THE JOURNAL OF TRANSPORT AND LAND USE http://jtlu.org
VOL. 11 NO. 1 [2018] pp. 19–30

Oregon's Transportation and Land Use Model Integration Program: A retrospective

Rick Donnelly
WSP

William J. Upton
Oregon Department of Transportation

Becky Knudson
Oregon Department of Transportation

Abstract: An ambitious and innovative integrated land use-transport modeling system has been developed in Oregon over the past two decades. This work, completed under the Transportation and Land Use Model Integration Program (TLUMIP), included the development of two generations of models and the data required to build and use them and spawned the development of two others that have continued independently. An outreach program and collaborative development of freight data and forecasts were also included, as well as system testing and applications. A brief description of the motivation behind TLUMIP and the resulting modeling systems are presented. Perhaps more interesting is the story behind the models, describing several major model design, institutional, and methodology issues that were overcome. Using an integrated model in practice also entailed addressing a wider range of analytical requirements and stakeholder expectations about usability, accuracy, and extensibility than typically considered in academic pursuits. The key lessons learned through development and use of the models are discussed, with the hope that they will inform the development of similar large-scale modeling systems.

Article History:
Received: April 15, 2017
Accepted: August 4, 2017
Available online: January 3, 2018

1 Introduction

The Oregon Department of Transportation (ODOT) developed a set of integrated economic-land use-transportation models over the past two decades to meet evolving analytical needs at the urban and statewide levels. The work was completed in several multi-year stages as part of their Transportation and Land Use Model Integration Program (TLUMIP), and included extensive model and data design and development, interim model applications, peer review, and program assessments. Most planning agencies have neither the need nor resources to develop as sophisticated a modeling system as Oregon. However, many of the important lessons learned apply to other large-scale model development projects, and are shared in the hope that they will enrich those efforts.

This paper is organized into six sections. The motivations for building an integrated modeling system are described in Section 2, while the design and architecture of first and second generation

Copyright 2018 Rick Donnelly, William J. Upton, and Becky Knudson.
doi: 10.5198/jtlu.2018.1210
ISSN: 1938-7849 | Licensed under the Creative Commons Attribution – NonCommercial License 3.0.

The Journal of Transport and Land Use is the official journal of the World Society for Transport and Land Use (WSTLUR) and is published and sponsored by the University of Minnesota Center for Transportation Studies.

systems are described in Section 3. The model implementation process is discussed in Section 4, followed by a summary of model assessment and acceptance testing in Section 5. These sections provide enough background to appreciate the lessons learned building and using the system in Section 6. The key lessons learned and likely next steps are summarized in the conclusion.

2 Motivations

Oregon has pushed the envelope in land use planning for several decades, particularly in the Portland region (Knaap and Nelson 1992; Walker and Hurley 2011). Many of the early motivations for TLUMIP revolved around quantifying the impacts of various land use strategies and synergies, as shown in the top part of Figure 1. Moreover, an effort was made at the outset of TLUMIP to identify the likely constituents of the model and the issues they were facing. These included legislative staff, decision-makers within the ODOT, and metropolitan planning organizations. Their needs dictated what Oregon wanted to accomplish, rather than simply desiring to have "better models" or advancing the state of the art. One key driver was the need to assess the land use and transportation effects of urban growth boundary legislation designed to protect Oregon farmland. Oregon would not likely have invested in such a sophisticated modeling system absent such requirements.

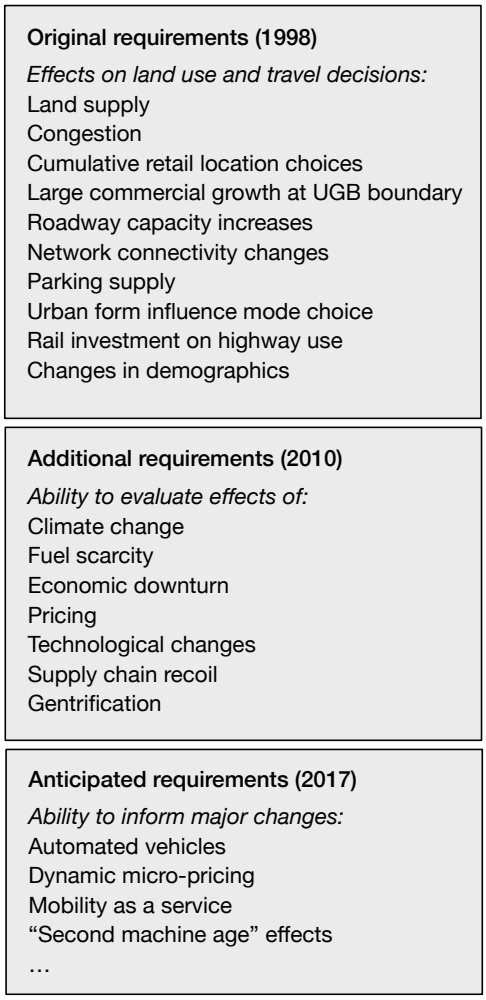

Figure 1: Analytical issues the modeling system must be capable of informing

At the outset, the interactions between land use and transportation were of high priority, and heavily influenced the design of the modeling system. A little over a decade later the policy focus had

shifted considerably. Applications of the first-generation model focused more on the interactions between the economy, jobs, and transportation than land use. Strategies for reducing greenhouse gas emissions had become an important topic, as did pricing and effects of the 2008-10 economic downturn.

At this writing, most of the previous issues remain, but several new ones have arisen. The ability to understand dynamic pricing remains a priority. Perhaps even more pressing is the need for information about the likely effects of connected and automated vehicles, which appear closer to reality than just a few years ago. Related to that are mobility services such as Uber and Lyft, and many others likely to enter this space.

The ability to model freight, its economic linkages and relationship to jobs, and visualize truck flow patterns turned out to be as important, and perhaps more so than modeling person travel flows. Most of the major applications of the model to date have involved freight, either wholly or in part.

Taken together, these requirements dictated the need for a highly flexible and scalable modeling framework that can quickly adapt to new needs. Oregon has invested in strategic visioning models and agile development to help meet these needs, and continues to seek new approaches that will enable them to do so better. The reminder that applications should dictate design rather than the other way around is a lesson worth repeatedly learning and has been key to sustained support for their program.

3 Model architecture

It was envisioned that at least two modeling systems would be created. The first generation system was to provide a proof of concept and interim operational capabilities. The first attempt at a dynamic disequilibrium model at the urban level was also undertaken. Based on those experiences a new modeling system was designed to include all the desired functionality, incorporate activity-based travel models, and enable each module to operate at the level of behavioral, temporal, and spatial resolution most appropriate for it.

3.1 First generation system

The first phase of TLUMIP focused upon building an interim model capable of meeting ODOT needs while the longer second generation model development was underway. The TRANUS platform (de la Barra 2005) was selected because of its well-understood foundations, graphical user interface, and previous quick implementations. Work also began in parallel on the design and original development of an urban model in Eugene-Springfield that evolved into UrbanSim (Waddell 2002).

A considerable amount of effort was devoted to assembling the data required to build both models. The goal was to create the initial models using secondary data sources, both to reduce their cost and implement them quickly. Oregon did not have a prior statewide model, necessitating the development of zone systems and networks from scratch, as well as data unique to land use-transport models. The latter included land pricing data, parcel and tract level databases, and economic input-output flows and coefficients. Surveys were conducted at truck weigh stations to better understand the relationship between tonnage by commodity and truckload equivalents.

The required socioeconomic and land use data were unusually detailed for a statewide model. They included 18 household types stratified by household size and income, 52 employment categories, 14 occupation categories, six residential floorspace types, 14 non-residential floorspace types, and 41 commodity classifications. There are few operational models known that attempt to operate at this level of detail.

A statewide household travel survey of 14,000 households was carried out in 1994-96. It was designed to support the development of small urban and MPO travel forecasting models, as well as the statewide travel models. The Oregon Household Activity Survey (OHAS) of around 18,000 households was subsequently conducted in 2009-11. Both surveys included households across the state,

with about 6,500 of them within the Portland-Vancouver metropolitan area as part of the OHAS. Oregon has otherwise relied upon secondary public or government data sources, or have purchased third-party data, to support their modeling needs. Like other states, they have borrowed data and model parameters from other places, to include long-distance travel survey data from Ohio's 2001-03 statewide household travel survey. The FHWA Freight Analysis Framework (FAF) data were used to depict commodity flows entering or leaving the state. These data were fused with Carload Waybill Sample data to better understand rail and intermodal traffic patterns.

3.2 Second generation model

Work began a decade ago on the current statewide model, after completing several applications with the first-generation prototypes. It was decided to focus on a single framework, rather than separate models at the urban and statewide scales. This was motivated in part by the anticipated use cases, as well as the practical necessity of focused the limited resources and talent available to the program. A modular approach was adopted, where each component would interact with others through data files and messaging. This was thought to facilitate the development of each module by different teams, as well as testing different approaches for each, without disrupting the work of others.

The resulting modeling platform is a hybrid, in that some models are aggregate equilibrium models, while others are microsimulated. The modeled area covers the state of Oregon and a halo of roughly 50 miles covering counties in the neighboring states of Washington, Idaho, Nevada and California. A multi-scale system is used to represent geography at scales appropriate for each part of the modeling system. The three scales currently used include 2,950 traffic analysis zones within the modeled area (alpha zones), 518 aggregations of them for activity allocation (beta zones), and six external markets.

A flow chart of the second generation statewide integrated model (SWIM2) is shown in Figure 2. The economic, land use and transportation models are tightly integrated, with the output of each model driving one or more downstream models as well as its own evolution. For example, travel costs and disutilities from the traffic assignment model are fed back to the land use modules in the subsequent simulation period, and land use changes feed back into the economic model. The modules currently implemented include:

- The New Economics and Demographics (NED) module determines model-wide production activity levels, employment, and imports and exports based on official Oregon state forecasts. This model is an aggregate formulation, and capable of producing different economic futures.
- The Synthetic Population Generator (SPG) module samples household and person demographic attributes (SPG1), and subsequently assigns the household to an alpha zone (SPG2).
- The Aggregate Land Development (ALD) module allocates model-wide land development decisions to alpha zones considering floorspace prices and vacancy rates.
- The Activity Allocations (AA) module uses economic input-output theory to locate production and consumption of land (floorspace), labor, and commodities in places that optimizes their accessibilities to inputs and markets. This results in estimates of both labor and commodities, and sets prices in exchange zones at the levels required to clear those markets.
- The Person Travel (PT) module generates activity-based person trips for each person in the synthetic population during a typical weekday and assigns a workplace by alpha zone. It includes both short and long-distance travel flows, which are combined with commercial vehicle flows in a trip list for traffic assignment.
- The Commercial Transport (CT) module maps interregional commodity flows into daily truckload equivalents, both within the modeled area and between it and external markets. It also generates internal (to the modeled area) truck tours using a microsimulation framework, to include trips made through distribution centers and intermodal drayage.

Macroscopic roadway and transit network assignments are carried out using a commercial transportation planning package for four periods of the day. Data and runtime parameters are passed be-

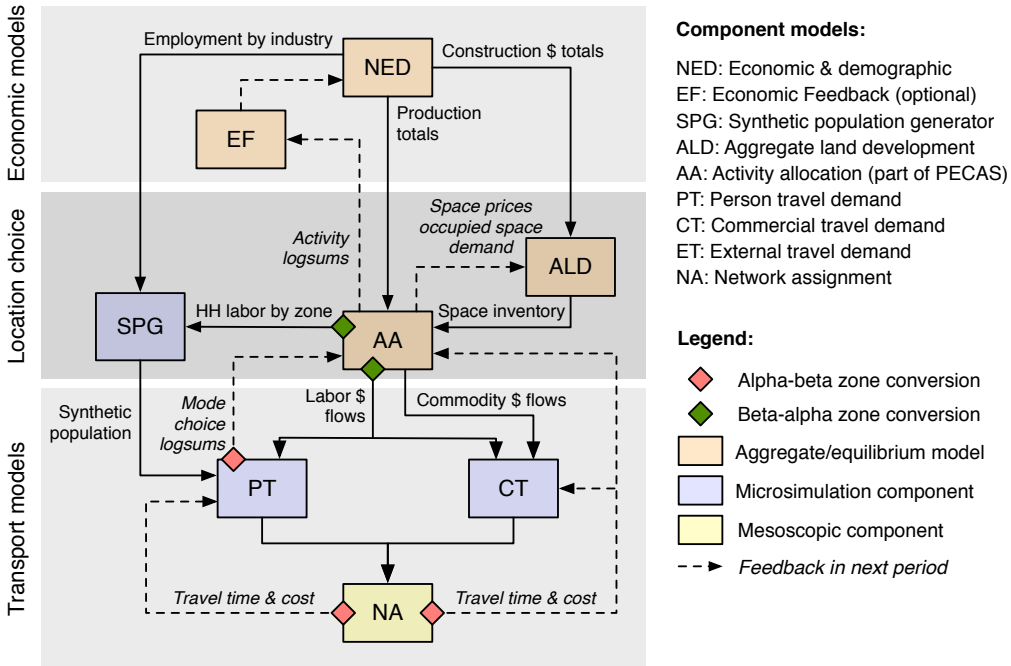

Figure 2: Schematic representation of the SWIM2 model components and key interactions

tween modules in data tables, mostly in text format to improve readability and portability. Log files and metadata are saved after to each simulation to describe the runtime properties and outcomes.

Still under development is an economic feedback (EF) module, which is a simplified dynamic feedback that adjusts the NED modules fixed model-wide economic forecast, considering the statewide composite location utilities by industry from the AA module. A detailed description of each of these modules and their development can be found in Donnelly (2017).

The economic and land use modules run in annual time steps through the simulation, while the transport models are typically run every third year to reduce overall run time for longer simulations. A complete model run, covering 30 years in annual steps, takes approximately eight hours per year using a cluster of six quad-core computers, making extensive use of process threading and distributed execution. Thus, a simulation spanning 35 years takes several days to complete.

4 Model implementation and applications

ODOT invested $9.1 million in TLUMIP over the course of the program. Slightly less than a third of the budget was devoted to the development of the first and second generation models. A larger amount was spent on data development, to include travel surveys, acquisition of third-party data (including data from other government agencies), and a commodity flow survey jointly undertaken with Portland Metro and the Port of Portland. Several applications of the model were carried out, both for model acceptance testing and the major studies described below. In addition, two ODOT staff members have been involved for about half of their time in model design, development, and implementation over the entire period that TLUMIP has been active.

4.1 System implementation and delivery

The SWIM2 system was implemented and tested in a staged manner. The preparation and testing of the software code, preparation of validation and calibration data and targets, and estimation and calibration of initial parameters were completed first. The cross-sectional calibration and validation of

the full model system (i.e., modules working together) were then carried out. Finally, the system was run through time, subjecting it different policies and stressors.

The design of the activity allocation (AA) module changed over the course of developing the SWIM2 platform. It started out as a microsimulation model, but difficulties obtaining data about firm behavior and transitions drove the design back towards an aggregate equilibrium formulation inspired, in part, by the MEPLAN model (Echenique *et al.* 2007). It later evolved into the Production-Exchange-Consumption Allocation System (PECAS), development of which was undertaken in several places in addition to Oregon (Hunt and Abraham 2015). As such, it is fair to describe TLUMIP as an incubator of innovative ideas to model land use changes, having already supported the initial development of UrbanSim.

The SWIM2 development started with traditional multi-year contracting, with a comprehensive model design at the outset. This was the typical approach to large-scale model development in the late 1990s, when TLUMIP was launched. Later experiences, and failures in some cases, led to the adoption of the agile development practices becoming more commonly used today. It is a concept from software development that proposes to start with, "the simplest thing that can possibly work." The components are then incrementally developed based on user and client feedback, continual review of requirements, and performance testing. Oregon has benefited from using this approach, despite adopting it while SWIM2 was in development. The questions asked by ODOT stakeholders changed while the program was in motion, and the team shifted development and implementation priorities to accommodate them. This would have required contract amendments and considerable delay had the work been traditionally structured, with deliverables and tasks defined up front.

4.2 Major applications

The first and second-generation models have been used in a variety of studies, most of which were high-priority analyses of planning, investment, and policy options being considered by ODOT at the time. Arguably the most significant application was the ODOT Bridge Limitations Study in 2004. A design flaw had been discovered that increased the risk of premature cracking in several dozen bridges across the state. The potential repair costs substantially exceeded the agency's budget. However, not addressing them would have resulted in weight restrictions or closures. ODOT's modeling staff proposed to the Department's management that the problem should be couched as a community connectivity and economic impact analysis, rather than engineering design issue. This shifted the message to impacts on jobs and gross state product, which the legislature better related to than engineering needs.

The statewide model was used to test a variety of prioritization packages and effects of bridge limitations. A $4.3 billion program was approved to address the worst cases first. The simulated solution was half the cost and avoided 90 percent of the economic impacts associated with the no investment scenario. Moreover, it resulted in prioritizing repairs for communities most affected by loss of connectivity and jobs, rather than the initial assumption that Interstate highway and high-volume state bridges should be addressed first.

The model has also been used in a freight bottleneck study, and is currently being used as part of the Department's "rough roads" assessment (Oregon DOT 2017). The latter focuses on the implications of inadequate maintenance of existing roadway infrastructure across the state. It has the potential to be as important to the Department as the bridge limitations study, and equally well suited for analyses of economic and social impacts.

These early applications were supported ODOT's mission, built support for the TLUMIP, and demonstrated the range of issues it could usefully inform. However, the insights it provided to the development team were likely larger. Applications are often seen as a distraction to model developers focused on creating new tools. They provides invaluable insight into operational requirements and usability issues, and played as large a role in redefining system requirements as conceptual, data, or methodological considerations. The need to adapt several times in response to experiences in appli-

cation led the team to a more agile development approach, discussed further among the major lessons learned.

5 Model assessment and acceptance testing

Bringing the TLUMIP models into practice at ODOT required demonstrating their accuracy and efficacy. As Waddell (2011) aptly notes, using integrated models for policy and planning analyses brings with it higher expectations than the conceptual considerations of researchers. This was challenging, for each component was a complex model in its own right, which only increased as the interactions between them through time were considered. There was scant precedence to inform expectations about model accuracy, or at what levels of spatial, behavioral, or temporal resolution to validate the system at. It was clear that a multiscale approach was required, as each of the primary systems – the economy, land use, and the transport system – evolved at different rates and were measured at different scales (Kok *et al.* 2001).

It became clear early in TLUMIP that an integrated land use-transport model could not achieve the same levels of accuracy that urban travel demand modelers were accustomed to. Those that held such expectations had not considered how travel demand fit into the larger modeling framework, its interactions with other systems, or the much larger number and diversity of calibration and validation targets. Workplace type and location choice, for example, were no longer parts of the travel model, but a function of labor consumption by firms and government in the AA module. A large number of other interactions beyond the scope of urban travel demand models had to be understood before practitioners from that realm grasped their implications for acceptance testing.

The acceptance testing framework shown in Figure 3 was developed in collaboration with a peer review panel. Their recommendations influenced the final design and specification of the models. A three-stage calibration process (S1-S3) was carried out for each component and for the system. The focus of this effort was on cross-sectional calibration for the base year of the model, for which calibration targets could be derived. The system level calibration (S3) proved difficult and time-consuming, for the parameter space of the combined models was vast, and doubts remained whether to focus on cross-sectional behaviors or how they change over time.

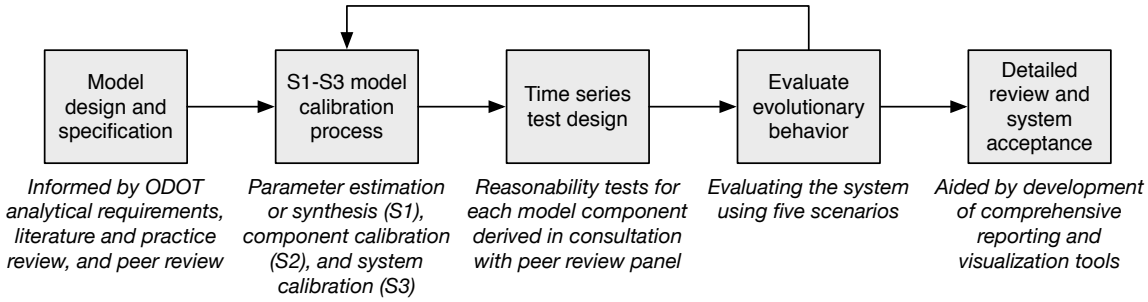

Figure 3: Validation and acceptance testing process used to assess the modeling system

A concerted effort was made to backcast to demonstrate that the evolutionary behavior of the model was reasonable. This effort failed, for many of the required model inputs were not available or of acceptable completeness and quality for past years, especially land consumption and pricing data. This did not reduce confidence in the results, but did necessitate rethinking the approach to validating evolutionary behavior. Five scenarios representing plausible futures were devised to test the model, with the goal of evaluating how the system responded to major changes. These included simulating the effects of different economic growth rates, an increase in multi-family housing stock in county within the Portland metropolitan area, collapse of an Interstate highway bridge, and imposition of

tolls on I-5 (the major north-south Interstate highway through Oregon). ODOT is also using the model in economic studies and to examine the effects of deteriorating highway infrastructure.

The testing of the evolutionary behavior of the model prompted a reassessment of the cross-sectional calibration (S2-S3) results, particularly for the activity allocation (AA) module. The choice of parameter estimates or runtime properties were found to not affect the cross-sectional calibration results in some instances, but did affect how the model performed through time. The need to consider validation from at least three perspectives became clear during this process. The statistical estimation results had to be reconciled with theory, as well as how well the model tracked past trends or expectations about how they might unfold in the future. Brown *et al.* (2005) noted this tension between theoretical and quantitative considerations, which influenced the design of the multi-phase process shown in Figure 3. However, the implications did not become clear until well into the model testing process in Oregon.

The validity of the SWIM system and its suitability for use in policy and planning activities was verified during the final review. However, the subjective nature of reasonability tests, coupled with uncertainties surrounding parameter estimates and the future, suggest that a definitive validation will always remain elusive. The practical implication for ODOT is that validation and reasonability testing will remain an ongoing activity rather than a milestone.

6 Lessons learned

The biggest challenges encountered during TLUMIP revolved around the process of developing robust and reliable implementations, rather than the technical details of the models themselves. There were ample challenges involving the software implementations, early versions of which relied upon distributing computing, and more recent versions on parallel processing. These seem typical of large-scale advanced models, as do the challenges of acquiring the data required for microsimulation modeling. These design and institutional issues unique to TLUMIP were harder to solve, for in most cases there was no optimal approach waiting to be discovered, or success stories to emulate.

6.1 Institutional considerations

The importance of champions within the agency, committed to the program over its lifespan, cannot be over-emphasized. This was identified as a key institutional requirement in earlier research on advanced models (Donnelly *et al.* 2010), but arguably nowhere more important than in multi-stage projects spanning several years. In such cases keeping the team, ODOT managers, and stakeholders all focused on the long-term vision and pay-offs is imperative. This was a larger task than expected, and cannot be substituted by consultants or external advisors. Oregon was fortunate to have two program managers who oversaw the program, and ensured continuity in funding, contracting, and oversight over the life of the program. It is highly unlikely that TLUMIP would have survived without their leadership.

Other champions played a less visible role, but were equally important to the program's success. Managers within ODOT's Transportation Planning Analysis Unit handled the program management and ensured funding over the lifespan of the project, quietly ushered the modeling system into the mainstream of ODOT planning, and ensured that it was applied for important analyses carried out by the Unit. This included turning away projects that were not a good match for the capabilities of SWIM2. Federal funding was used in addition to state funds, with several ODOT Directors making the case for continued funding to the U.S. Department of Transportation. The effectiveness of these champions was enhanced by their positions within the agency, for they controlled the resources required for program success.

An equally important partner was a peer review panel of internationally prominent experts in integrated land use-transport and travel demand modeling. The membership fluctuated between six and

seven members throughout the life of the program, with most members serving for the entirety of the program. Assembling them at the outset allowed them to influence the program direction and subsequent model designs. They assumed ownership of those recommendations and held the development team accountable for progress on them. Perhaps more importantly, the panelists shared a wealth of knowledge and experience, from both successes and failures, that would be difficult to assemble on the development team.

6.2 Design issues

The development priorities for TLUMIP were guided by current and anticipated analytical requirements rather than desired model features. A progression of key design decisions drove this thinking, as illustrated in Figure 4. That is, the context of the analyses ("stories to be told with the model") and forecasting requirements drove the TLUMIP agenda, rather than a desire to advance the state of the art. The differences were subtle, but important, for the context and forecasting requirements were unique, whereas a range of tools and data might plausibly have satisfied those requirements. This kept the focus upon meeting analytical needs instead of the most elegant or conceptually appealing model formulation. Maintaining this focus resulted in the team asking how any proposed revisions or additions to the model design better met the requirements. This did not always prevent "feature creep" and additional model complexity, but it certainly reduced it.

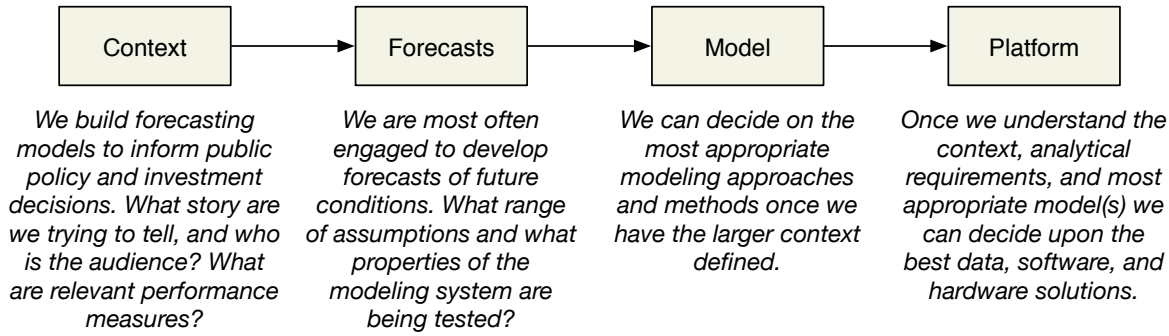

Figure 4: The context and forecasting requirements dictate model design and platforms

Both the first and second generation models began with an extensive model design process, followed by detailed work plan. This "big design up front" (BDUF) followed by several years of development is typical, but becomes more inflexible as the project progresses. This approach worked well in the first generation, where the timeline was relatively short and products well defined. The second generation included the original development of all model components, which carried much higher technical and schedule risk, uncertainties about model formulation and estimation results, and increased complexity.

The team adapted the principles of agile software development (Martin 2003), which emphasizes starting with "the simplest thing that can possibly work," and then making incremental improvements in short sprints of 4-12 weeks' duration. A new operational version of the model is produced at the end of each sprint, accompanied by formal team review of the deliverables and next steps.

An ideal progression of agile model development is shown in Figure 5, where each column represents a specific module or functionality, and the rows show several interim versions of the model. Unfortunately, work on the second generation model was already underway when agile development gained a following, so had to be gradually worked into the program. Even then the uptake was mixed, for adjusting thinking and work habits proved more challenging in some cases than adjusting the work program.

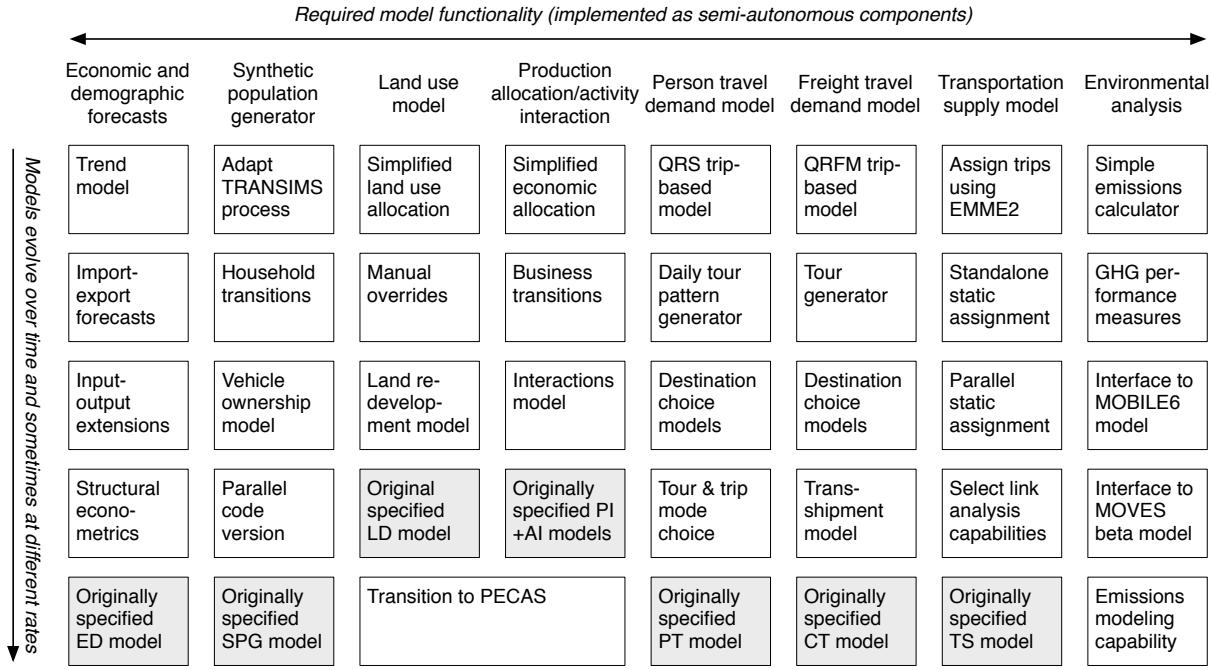

Figure 5: An agile development approach to facilitate incremental development of the system

The complexity of the model increased as its development proceeded, which was perhaps inevitable. Each of the SWIM modules encompasses functionality often found in standalone packages. Advanced person and commercial vehicle travel models are both included in SWIM, for example, which are often complex and complicated when applied in isolation. Many model parameters and run-time properties were required for each module, which led to what was called "parameter storm." The parameter space in such modeling systems is vast, frustrating attempts to fully understand the highly complex interactions of parameters that appear simple when viewed in isolation. The uncertainties arising from such interactions are compounded rather than attenuated over time (Næss *et al.* 2015; Zhao and Kockelman 2002). The challenges posed by parameter storm in SWIM are still being tackled in practice.

These design changes all worked together to enable ODOT to meet the biggest design challenges – continually changing analytical requirements and the need to incorporate new data and methods. The questions that policy-makers faced, summarized in Figure 1, changed faster over the lifetime of the program than model development could keep up. This suggests the need for a more flexible and quickly adaptive modeling framework than typically found in practice, for the challenge remains today.

6.3 Methodological challenges

The original specification of the SWIM system consisted of several modules whose parameters were estimated and selected using econometric approaches. Estimation datasets were assembled for each behavioral model, and contained variables and measurements from both chosen and unchosen alternatives. Discrete choice models were formally estimated using these data, and calibrated to observed conditions. This process is widely used in transport modeling (Hess *et al.* 2007). Many of these models had mediocre fits despite incorporating dummy variables and sometimes large numbers of constants. It was also challenging to explain the interaction of the variables to policymakers. These shortcomings, coupled with the effort and resources required to estimate and calibrate complex formulations, led the team to explore alternative approaches.

Many of the behavioral models retain their original econometric basis, while other approaches were used in others. The aggregate land development model (ALD in Figure 2) is combined rule-based heuristics with linear regression models, while the commercial travel (CT) model combines aggregate allocation methods and Monte Carlo simulation. They are more theory-based, with emphasis upon making explicit assumptions about how behavior is expected to change over time. In a sense, the emphasis changed from replication of cross-sectional behavior to more intuitive formulations, from which a story could be told of how and why behavior changes over time. Moreover, the alternative approaches have been designed with the goal of reducing parameter storm, seen as a roadblock for cost-effective implementation of the model.

7 Conclusion

TLUMIP as a specific program has recently ended, but integrated modeling has become mainstream in Oregon as a result. It represents one of the most ambitious land use-transport models ever developed. Over a dozen people made major contributions to the system and its application, and that many more helped in significant ways. The unwavering support of ODOT, from its Transportation Planning Analysis Unit managers all the way to the Director, as well as federal support and substantial contributions of the peer review panel, were essential to the success of TLUMIP. Their collective contributions and experience formed the basis for this retrospective, for most of the lessons worth learning from TLUMIP came from overcoming challenges rather than brilliant insights.

The SWIM system is as complicated and complex as it is powerful and comprehensive. Perhaps these opposing properties cannot be avoided (Acheampong and Silva 2015; Miller *et al.* 2011), given the diversity and depth of the systems being simulated. This, coupled with the analytical issues the model needed to inform, shaped the design of the system. Indeed, the models evolved during development to accommodate changes in requirements, straining the viability of the traditional approach to developing a single model from a detailed model specification. As a result, the institutional and design issues encountered during TLUMIP were more significant, with more long-term implications, than the methodological, data, or software considerations. The imperative of a champion for the program within ODOT became quickly apparent and remained so throughout the program. The likelihood of success in a program like TLUMIP without the champion seems remote.

The work in Oregon remains a work in progress, and perhaps will always be so. The constantly changing policy environment, the level of questions posed by decision-makers, their zeal for pushing the envelope in land use-transport modeling, and upper management support ensures that Oregon will remain at the forefront of statewide modeling. The investment they have made in their staff is particularly evident, giving them the ability to do many things that many agencies cannot. Pushing the envelope in land use-transport modeling was rewarding, but moving it into practice in the hands of experienced and invested users is the greatest and most enduring achievement of TLUMIP.

References

Acheampong, R. A. and E. A. Silva. 2015. Land use-transport interaction modeling: A review of the literature and future research directions. *Journal of Transport and Land Use*, 8(3):11–38. doi: http://dx.doi.org/10.5198/jtlu.2015.806.

Brown, D. G., S. Page, R. Riolo, M. Zellner, and W. Rand. 2005. Path dependence and the validation of agent-based spatial models of land use. *International Journal of Geographical Information Science*, 19(2):153–174. doi: http://dx.doi.org/10.1080/13658810410001713399.

de la Barra, T. 2005. *Integrated land use and transport modelling: decision chains and hierarchies*. Cambridge, United Kingdom: Cambridge University Press.

Donnelly, R., ed. 2017. *SWIM v2.5 Model Development Report*. Oregon Department of Transportation. URL https://github.com/tlumip/model-dev-report/swim25mdr.pdf.

Donnelly, R., G. Erhardt, R. Moeckel, and W. A. Davidson. 2010. Advanced practices in travel forecasting. NCHRP Synthesis 406, National Cooperative Highway Research Program, Transportation Research Board, Washington, D.C.

Echenique, M., A. Flowerdew, J. D. Hunt, I. Skidmore, and D. C. Simmonds. 2007. The MEPLAN models of Bilbao, Leeds and Dortmund. *Transport Reviews*, 10:309–332.

Hess, S., A. Daly, C. Rohr, and G. Hyman. 2007. On the development of time period and mode choice models for use in large scale modelling forecasting systems. *Transportation Research Part A*, 41(9):802–826. doi: 10.1016/j.tra.2007.04.001.

Hunt, J. D. and J. E. Abraham. 2015. Design and implementation of PECAS: a generalized system for the allocation of economic production, exchange and consumption quantities. In M. Lee-Gosselin and S. Doherty, eds., *Integrated land-use and transportation models: behavioural foundations*, pp. 253–274. Amsterdam: Elsevier.

Knaap, G. and A. C. Nelson. 1992. *The regulated landscape: lessons on state land use planning from Oregon*. Cambridge, MA: Lincoln Institute of Land Policy.

Kok, K., A. Fallow, A. Veldkamp, and P. H. Verburg. 2001. A method and application of multi-scale validation in spatial land use models. *Agriculture, Ecosystems & Environment*, 85(1–3):223–238. doi: 10.1016/S0167-8809(01)00186-4.

Martin, R. 2003. *Agile software development: principles, patterns, and practices*. Upper Saddle River, NJ: Prentice-Hall.

Miller, E., B. Farooq, F. Chingcuanco, and D. Wang. 2011. Historical validation of integrated transport-land use model system. *Transportation Research Record*, 2255:91–99. doi: 10.3141/2255-10.

Næss, P., M. S. Nicolaisen, J. Andersen, and A. Strand. 2015. Forecasting inaccuracies: A result of unexpected events, optimism bias, technical problems, or strategic misrepresentation? *Journal of Transport and Land Use*, 8(3):39–55. doi: http://dx.doi.org/10.5198/jtlu.2015.719.

Oregon DOT. 2017. Rough roads ahead: the cost of poor highway conditions to oregon's economy. URL https://www.oregon.gov/ODOT/COMM/Documents/RoughRoads2014.pdf.

Waddell, P. 2002. UrbanSim: modeling urban development for land use, transportation, and environmental planning. *Journal of the American Planning Association*, 68:297–314.

Waddell, P. 2011. Integrated land use and transportation planning and modelling: Addressing challenges in research and practice. *Transport Reviews*, 31(2):209–229. doi: 10.1080/01441647.2010.525671.

Walker, P. and P. Hurley. 2011. *Planning paradise: politics and visioning of land use in Oregon*. Tucson, AZ: The University of Arizona Press.

Zhao, Y. and K. M. Kockelman. 2002. The propagation of uncertainty through travel demand models: an exploratory analysis. *Annals of Regional Science*, 36(1):145–163.