally all the residents of the region being modeled, as well as residents of nearby regions. One way of framing our goal of opening up the planning process and use of models to wider participation is to move more people from the indirect to the direct stakeholder category: to provide meaningful direct access to the results from the models, and ultimately to the tools for simulating other possible plans and futures.

Early in our conceptual investigations, we made a sharp distinction between explicitly supported values (i.e., ones that we explicitly want to support in the model system) and stakeholder values (i.e., ones that are important to some but not necessarily all of the stakeholders). Next, we committed to several key moral values to support explicitly: fairness and more specifically freedom from bias [15], representativeness, accountability, and support for a democratic society. In turn, as part of supporting a democratic society, we decided that the system should not a priori favor or rule out any given set of stakeholder values, but instead should allow different stakeholders to articulate the values that are most important to them, and evaluate the alternatives in light of these values. We identified comprehensibility, and subsequently legitimation and transparency, as key instrumental values to be supported explicitly as well.

In terms of specific Value Sensitive design research activities, one project involved carefully documenting and presenting the indicators that portray key results from the simulations [4] using a web-based Indicator Browser. We also developed the Indicator Perspective framework, which allows a set of organizations with a diverse set of positions and interests to put forth different perspectives on what is most important in the results from UrbanSim, and how it should be interpreted. Our design for the Indicator Browser and Indicator Perspectives addresses a number of challenges, including responding to the values and interests of diverse stakeholders, and providing relatively neutral technical information and at the same time supporting value advocacy and opinion. A third activity has been the development of “Personal Indicators” that allow individuals and households to ask how different alternatives will affect them personally, for example, in terms of housing costs or individual commute times [6, 7]. Together, these projects represent a significant set of steps toward making the modeling components of the planning process more open and participatory, supported by a solid theoretical foundation.

Software Engineering

The significant performance demands of the system that arise due to the level of detail being represented, the complexity of the models, the close scrutiny of the simulation results, and the need to support system evolution and change, all give rise to pressures for good software engineering and software development practices. The ITR grant, with its funding level and five-year term, enabled us to hire two software engineering professional with substantial industry experience, who were then able to provide much stronger guidance on our software development process, as well as being able to work closely with the students involved with the software side of the project over the years, including a large number of computer science and computer engineering undergraduates.

We put into place an agile software development process [1], which relies on small, incremental development steps, automated testing, and responsiveness to change. *Unit tests* — small, self-contained tests that exercise a particular part of the code — play an essential role in the process [17, 21]. Following a modified eXtreme Programming approach [2], we carry this further and use a test-first development strategy, in which the software developer first writes the unit tests, and then writes the code to be tested. We also rely on an automatic build system that runs all the tests whenever new code is checked into the code repository. A novel feature of our development environment is a set of traffic lights (real ones), placed in the hallway of our lab and in developers’ offices. When the light is green, the checked-in code has passed all tests; when it is yellow, the tests are currently being run on newly-submitted code; and when it is red, one or more tests have failed. A green light serves as a reassuring status indicator of the state of the software, while a red light is a powerful social incentive to fix the software problem, as a top priority. The traffic light and other novel features of our development methodology are described in a paper in the 2003 Agile Development Conference [13], while another novel feature of our process, the use of ‘today’ messages as a coordination tool for software developers, is discussed in more detail in reference [5]. Finally, a recent paper [23] describes a novel extension to the unit testing framework to handle testing stochastic models (such as many of the UrbanSim models).

Opus and UrbanSim 4

In 2002 we rewrote UrbanSim again (still in Java), resulting in UrbanSim 3, which was our production system for several years. The use of Java as an implementation language for UrbanSim was a good choice in many respects, but there were some problems as well, in particular, inaccessibility of the code to domain experts, and performance issues. Regarding accessibility, our hope was that, by using suitable abstractions and a very clear coding style, that modelers (domain experts), as well as software engineers, could experiment

with the models. Alas, due to the inherent complexities of the Java implementation of UrbanSim 3 and the Java development environment, this did not happen. There were also some intractable modeling problems (regarding dispersion of housing and employment) that we were having great difficulty with in the Java version — but which we were able to quickly resolve in some experimental code in the Python programming language. We also found that, unlike Java, the domain experts (the modelers) *were* willing to read and modify Python code. Python itself has poorer performance than Java. However, in recent years, a number of packages have been developed, including Numeric and numarray, that can be called from Python and that allow one to manipulate large arrays and matrices very efficiently. UrbanSim simulations usually process a huge amount of data, which in the Java version was done using iterations written in Java, and is now handled using calls to these efficient packages.

Another problem was the existence of a set of independent research projects at different institutions, all working in the domain of land use and transportation modeling, but that had incompatible platforms and implementations, making it hard to share and build on each other’s work. A growing consensus emerged among researchers in the urban modeling community that a common, open-source platform would greatly facilitate sharing and collaboration. In response, an international group of researchers, including groups at the University of Toronto, the Technical University of Berlin, and ETH (the Swiss Federal Institute of Technology in Zurich), began collaborating with the UrbanSim team to develop a new software architecture and framework — Opus, the Open Platform for Urban Simulation [29].

At the same time that we were designing and implementing Opus, we began implementing UrbanSim 4 using the new platform (using UrbanSim as one of the test cases for the design and implementation of Opus). We continued using our agile, test-driven development methodology with Opus/UrbanSim. We now have a well-functioning system, written in Python, that makes heavy use of efficient matrix and array manipulation libraries. Opus and UrbanSim 4 contain far less code than the previous implementation, yet implement a much more modular and user-extensible system that also runs faster. UrbanSim 4 also incorporates key functional extensions such as integrated model estimation and visualization. In terms of our taxonomy (Section 1.2), these extensions are process refinements that address technical and engineering issues. Model estimation, for example, was formerly an error-prone process involving using external econometric packages, which required several months of skilled staff time. With the integrated estimation capabilities, models can now be estimated in less than a day. The integrated visualization facilities don’t provide quite such a dramatic gain in efficiency, but nevertheless make it much easier to produce simple visualizations of indicator results for diagnosis and policy application.

Another component of the Opus work — our data storage representation — illustrates the ongoing tension between ease of use and transparent access

on the one hand, and efficiency on the other. In UrbanSim 2, input data for the simulation was held in a single large binary file that used a format specific to UrbanSim (and similarly for output data). This was efficient but not particularly transparent. In addition, we regularly changed the format to meet new requirements, introducing versioning problems. For UrbanSim 3, we moved to a SQL database, using MySQL, an efficient, widely-used open source database. The database schema was carefully documented in our web-based documentation. Indicators were defined using SQL queries, which were part of “live” technical documentation [4]. This greatly enhanced transparency, but at a cost in efficiency. In particular, computing a set of indicators could take longer than the run time for the simulation itself. In Opus and UrbanSim 4, we support a variety of data storage formats, including a SQL database, binary files (using the format defined by the numarray package), tab-delimited files, and others. Indicators can now be defined using Opus variable definitions in addition to SQL queries. This is currently a considerably less transparent format, but indicator values can now be computed in minutes rather than hours. We expect to be able to design ways to increase the transparency of this more efficient approach by developing appropriate user interfaces for browsing and computing indicators.

We released Opus and UrbanSim 4 in June 2006. UrbanSim is currently in use in projects in the U.S. for the metropolitan regions around Detroit, El Paso, Honolulu, Houston, and Seattle, and internationally in Amsterdam, Paris, Tel Aviv, and Zurich. Additional projects have been launched in Burlington, Durham, Phoenix, and San Francisco, and internationally in Melbourne, Australia. Many of the current projects, as well as all of the new ones, are now using UrbanSim 4. In July 2006 we hosted the second UrbanSim Users Group meeting, which included 35 participants from across 18 Chapter 22 the U.S., as well as Netherlands and Israel. The primary purpose of this meeting was to assist current users in transitioning to UrbanSim 4.

Uncertainty

We recently started a project to provide a principled statistical analysis of uncertainty in UrbanSim, and to portray these results in a clear and useful way to the users of the system. In this work, we are leveraging a promising technique, Bayesian melding, which combines evidence and uncertainty about the inputs and outputs of a computer model to yield distributions of quantities of policy interest. From this can be derived both best estimates and statements of uncertainty about these quantities. This past year we have had some initial success in employing this technique, applying it to calibrate the model using various sources of uncertainty with an application in Eugene-Springfield, Oregon. These results are reported in a recent journal article [24]. Once the analysis infrastructure is in place, we will extend the Indicator Browser and other interaction and visualization tools to include appropriate depictions of uncertainty. We also intend to extend the uncertainty analysis to address si-

multaneously the problems of developing data for use in modeling (which is currently a very tedious and error-prone process) and analyzing the role of uncertainty in input data on model estimation and on simulation results.

2.4 Applying UrbanSim

An important part of our research agenda has been close collaborations with a number of regional government agencies in applying UrbanSim to their regions. In this subsection we call out the Salt Lake City and Puget Sound applications in particular. These two cases present substantially different aspects, both in terms of characteristics of the region, and of the legal and political context. The Salt Lake City application is described in detail in reference [30], and more briefly in reference [28]. The PSRC application is being studied, particularly with respect to the PSRC/UW collaboration, the decision-making regarding adoption, and their implications for use of the model in practice, as a component of Ruth Förster’s forthcoming Ph.D. dissertation [12].

During 2002–20003, we worked intensively with the Wasatch Front Regional Council in applying UrbanSim to Salt Lake City and the surrounding areas, building on earlier work with the agency. The Greater Wasatch Area, containing 80% of Utah’s population and centered on Salt Lake City, is a rapidly growing metropolitan area, with population predicted to increase by 60% in the year 2020. In order to deal with projected increases in population and travel, Utah officials developed a series of transportation improvement plans, one component being the Legacy Parkway, a controversial four-lane highway extending 14 miles north from Salt Lake City. The highway project precipitated a lawsuit, which ultimately resulted in a settlement that included terms requiring the Wasatch Front Regional Council to test the integration of UrbanSim with their regional travel model system, and if successful, to bring this into their operational use in transportation planning.

The assessment of the integration of UrbanSim with the regional travel model system was launched in 2003 with the formation of a Peer Review Panel, consisting of technical experts in land use and transportation modeling, along with a Management and Policy Committee and a Scenarios Committee. A very tight schedule was specified in the out-of-court settlement, requiring that the entire review be completed by the end of 2003. The Peer Review Panel decided on a validation of the combined UrbanSim – Travel Model System using a series of tests. One test was to model the effects of the existing long-range plan (LRP) in 2030. There were also five other “sensitivity tests,” each of which involved making a simulated major change to the adopted plan and assessing the results, for example, removing a major highway link included in the LRP, removing a major transit link, or adding a significant land use policy such as an urban growth boundary. The overall evaluation focused on issues of model validity and usability.

The Peer Review Panel concluded that UrbanSim produced credible results for tests involving large changes (e.g. the urban growth boundary test),

but not for relatively smaller changes (e.g. removing a major highway link). It was during this process that we discovered the dispersion problem mentioned earlier. Their summary assessment supported the implementation and application of UrbanSim by the WFRC, with the understanding that important refinements and improvements were needed. Subsequently, during 2004, staff of the WFRC made additional efforts to improve the data and the model specification, and the agency decided to move the model into operational use.

Another important collaboration has been with the Puget Sound Regional Council (PSRC), to apply UrbanSim to that region (which includes Seattle and other surrounding cities). In this project, a collaborative agreement was developed between the PSRC and the Center for Urban Simulation and Policy Analysis (CUSPA), to develop the database and apply UrbanSim to the Central Puget Sound region containing Seattle. A Technical Advisory Committee, consisting of planners and analysts from cities and counties in the region, was engaged with the process to review the development of the data and model, and to provide refinements to the data and feedback on the model development. The process was managed closely by two staff members of the PSRC, and during the project two PSRC staff members were hired by CUSPA for the intensive database development effort, and then re-hired by PSRC afterwards. The project was funded by the PSRC and structured in annual contracts with clear work scope for each period. The first two years were essentially spent on database development and refinement, and the third on model estimation and testing.

During this period, as we learned from both the Salt Lake City project and the Puget Sound project about several problems in the model specification and software implementation in UrbanSim 3, we confronted a difficult choice of whether to attempt to resolve these problems in the production system (UrbanSim 3 in Java), or invest all of our effort on completing more quickly the conversion to the more modular Python implementation in OPUS and UrbanSim 4. We had already learned that we could more readily solve modeling problems using the incomplete Python version that we had been unable to address in the UrbanSim 3 code base (mainly due to the complexity of debugging the code). At the same time, there were considerable risks to attempting to rapidly complete and test the conversion to Python without putting the PSRC project well behind schedule. Ultimately, after extensive consultation with the PSRC, we made the decision to freeze further investment in the UrbanSim 3 code, other than minor maintenance, and to put all our effort into the new platform. Since that decision, we have fully implemented the PSRC model application using the new UrbanSim 4 code in OPUS, and have done extensive testing with the system. It may have added as much as a year to the project schedule over what had been expected, but it is not clear that the schedule would have been any earlier if we had continued working on the UrbanSim 3 platform to attempt to resolve the problems we faced.

PSRC staff worked with CUSPA to develop criteria for evaluating the model results and to determine when it would be ready to put into produc-

tion use. Unfortunately, longitudinal data were not available to undertake a historical validation as was done in Eugene-Springfield [27], so the focus of the evaluation was shifted to sensitivity analysis and comparison with previous results. The PSRC has used the DRAM/EMPAL model to prepare land use allocations for use in its transportation planning process for many years, and though they have numerous concerns about it, they have managed to find workable solutions by overriding the model results based on local review procedures using their Technical Advisory Committee and a Regional Technical Forum. The assessment of UrbanSim involved comparing the predictions from 2000 (the base year) to 2030 using UrbanSim with the heavily reviewed and adjusted predictions of DRAM/EMPAL. Since the two models use different levels of geography (DRAM/EMPAL uses approximately 200 Forecast Analysis Zones and UrbanSim uses approximately 800,000 grid cells), UrbanSim data were aggregated to levels that could be compared to the previous results. The initial focus was on population and employment predictions, since these are the data used in the travel model system. A set of indicators was selected for use in diagnosing the performance of the model and to identify issues and problems, in a process of negotiation among the participants.

Based on this, the UW/PSRC team classified the issues into critical, significant, and cosmetic. (Critical issues were ones that had to be resolved before the PSRC would be comfortable placing the model system into operational use. Significant issues were ones that PSRC regarded as important, but would not block putting the system into operational use.)

In addition to comparing the predictions of UrbanSim with previously adopted forecasts, a series of sensitivity tests were conducted to determine if the results from the model were sensitive in the direction and magnitude experts would expect when an input was changed significantly. For these tests, we compared results for the baseline scenario, which was based on currently adopted land use policies and the adopted Regional Transportation Plan, with results for scenarios that included doubling highway capacity in Snohomish County on the north end of the study area, relaxing the Urban Growth Boundary, and constraining development capacity in King County (the central county in the study area). These tests were selected to probe both the scientific robustness of the model and its sensitivity to policies of interest in the region. This can thus be seen as addressing both the accuracy and policy sensitivity aspects of model validity, as discussed in the taxonomy presented in Section 1.2.

The land use policies showed the expected effects, but the highway scenario showed less sensitivity than was expected, in the sense that it did not produce much redistribution of population and employment into Snohomish County. This was consistent with a pattern in the comparison of the UrbanSim baseline scenario against previously adopted forecasts, which showed considerably lower growth rates in Snohomish County than the earlier results did. These remain on the agenda for further examination to determine whether these results are plausible or are an artifact of the model specification or input data.

In spite of these remaining issues, the PSRC is moving ahead with the process of bringing UrbanSim into operational use, and is preparing for a final phase of work to do so during 2006-07.

In addition to the Salt Lake City and Puget Sound applications, we have worked with other agencies in applying UrbanSim in the urban areas around Detroit, Eugene, Honolulu, and Houston. There have also been research and pilot applications in Amsterdam, Paris, Phoenix, Tel Aviv, and Zurich. We have also been working actively to form a user community. The first UrbanSim Users Group Meeting was in San Antonio, Texas, in January 2005. This meeting attracted some 30 participants from Metropolitan Planning Organizations around the country, a number of academic researchers, and one participant from Europe. The second UrbanSim Users Group meeting will be held in July 2006, with one important goal being to help existing UrbanSim 3 users to migrate to UrbanSim 4. We hope to increasingly engage model users and other research groups in the collaborative development of Opus and of UrbanSim itself, by extensions and refinements to the current model system, and the addition of new tools and models.

3 Conclusion and Future Work

In this chapter we have presented an introduction to the domain of urban modeling and to some of the uses and controversies around employing these models to inform public decision-making, including a taxonomy of refinements to urban models and to the process of applying them. We then presented a case study of the UrbanSim model system, including principal areas of research and some applications to planning activities in different regions. This domain represents a significant opportunity for digital government research: hard technical problems, unmet demand from government users, and important issues around supporting a more democratic planning process.

There is work that remains to be done. Most importantly, our goal of producing a system that is in routine and widespread use in informing the planning process is not yet achieved. UrbanSim is being transitioned to operational use in a number of regions, and there are a fair number of additional research applications of the system. However, it is not yet in routine policy use. Beyond that, the development of the Opus platform should enable a rich set of collaborations among researchers world-wide, including the development of open-source travel models and environmental models closely integrated with the land use models in UrbanSim. Finally, we have touched on two other major open areas of ongoing research: first, increasing access to the results of modeling for a wide range of stakeholders, and ultimately to simulating additional alternatives; and second, providing a principled modeling of uncertainty in land use and transportation simulations.

4 Additional Information

4.1 Acknowledgments

We would like to thank all of the UrbanSim research team members and collaborators. This research has been funded in part by grants from the National Science Foundation (EIA-0121326 and IIS-0534094), in part by a grant from the Environmental Protection Agency (R831837), in part by a partnership with the Puget Sound Regional Council, and in part by gifts from IBM and Google.

4.2 Suggested Readings

- Paul Waddell and Gudmundur F. Ulfarsson, “Introduction to Urban Simulation: Design and Development of Operational Models,” *Handbook of Transport, Volume 5: Transport Geography and Spatial Systems*, P. Stopher et al., eds., Pergamon Press, 2004. Available from www.urbansim.org/papers. An introduction to urban simulation (interpreted broadly to mean operational models that attempt to represent dynamic processes and interactions of urban development and transportation).
- Paul Waddell, “UrbanSim: Modeling Urban Development for Land Use, Transportation, and Environmental Planning,” *Journal of the American Planning Association*, Vol. 68 No. 3, Summer 2002, pages 297–314. An early but accessible overview of the UrbanSim system.
- Paul Waddell and Alan Borning, “A Case Study in Digital Government: Developing and Applying UrbanSim, a System for Simulating Urban Land Use, Transportation, and Environmental Impacts,” *Social Science Computer Review*, Vol. 22 No. 1, February 2004, pages 37–51. As the title suggests, presents UrbanSim as a case study in digital government research. Emphasizes the various tensions between conducting university research and applying it to a regional government application in Salt Lake City in a highly visible, contested domain.
- Batya Friedman, Peter H. Kahn Jr., and Alan Borning, “Value Sensitive Design and Information Systems: Three Case Studies,” in *Human-Computer Interaction and Management Information Systems: Foundations*, Ping Zhang and Dennis Galletta, eds., M.E. Sharpe, Armonk, NY, 2006. A description of the Value Sensitive Design theory and methodology, including an UrbanSim case study.

4.3 Online Resources

- <http://www.urbansim.org> The UrbanSim website, including papers, downloads, and other information.
- <http://www.urbansimcommons.org> The UrbanSim Commons is a companion site, designed and maintained by UrbanSim users.

- <http://www.ischool.washington.edu/vsd> The home page for the Value Sensitive Design project.

4.4 Questions for Discussion

Here are a number of questions that would be suitable for classroom discussion (as well as future research).

- Consider D. B. Lee’s 1973 critique “Requiem for Large-Scale Models” [18]. Which of the problems raised by Lee have been addressed by current models, and which are still open problems? Of the open problems, which will it be feasible to address in the next decade?
- What issues arise in applying UrbanSim to urban regions in countries other than the United States, including both developed and developing countries? Consider issues of alternate land use laws and ownership, transportation patterns, data availability, and others.
- There are huge uncertainties about some of the exogenous assumptions used by UrbanSim, including the future price of oil, the macroeconomy, possible technological shifts in transportation or telecommuting, and others. How should these uncertainties best be communicated to users of UrbanSim? What are appropriate controls for allowing them to change these assumptions?
- Suppose UrbanSim were being designed for use by authoritarian governments rather than in a democratic context. Are there things that the designers should do differently? If so, what?

References

1. K. Beck. *Extreme Programming Explained: Embrace Change*. Addison-Wesley, Reading, Mass., 2000.
2. K. Beck. *Test-Driven Development – By Example*. Addison-Wesley, Reading, Mass., 2003.
3. E. Beimborn, R. Kennedy, and W. Schaefer. Inside the blackbox: Making transportation models work for livable communities. Citizens for a Better Environment and Environmental Defense Fund, 1996. Available from Environmental Defense Fund, Washington D.C., <http://www.edf.org>.
4. A. Borning, B. Friedman, J. Davis, and P. Lin. Informing public deliberation: Value sensitive design of indicators for a large-scale urban simulation. In *Proc. 9th European Conference on Computer-Supported Cooperative Work*, Paris, Sept. 2005.
5. A. B. Brush and A. Borning. ‘Today’ messages: Lightweight support for small group awareness via email. In *Hawaii International Conference on System Sciences (HICSS 38)*, Jan. 2005.
6. J. Davis. Household indicators: Design to inform and engage citizens. In *CHI 2006 Work-in-Progress Papers*. ACM Press, Apr. 2006.

7. J. Davis. *Value Sensitive Design of Interactions with UrbanSim Indicators*. PhD thesis, Dept. of Computer Science & Engineering, University of Washington, 2006.
8. T. de la Barra. *Integrated Land Use and Transportation Modeling: Decision Chains and Hierarchies*. Cambridge University Press, 1995.
9. R. Dowling, R. Ireson, A. Skabardonis, D. Gillen, and P. Stopher. Predicting air quality effects of traffic-flow improvements: Final report and user's guide. Technical Report 535, National Cooperative Highway Research Program, Transportation Research Board, National Research Council, 2005.
10. A. Downs. *Still Stuck in Traffic: Coping with Peak-hour Congestion*. The Brookings Institution, Washington, D.C., 2004.
11. M. Echenique, A. Flowerdew, J. Hunt, T. Mayo, I. Skidmore, and D. Simmonds. The MEPLAN models of Bilbao, Leeds and Dortmund. *Transport Reviews*, 10(4):309–322, 1990.
12. R. Förster. *Bridging the gap between scientifically-based information and its application in planning for sustainable regional development: The role of computer models*. PhD thesis, Department of Environmental Engineering, ETH Zurich, 2007. In preparation.
13. B. Freeman-Benson and A. Borning. YP and urban simulation: Applying an agile programming methodology in a politically tempestuous domain. In *Proceedings of the 2003 Agile Development Conference*, Salt Lake City, Utah, June 2003. Available at <http://www.urbansim.org/papers>.
14. B. Friedman, P. H. Kahn Jr., and A. Borning. Value Sensitive Design and information systems: Three case studies. In *Human-Computer Interaction and Management Information Systems: Foundations*. M.E. Sharpe, Armonk, NY, 2006.
15. B. Friedman and H. Nissenbaum. Bias in computer systems. *ACM Transactions on Information Systems*, 14(3):330–347, 1996.
16. M. Garret and M. Wachs, editors. *Transportation Planning on Trial: The Clean Air Act and Travel Forecasting*. Sage Publications, Thousand Oaks, CA, 1996.
17. A. Hunt and D. Thomas. *Pragmatic Unit Testing*. The Pragmatic Programmers, LLC, 2003.
18. D. B. Lee. Requiem for large-scale models. *Journal of the American Institute of Planners*, 39:163–178, 1973.
19. I. Lowry. A model of metropolis. Technical Report RM-4035-RC, Rand Corporation, Santa Monica, CA, 1964.
20. E. Miller, D. Kriger, and J. Hunt. Integrated urban models for simulation of transit and land use policies. Final Report, Project H-12. TCRP Web Document 9, Transit Cooperative Highway Research Project, National Academy of Sciences: Washington, DC. <http://books.nap.edu/books/tcr009/html>, 1999.
21. R. E. Noonan and R. H. Prosl. Unit testing frameworks. In *SIGCSE '02: Proceedings of the 33rd SIGCSE Technical Symposium on Computer Science Education*, pages 232–236, New York, 2002. ACM Press.
22. S. Putman. *Integrated Urban Models: Policy Analysis of Transportation and Land Use*. Pion, London, 1983.
23. H. Ševčíková, A. Borning, D. Socha, and W.-G. Bleek. Automated testing of stochastic systems: A statistically grounded approach. In *ISSTA '06: Proceedings of the 2006 International Symposium on Software Testing and Analysis*, pages 215–224, New York, 2006. ACM Press.

24. H. Ševčíková, A. Raftery, and P. Waddell. Assessing uncertainty in urban simulations using Bayesian melding. *Transportation Research B*, 2006. Accepted subject to modest revision - draft available from www.urbansim.org/papers/BMinUrbansim.pdf.
25. F. Southworth. A technical review of urban land use-transportation models as tools for evaluating vehicle reduction strategies. U.S. Department of Energy, 1995.
26. P. Waddell. A behavioral simulation model for metropolitan policy analysis and planning: Residential location and housing market components of UrbanSim. *Environment and Planning B: Planning and Design 2000*, 27(2):247–263, 2000.
27. P. Waddell. UrbanSim: Modeling urban development for land use, transportation, and environmental planning. *Journal of the American Planning Association*, 68(3):297–314, Summer 2002.
28. P. Waddell and A. Borning. A case study in digital government: Developing and applying UrbanSim, a system for simulating urban land use, transportation, and environmental impacts. *Social Science Computer Review*, 22(1):37–51, 2004.
29. P. Waddell, H. Ševčíková, D. Socha, E. Miller, and K. Nagel. Opus: An open platform for urban simulation. Presented at the Computers in Urban Planning and Urban Management Conference, London, June 2005. Available from www.urbansim.org/papers.
30. P. Waddell, G. Ulfarsson, J. Franklin, and J. Lobb. Incorporating land use in metropolitan transportation planning. *Transportation Research Part A: Policy and Practice*, 2006. In press – draft available from www.urbansim.org/papers.
31. P. Waddell and G. F. Ulfarsson. Introduction to urban simulation: Design and development of operational models. In P. Stopher, K. Button, K. Haynes, and D. Hensher, editors, *Handbook of Transport, Volume 5: Transport Geography and Spatial Systems*. Pergamon Press, 2004.