How is an Activity-Based Model Set Developed?

Presented at the Chicago Metropolitan Agency for Planning
Symposium on Developing and Implementing an Activity-Based Travel Demand Model
August 27, 2008

John L Bowman, Ph.D.
John_L_Bowman@alum.mit.edu
JBowman.net

INTRODUCTION

(page 1) This symposium is a great idea, and represents an important step toward activity-based models for CMAP.

Thank you for the opportunity to contribute by sharing my thoughts about how activity-based models are developed.

(page 2) To make sure we're on the same page, I'll start by describing an activity-based model system.

(page 3) By the time I'm done, this figure should be very familiar to you.

(page 4) The biggest part of my presentation will be devoted to looking at the development tasks

(page 5) The tasks can be divided into two groups: build (on the left) and tune (on the right). I'll spend time discussing each of the 12 tasks shown here.

(page 6) Then we'll look together at three basic approaches that have been used for completing these tasks.

(page 7) Invent, adapt, and adopt.

(page 8) Then we'll look at the several key roles that are needed for completing the tasks.

(page 9) -----

(page 10) And finally, based on my assessment of past and ongoing projects, I'll mention what I think are management keys to success….

(page 11) -----

(page 12) After I prepared all the slides for this presentation, I decided to wrap-up by giving a few specific suggestions. I mention this now, so that you won't be surprised when I slip in the postscript at the end of my remarks.

(page 13) But before I dive into the details of this outline, I want to spend a minute clarifying the context of my remarks.

What I have to say is based on the activity-based model systems shown on this timeline. This list includes only:

--US projects

--for metropolitan regions

--where a sponsor expressed an intent to implement and use the model, and

--development has started.

These are the most relevant projects given who CMAP is and your desire to learn from the experience of others who have gone down this path.

ACTIVITY-BASED MODEL SYSTEM

(page 14) So, let's start by looking at what one of these activity-based model sets looks like.

(page 15) To do this I start by showing the framework of a trip-based model system, since

      --that is where you and just about everybody else starts.

      --AND much of the framework and components are the same.

(page 16) The transport model system takes inputs from the urban forecasting and planning process:

(page 17)   --transport networks representing potential future scenarios

(page 18)   --zonal attributes representing forecasts of size and distribution of employment and population in the region

(page 19)   --and socioeconomic attributes of the population.

(page 20) It then

      --predicts zone-to-zone trip flows

(page 21)   --and assigns those trips to the network

(page 22)   --this occurs iteratively to achieve a final result where

--inputs to the trip demand models are consistent with

--predicted link flows and travel speeds coming from assignment of the demand they generate

(page 23) In some cases, accessibility measures from the transport model system serve as input to the land use forecasting models

(page 24) The trip demand includes

        --trips associated with special generators, such as airports

(page 25)   --trips carried out by persons who live outside the region

(page 26)   --trips conducted by residents of the region

(page 27)   --and commercial movements

(page 28) These trips are usually modeled by separate demand models,

        --and the results are combined into trip matrices for highway and transit assignment.

Here is where we can start talking of the difference introduced by an activity-based model….

(page 29) In an activity-based model system, only the trips of residents are modeled using an activity-based model.

(page 30) These are produced by a household travel demand simulator.

--It generates a synthetic population representing the future population of the region,

--and predicts activities and trips for every member of each synthetic household.

(page 31) In some cases, the activity-based model uses a large database of parcel attributes instead of only zonal attributes.

        --It takes a lot of resources to generate the parcel database

                --for model development

                --and for forecast scenarios

        --But it is very desirable because it provides substantially better information about

--transit accessibility,

--travel times for non-motorized and short trips,

--and destination attractiveness in heterogeneous zones.

(page 32) When parcel data is available the models use information about

--each individual parcel

--as well as information about the nature of a buffer zone surrounding the parcel

(such as the number and type of intersections and parking spaces)

Parcel buffer attributes represent neighborhood effects more realistically than zonal attributes because they are centered on the parcel.

(page 33) The main output of an activity-based travel demand simulator is a detailed itinerary for every person in every household

(page 34) It comes in the form of tables or lists that identify attributes

--of each household

--each person in that household

--each person's day

--each travel tour in their day

--and each trip on that tour

(page 35) These trips are then aggregated and combined with the other predicted trips into matrices for assignment.

So, at this point we see all the components of an activity-based model set needed for transport forecasting.

BUT some additional things are needed to unlock its power ….

(page 36) In order to leverage the simulator's ability to predict travel by time of day

--it may be necessary to enhance the highway and transit assignment components to assign trips for more time periods.

(page 37) The activity-based model output can also be aggregated in a great variety of ways,

--in order to extract information that is not available from trip-based models,

--ESPECIALLY the impact of a policy on various population subsets.

(page 38) So, to get at this information, it is desirable to create a reporting and query system that can be used to

   --produce standard reports

   --generate custom queries

   --and display results on GIS maps.

(page 39) Finally, the activity-based simulator has the potential to supply better measures of accessibility as input to land use models.

The existing activity-based models have not yet exploited this potential, but they are rapidly heading in that direction.

DEVELOPMENT TASKS

(page 40) Now that I have described the goal that we are aiming for, let's look at the development tasks required to get there.

(page 41) Design is the first, and most important, task.

If all subsequent tasks are done perfectly,

--but the design is flawed, then the forecasts will be unreliable,

--or if the design is limited, then the forecasts may not be useful.

--or if the design is inflexible, then the model system will be difficult to enhance.

(page 42) I want to highlight several important aspects of the design.

(page 43) The model system includes several model components on each of several levels

(page 44) At the top is the population synthesizer. It provides a synthetic population matching the best available estimate of the geo-demo-graphic distribution of the region's population.

After that come the models predicting outcomes that are longer term than a day.

(page 45) These are followed by models that predict the one-day pattern of activity for each person,

   --and perhaps for each household.

(page 46) Tour models predict the attributes of each tour in the day,

   --including destination, travel mode and time of day.

(page 47) Finally come the models that predict details of each trip and stop on the tour.

(page 48) Including a lot of related model components makes it possible to simulate itineraries that appear to be realistic.

But in order for them to **behave** realistically, they must be appropriately integrated.

(page 49) A basic integrating feature of the model systems is a model hierarchy.

Models lower in the hierarchy need to treat the outcomes of higher models as given.

(page 50) For example:

--If a person is a worker with a non-home usual work location,

(page 51) then their day is likely to be a work day.

If their day is a workday,

(page 52) then a work tour is likely to occur.

If there is a work tour, then it is more likely to go to their usual workplace than elsewhere,

And non-work tours cannot be made at times of the day that have already been claimed by a work tour.

(page 53) And Intermediate stops are likely to occur near their work destination.

In order to properly achieve this downward integration, the design must specify how to do it in three subsequent tasks:

(page 54) --in preparation of the survey data for model estimation

--in model estimation itself

--and in programming of the simulator software.

And it must be done consistently in all three steps or the model system won't give reliable forecasts.

(page 55) Models higher in the hierarchy need to be sensitive to conditions affecting lower level models. This is upward integration.

(page 56) For example,

--If transit service in a certain corridor is significantly improved for trips occurring during the peak period,

(page 57) then

--tours by transit at that time of day should increase among persons living and working in that corridor

--and their tours in the corridor by other modes should probably decrease.

(page 58)   --more people will arrange their day to include a transit commute,

--but non-work auto tours might increase

    --as workers run errands after work that they used to do by car on their work commute.

    --and more non-workers have a car available during the day

(page 59)   --Workers in the area may be more likely to buy a transit pass,

    --and more households in the area should eventually own less cars

The preferred method of capturing these kinds of affects is through the use of logsums

    --these measure the expected utility across lower level alternatives

    --and are used as explanatory variables in the upper level models

A big problem is that Logsum calculations take a lot of time

    --so compromises are made in order to speed up run times,

    --and creative techniques are developed to reduce the bad impact of such compromises.

(page 60) As with downward integration, the design must specify how to implement upward integration consistently:

--in preparation of the survey data for model estimation,

--in model estimation itself,

--and in programming of the simulator software.

(page 61) Because of the need for consistent models and software, the simulator should be designed during this task.

(page 62) Here is a snippet of pseudocode from the design of one of the existing simulators.

(page 63) Finally, the design needs to cover the entire transport model system, in which the activity-based simulator operates.

(page 64) Of special importance are two aspects:

--equilibration: as with a trip-based model system, the demand and assignment must be iterated to achieve consistency in both parts of the system

--and performance: an activity-based model system can have more computation than its trip-based counterpart, so the design needs to include distributed processing of both the demand and assignment portions of the model system.

(page 65) The last aspect of design that I want to emphasize is that you will have major decisions to make at this stage.

Although the model systems in use today are strikingly similar, they also have some major differences. You will need to evaluate these differences in light of YOUR NEEDS and decide between them.

--Will you model spatial choices using parcels or similar small geography (as Sacramento did), or will you use a more traditional zone-based approach?

--Will you explicitly model joint decisions and actions of two or more household members (as pioneered in Columbus), or will you rely on implicit modeling of these interactions

--and there are other differences among existing model systems that you should consider.

Beyond that, the microsimulation framework is well-suited to innovation.

--Most of the projects so far have introduced substantial innovations in order to achieve specific objectives,

--and there are still many advances to be made, in the areas of

     --Parking (for parking subsidies, park&ride, and park&walk)

     --Vehicles (with vehicle type modeling, vehicle choice for trips, and implementation of vehicle time and space constraints)

     --pricing,

     --and Transit (such as pass-holding, subsidies)

What specific needs do you have that may require innovative features?

And do you want to risk the potential schedule delay and cost that come with innovation?

(page 66) Having spent considerable time on the design, because of its importance,

I will move quickly through the remaining eleven tasks in the development process,

(page 67) starting with the preparation of data for model estimation.

(page 68) In most cases, the design is actually done AFTER the household survey has been designed and at least partially implemented. This seems to be your case as well, although I understand that you have involved activity-based modelers in the survey design, and that was a good decision on your part.

Nevertheless, after the design is complete, then you will need to prepare the household survey data for model estimation,

--to represent tours and trips in accordance with the design

--to identify the choice set and outcome for each choice

--to properly restrict choice sets based on higher level outcomes

--and to associate level of service and spatial attributes with each choice alternative

(page 69) Level of service data will need to come in the form of skims from the existing model system.

You may decide to make some enhancements so you can generate skims that are as consistent as possible with the new model design, with regard to

--the mode definitions,

--and variations in level of service by time of day.

(page 70) For each zone, or parcel if you decide to go that way, you will need to have estimates of

--employment, school enrollment and housing units in various categories,

--as well as network attributes such as intersections, transit access, and parking.

(page 71) Many books, courses and research papers have been devoted to the estimation of models, so I will not go into any detail here,

(page 72) Other than to say that, for each model that is estimated using local data, it must be specified, estimated and tested in an iterative process, using the household survey data.

(page 73) Building the activity-based simulator is a major task

(page 74) And at the end of the development process, the simulator is the tangible result that

--takes the form of an executable program

--and embodies the activity-based portion of the model system

(page 75) It creates the synthetic population, applies all the models, and constructs a detailed one-day itinerary for each person.

In doing this it enforces the hierarchy, time & space constraints, and upward integration prescribed by the design.

(page 76) You can start building the simulator once you have your design.

(page 77) However, by coordinating this task with the prior tasks, the simulator can use code that was written for

      --preparing the estimation data

      --and estimating the models.

Re-using this code has the big advantages of saving time and avoiding inconsistencies in the final product.

(page 78) Once the simulator is complete, it can be used to do a preliminary validation of the activity-based portion of the model system.

(page 79) To do this, it is run for the year of the household survey, and its results are compared to aggregate statistics from the survey.

--It may be necessary to adjust calibration constants in component models,

--or perhaps to re-estimate one or more of the component models.

--And it is a natural time to discover and get rid of bugs in the simulator.

(page 80) Since the activity-based simulator only forecasts trips of residents, the other components of the transport model system must also be built.

(page 81) Most of these components probably already exist. During the design you will have decided whether to use them as they are, or to enhance them.

(page 82) However, a program will be needed to aggregate the trips from the activity-based simulator into matrices.

(page 83) And you will probably need to enhance your assignment and skimming procedures.

(page 84) Although you might have done that already in order to prepare skims for model estimation.

(page 85) Once you have a validated activity-based simulator and the other components of the model system, you're ready to put them together….

(page 86) …into the integrated transport model system.

(page 87) This involves

        --installing hardware if necessary,

        --assembling the scripts that run all the components,

        --and testing and tuning the system.

Here the focus of attention is convergence and performance.

It may be necessary to adjust the iteration scheme and the distribution of processes among machines.

In the end, you need to achieve two conflicting objectives simultaneously:

--link flows and speeds coming from assignment must match those used to generate the final skims used as input in the demand model components

--and the model must complete its job in an acceptable elapsed time.

This is a challenge, and none of the existing model systems has fully achieved of these twin objectives.

(page 88) The last two tasks on the build side of this diagram could be put off until you need them, but there are advantages to planning ahead for them because they're important, and useful for tuning the model system.

(page 89) The first one is a reporting and query system, used to extract information from the output of the activity-based simulator.

(page 90) If it is developed soon enough, it can be used to extract model results for base year model system validation.

(page 91) It extracts information from

--the output of the activity-based simulator,

--and the convergent networks,

(page 92) The objective is to make it easy to aggregate the results as needed, and get them into

      --standard reports,

      --custom queries,

      --and GIS-based display forms

(page 93) The last build task is to implement procedures that generate input data for forecast scenarios.

(page 94) These will be useful near the end of the project for model sensitivity tests.

(page 95) This task is especially important if you decide to use parcel information in the models. The Sacramento project has shown that generating forecast scenarios at the parcel level, while extremely valuable, can also be a lot of work, and automated procedures are needed to assist with that process.

(page 96) [pause] Now we're ready to move to the model system tuning tasks.

First is the preparation of data for validating the ability of the model to replicate base year conditions.

(page 97) If the base year for model forecasts is the same as the HH survey year, then some of this will have been done when data was prepared for model estimation.

(page 98) Otherwise, the base year network, spatial and population data must be assembled.

(page 99) And this must be supplemented by fairly standard validation data.

One of the main differences may be that, in order to validate the temporal aspects of the model system, the validation data should be as detailed as possible by time of day.

(page 100) With this data in hand, base year validation can occur.

(page 101) If the base year is different than the estimation year, then the activity-based model must be re-validated on the base year.

(page 102) Beyond that, base year validation is much like that for a trip-based model,

--but validation should be by time of day.

If significant discrepancies occur, then calibration constants or re-estimation of models may be necessary.

(page 103) Finally, the model is ready for forecasting, and the final task is to do some forecast sensitivity tests.

(page 104) Elasticities can be checked for reasonableness.

If problems are encountered, you might need to decide between

--working out the problems before releasing the model for prime time use, on the one hand,

--and releasing the model and queuing up a project to improve it, on the other.

At this stage, the model can also be used to

       --train users,

       --and familiarize clients with its capabilities.

BASIC BUILD APPROACHES

(page 105) What I have just completed is a description of the development process if you are developing a model system from scratch.

But there are important variations on this theme that I am calling basic build approaches.

(page 106) And I've identified three basic approaches: invent, adapt and adopt.

Another way to think about it could be to put invent on one end of a spectrum of approaches, adopt on the other end, with adapt occupying everything in between.

(page 107) Among all the US projects I showed you earlier,

(page 108) I've classified only Metro, NYMTC, and MORPC in the Invent category. (although, in reality, they all incorporated prior work of others)

(page 109) In contrast, San Francisco, Sacramento, Denver, and now Seattle, are clear adaptations of the early work done in Portland. And later projects in this line have tended to adapt from the most recent adaptation.

(page 110) Atlanta and MTC are implementing adaptations of the MORPC model.

(page 111) And the Tahoe area model is the only one that can be classified as an adoption, taken straight from MORPC.

(page 112) As I am using the term here,

Inventions carry out all aspects of the development process as I described it earlier.

Some of them borrow significant aspects of design from earlier research and models developed in academic settings, but they are not based primarily on any one operational model system.

(page 113) On the other hand, even those I classify as inventions may adapt the non-activity-based components and the integrated model system from an existing trip-based model system.

(page 114) Adaptations borrow heavily from other model systems for the development of the activity-based simulator.

In most cases this includes not only the design, but also

--the programs used for preparing estimation data,

--the programs used to estimate the models,

--and the activity-based simulator program.

Adaptations usually do most of all these steps, and usually introduce innovations to achieve important objectives, but they borrow heavily from a prior model system in the process.

(page 115) An adoption, on the other hand, takes an existing simulator, incorporates it into a full model system, and validates it using local data.

This bypasses the costly steps of preparing data, estimating models and programming a simulator.

It also eliminates the need for a regional household survey.

(page 116) Significantly, nearly all the model systems started in the last several years are adaptations or adoptions

(page 117) Only Metro has decided to embark on a new design, and this work is just beginning.

DEVELOPMENT ROLES

(page 118) Now let's look at the development tasks from the perspective of the roles assigned to various experts.

(page 119) Each activity-based model development project has been unique, including the way that roles have been assigned.

However, there seems to be a pattern of assigning responsibilities to four primary roles, with each role tending to have more than one critical area of expertise.

(page 120) The activity-based developer is responsible for the overall design,

--and also for the development of the activity-based models

--and simulator.

This requires knowledge of activity-based model system design, econometric skills, and programming skills.

(page 121) But because the activity-based model fits into a traditional model framework, an expert in the region's trip-based model system takes primary responsibility for the trip-based components,

(page 122) and, with advice from the activity-based developer,

--may also take primary responsibility for implementing and tuning the integrated transport model system.

(page 123) For the tasks of dealing with the inputs and outputs required by the model system, one or more persons with extensive skills in GIS, database and graphical user interfaces are needed.

(page 124) And finally, a person with extensive application experience should be heavily involved in model system validation and sensitivity testing.

This leads naturally to their ongoing role of Model System Expert from the standpoint of internal users and external clients.

(page 125) I think it is important to note that in all of the model systems now in use, and in most of the projects under way,

--an external consultant filled the role of the activity-based developer,

--and fulfilled other roles as needed.

And in all the completed projects, the developer has had a repeated or ongoing role after implementation.

(page 126) It is also important to note that DRCOG has deliberately chosen to share the activity-based developer role with the external consultant.

DRCOG

--assisted heavily in the design,

--estimated some of the models,

--and augmented the survey data preparation as needed for model estimation.

And perhaps most significantly, they chose to be the primary developer of the activity-based simulator and the model system software as a whole.

I'm glad that Erik Sabina is here and I hope that he shares with you from the DRCOG experience in this regard.

MANAGEMENT KEYS TO SUCCESS

(page 127) I have reflected from time to time on the various projects that I have been directly involved with or observed from a distance, trying to assess what causes some projects to be more successful than others. From this I have identified three keys to success.

(page 128) Before describing them I should give you my working definition of a successful project, in bullet form.

The key words here are implemented, sound, and used.

(page 129) In addition, the most successful projects experience

--cost effective development

--and timely development

These two factors tend to go hand in hand. Projects that string out over a long period of time tend to waste a lot of money.

Because existing model systems have not yet tapped all the potential of the approach, successful projects

--introduce useful innovation

--and provide a foundation for ongoing enhancement.

(page 130) Given this definition, I will share three things that appear to be keys to success, beginning with a sound design.

(page 131) For a sound design the basic framework needs to be workable.

--Nearly all of the projects started with a basic framework that had been earlier shown to work.

(page 132) Although I called the first Metro project an invention, it was actually based on a working prototype that had been built in the lab at MIT.

A workable framework

--helps assure that a sound result is possible,

--it provides a vision for what is possible,

--it fosters confidence in the sponsor,

--and it gives the developers something to build upon.

(page 133) You will have some major design decisions to make, and they should be made up front. It is not a good idea to proceed into development without a clear design architecture.

(page 134) The design needs to be comprehensive and integrated.

Each of the important component models should be identified, along with an understanding of its basic model form

The methods of integrating all the model components must be sound and clear. This includes

--the integration among the activity-based model components

--and the integration of the activity-based simulator with all the other components of the entire transport model system,

--including the equilibration of demand models with traffic and transit assignment.

(page 135) And the design must be implemented consistently

--in the preparation of data for model estimation

--in the model estimation itself

--and in the simulator software.

(page 136) The second key to success is capable innovative developers.

(page 137) Here I'm not just talking about the activity-based model developer, but people in all four of the main project roles.

(page 138) This enhances the technical soundness of the product.

(page 139) Also, it allows the innovations to occur.

Although it is now possible to implement an activity-based model system just like one that exists, innovation is still the norm, rather than the exception, and can help achieve important objectives.

(page 140) Paying attention to the development roles of those who

      --deal with the inputs and outputs,

      --and those who run the models,

can significantly improve the usability and usefulness of the model system,

(page 141) and assure that the necessary follow through occurs successfully:

      --breaking in new features,

      --seeing the need for and initiating ongoing improvements,

      --and dealing effectively with model users and clients.

(page 142) The first two keys refer to aspects that I spoke of in my presentation. The third key is sustained sponsorship. I didn't speak of this earlier, but it is just as important for project success.

(page 143) Without a motivated sponsor, the necessary funding stream won't materialize, or it will dry up and kill the project just when it is needed most.

(page 144) Each successful project has had a motivated sponsor, and each sponsor had its own motivations.

--SFCTA had specific unmet information needs

--NYMTC needed a model and the region was too complex for a 4-step model

--MORPC had a consultant that was willing to fund the extra cost.

(page 145) At the beginning, the project needs an instigating advocate to secure sponsorship and get the ball rolling.

Most of the projects can point to a specific person who took this role at the outset.

--Keith Lawton at Metro was the first of these advocates

--And Gordon Schultz served that role with NYMTC, ARC and MORPC

(page 146) And as the project progresses, it needs an internal champion to make sure that momentum is maintained through implementation and beyond

(page 147) Let's look at two examples that cannot be declared successes, at least not yet. One of these models is no longer being used, and another has been in development for a very long time.

(page 148) First, "Why didn't Metro keep using their model?"

(page 149) The problem had to do with sponsorship. When the time came to calibrate and validate the model,

(1) the MPO was struggling financially,

(2) And substantial federal funds became available from TranSIMS.

(3) so the development money and staff resources were dedicated to the TranSIMS project,

(4) and the model calibration and validation work was never funded.

(page 150) Second, "Why is it taking ARC so long?"

(page 151) Again, the answer is not technical. Again it has to do with sponsorship.

(1) In this case, a cautious ARC top management chose to invest at a very slow rate in the development effort.

(2) They also became pre-occupied with a major geographic expansion of their MPO region from 13 to 20 counties,

(3) and didn't commit adequate resources to development until this year.

(page 152) So, among the projects that we are considering, although I think all three keys are ESSENTIAL

Lack of sustained sponsorship has the factor that has prevented success.

POSTSCRIPT—A FEW SUGGESTIONS

(page 153) So now we are ready for the postscript that I warned you about…a few suggestions for you to consider as you map out your future.

I don't know your exact situation, so some of these suggestions might be inappropriate, but I hope they at least provide food for serious thought and discussion.

(page 154) First, adapt, rather than inventing or adopting.

You will be further ahead building on the success and investment of others than starting from scratch.

Try to leverage the existing designs, models, software and knowledge of others.

(page 155) Second, don't wait on your survey data.

Start your design now.

If you get that done and the survey data still isn't available, you can make any needed upgrades to your assignment and other non-activity-based models that need to change

(page 156) Third, seriously consider using parcel data.

If your region has a good parcel database, use it.

If it is moving that way, aim to use it eventually.

But don't hold up activity-based design and development waiting on parcel data

(page 157) 4. Innovate with care, if you need to.

If the existing designs won't achieve an important objective very well, then try to figure out how to enhance the design to achieve it.

For example, if parking capacity is a major issue downtown or at park and ride locations, then maybe you can work up a design that incorporates explicit modeling of parking lot fill status by time of day into the activity-based household travel demand simulation.

But get a satisfactory design before committing to an innovation

And avoid more than one, or maybe two, major innovations, especially on your initial implementation.

(page 158) Fifth, implement promptly, and then enhance.

If possible, plan to have your first model system built and functioning within two years after your household survey data is cleaned and ready.

Stretching it out much longer wastes money, hurts quality and loses momentum.

Existing implementations have shown that experience with an operating model provides a good basis for identifying and making further enhancements.

Wrap-up

(page 159) This brings my comments to a conclusion.

I hope that I have provided some insights, language, and food for thought that will help as you proceed down your path toward activity—based model systems.