Activity-Based Models:
Approaches Used to Achieve Integration Among Tours and Trips Throughout the Day

Presented at the European Transport Conference
Leeuwenhorst, The Netherlands, October 6-8, 2008

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

Mark A. Bradley
mark_bradley@cox.net

(page 1)

(page 2) In developing the so called activity-based models, our objective is to more realistically predict the affect of land use and transport policy on travel behavior.

(page 3) And we try to do it through a more realistic representation of that behavior.

(page 4) This includes modeling an entire day of activity and travel as one integrated outcome.

(page 5) Trip-based models deal with separate one-way trips.

Tour based models model separate tours.

We attempt to integrate the representation of travel demand across an entire day.

(page 6) Trying to realistically connect tours throughout the day makes time-of-day models an important component of the model system,

(page 7) assigning time periods to each activity episode and trip in the day's itinerary.

(page 8) Because of the household's importance in shaping each individual's activities and travel, some modelers have also emphasized including household aspects of personal travel,

(page 9) such as

coordinated schedules

joint tours

household activities that get assigned to individuals to carry out

and shared rides where one member of the household chauffeurs or escorts another.

(page 10) Here is a diagram of an activity-based model system

--The activity-schedule models are implemented in a simulator that generates an itinerary of trips for each person in the population.

--These trips are aggregated into OD matrices,

--combined with commercial traffic and other trips,

--and assigned in the highway and transit assignment procedures.

--This process is iterated to an equilibrated solution.

The focus of this presentation is the models in the gray box, not the traffic assignment or the overall system.

That box looks pretty simple. But it actually contains many component models.

(page 11) And these are strung together in sequences, as shown by this diagram that depicts the generation of tours in the current New York model. And this diagram doesn’t even deal with the modeling of destinations, modes and times of the travel on the tours.

(page 12) But wait a minute! I just said we try to model the entire day of activity and travel as one integrated outcome.

How can a sequence of models capture the correlations that exist among the sequentially modeled components?

(page 13) My thesis, and the motivation for this paper, is that,

in order to do their job

activity-based model systems must be effectively integrated.

(page 14) So, I studied the four activity-based models now in use for regional travel forecasting in the United States, to see how they are integrated.

These are the models in San Francisco, New York, Columbus and Sacramento.

I will refer to them by the initials of the agencies that own and use them: SFCTA, NYMTC, MORPC and SACOG.

One note I should make about SFCTA. It is the first of these model systems, and has been recently going through some major enhancements. My presentation is not informed by these enhancements.

(page 15) I will describe the integration employed by these models using the terminology of horizontal and vertical integration, with vertical split into downward and upward.

(page 16) As I use the term horizontal integration, it means a joint model of two or more model components.

To use a familiar example of mode and destination choice,

a model that is estimated and applied to simultaneously model mode and destination choice is horizontally integrated over those two dimensions.

It could be one of several model forms; for purposes of today’s discussion it doesn’t matter.

(page 17) In contrast, vertical integration ties together two models that are estimated and applied sequentially.

Here I depict destination choice before mode choice.

(page 18) Downward vertical integration occurs by constraining the subsequent mode choice model by the outcome of the prior destination choice model.

(page 19) And upward vertical integration occurs by accounting in the destination choice model for the quality of the alternatives that are available for consideration in the subsequent mode choice model.

So, the destination choice is modeled first, but not blindly! It is affected by the factors that affect the mode choice.

(page 20) The first of five aspects of integration that I will describe is integration of the three core tour models of destination, mode and time-of-day choice.

(page 21) MORPC uses full downward integration and partial upward integration.

Upward integration is implemented using logsums.

However, in the destination choice model,

--instead of taking a logsum across all possible time periods,

--it uses an assumed time of day that is the same for all similar cases in the population.

It does this because it is time consuming and impractical to calculate the logsums across both time and mode dimensions.

But it comes at a cost.

Suppose that it uses peak period for the logsum for all work tours. Now suppose that a forecast scenario introduces off-peak congestion to an area that is already fully congested during the peak period, and doesn't get worse in the scenario. Unfortunately, the logsum won't reflect the change, and destination choice will be unaffected.

(page 22) SACOG also uses partial upward integration, but introduces simulation of time-of-day in an effort to overcome the drawback just illustrated.

The destination choice also uses a mode choice logsum and an assumed time of day. However, instead of using the same assumed time of day for all similar cases,

it draws a time of day according to a probability distribution that applies to those cases.

In most work cases, it draws peak-period times, but for some cases it draws other times. At an individual level, the destination choice is responsive only for one time period, but across the population there is responsiveness to all time periods.

This approach isn't perfect, but short of full upward integration with logsums, it is the best technique I have seen for integration among these three model components.

(page 23) The second aspect of integration we'll look at is integration between long-term and within-day model components.

The first activity-based models modeled only auto ownership as a long-term choice. They modeled work location as a within-day choice. This fails to recognize the effect that going to work the same place most days has on people's choice of daily activity pattern.

(page 24) SFCTA was the first to model usual work location as a long-term choice, and it does it as shown here.

Downward integration conditions auto ownership on workplace choice and conditions the daily activity schedule on the long-term choices.

For upward integration it uses a work tour mode choice logsum to make the work location choice sensitive to level of service. Here again, however, it uses a non-simulated assumed time of day.

Auto ownership doesn't use person-specific logsums. Instead it uses an ad hoc zonal accessibility measure. The measure is sensitive to auto and transit travel time, but it is not sensitive to other aspects of auto and transit impedance, such as tolls; it is not sensitive to time-of-day factors; and it does not depend on characteristics of the household.

SACOG is like SFCTA except it also models long-term school location, and it uses mode choice logsums and mode-destination logsums in all three models.

NYMTC and MORPC model only auto ownership as long-term choice.

(page 25) The third aspect of integration is downward integration among tours and stops during the day.

(page 26) Given that tours and stops are modeled in a sequence, the objective of this downward integration is to realistically limit the choice sets and condition the utility functions of subsequent destination, mode and time-of-day models in light of the prior modeled outcomes.

(page 27) SFCTA and NYMTC are both pretty weak in modeling time of day. They use so few time periods that the prior time-of-day outcomes provide little meaningful information to constrain subsequent time-of-day choices.

MORPC was the first to introduce a model with one-hour time periods, and to use a detailed accounting of time usage to constrain the timing of subsequently modeled tours.

SACOG reduced the size of the time periods to a half hour and extended the use of time accounting to constrain time-of-day choice for intermediate stops on the tours.

(page 28) This slide depicts the sequence that SACOG uses to model the intermediate stops on each half-tour.

The half-tour is conditioned by the prior-modeled pattern, tours and half-tour.

Within the half-tour, the number and purpose of stops is modeled first, then the stop location, trip mode and time is modeled for each stop, beginning with the stop nearest to the tour's destination.

(page 29) Before modeling the half-tour, we already know several things about the tour, including the purpose, destination, mode, and arrival and departure times.

These are taken into consideration in modeling the half-tour.

(page 30) After modeling the number and purpose of stops on the first half-tour, we know the number, purpose and sequence of the stops,

and are ready to model the details of the stop nearest to the destination.

(page 31) First the location is modeled, taking into consideration the tour mode, as well as the impedances associated with detouring to make the stop, in this case on the way to work.

Then the trip mode is modeled, constrained by the tour mode choice and the mode availabilities at the location.

The stop location, trip mode and arrival time at the tour destination determine the departure time from the stop. This leaves the arrival time to be modeled.

This process is repeated for each stop on the half-tour, working back to the tour origin.

Thus SACOG's approach uses a fairly rigorous downward integration for location, mode and timing choices.

It would be more effective if the time periods were smaller. Or, another approach might use Monte Carlo simulation of a specific time after each time period is modeled, and use the specific time to condition subsequent choices.

(page 32) One desirable aspect of downward integration that none of the models uses would be to account for vehicle schedules in the same way that MORPC and SACOG account for each person's schedule.

This would involve assigning one of the households' vehicles every time a person chooses an auto drive mode, and blocking out that vehicle's availability for other uses for the duration of that tour.

This could be preceded by a long-term model that identifies not only the number of household vehicles, but also the types. And the mode choice could then include choice of a specific vehicle.

This type of vehicle modeling and downward integration would enable the model system to associate a specific vehicle type with every modeled trip, and identify whether the vehicle engine was cold or hot when it started. This would significantly improve the information available for emissions and air quality modeling.

(page 33) Our next aspect of integration is integration among household members.

All four models model auto ownership as a household choice.

Beyond that, SFCTA and SACOG use individual models exclusively, and don't integrate them explicitly among household members. They rely on household characteristics to capture any household effects on individual outcomes, ignoring correlations among within-day outcomes of household members.

NYMTC and MORPC both integrate within-day choice among household members. MORPC's is more sophisticated than NYMTC's, so I will describe how MORPC handles it.

(page 34) MORPC uses an extensive sequence of downwardly integrated models. This diagram depicts the entire sequence as five boxes, each of which is itself a sequence of models. Integration among household members is emphasized in linked activity patterns, non-mandatory tour generation, and non-mandatory tours.

(page 35) For each person, a simple activity pattern is modeled.

These are implemented in a prescribed sequence by person type, and the pattern choice of subsequent persons is conditioned by the pattern choices of prior-modeled persons. The main thing this accomplishes is that it captures the tendency for multiple people in the household, on any given day, to stay at home together, or to not go to work or school.

(page 36) Given the patterns, the mandatory tours are modeled for each person. Destination, timing and mode for each tour are integrated as I described earlier.

(page 37) Then a sequence of downwardly integrated models generates non-mandatory tours.

First, a horizontally integrated model generates joint tours for the household. These tours are shared in their entirety by two or more household members. Subsequent downward integrated models assign household members to each of these tours.

Then, maintenance tours are similarly generated for the household and assigned to individual household members to be carried out.

Finally, discretionary tours are generated for each person in the household.

(page 38) After the non-mandatory tours are generated, the destination, timing and mode are modeled for each non-mandatory tour, again using the integration techniques I described earlier.

Finally, simple intermediate stop models and procedures identify any stops and their details for each tour.

(page 39) This slide shows two additional possible household integration features that none of the four model systems uses.

The first idea is to horizontally integrate the simple pattern choice among all household members. This feature is being used in new models now under development, and might make it easier to upwardly integrate the tour detail models with the household pattern choice.

The second idea is to model shared partial tours among two or more household members. Shared partial tours occur in very complex ways, and would be very difficult to model. However, they occur a lot, and modeling them correctly might improve the accuracy of the model's response to transport policies.

(page 40) The last aspect of integration is horizontal and upward integration among tours and stops, which I consider to be very important

(page 41) for achieving the modeling objective

(page 42) Let's look back at the MORPC model sequence, now focusing attention on horizontal and upward integration.

The generation of mandatory tours (in the linked activity patterns), the non-mandatory tours, and the intermediate stops are separated, forcing the model system to rely entirely on downward and upward integration to capture the important correlations among them. The downward integration is pretty rigorous, but upward integration is lacking. Why do I say this?

(page 43) In the individual pattern model, zonal accessibility indices measuring walk accessibility to jobs and retail influence the pattern choice.

But pattern choice is insensitive

--to auto and transit accessibility

--to accessibility that varies by time of day

--and to accessibility for secondary tours and intermediate stops.

So, when these things change, they have no effect on pattern choice.

(page 44) Similarly, the generation of maintenance and discretionary tours is insensitive to

--accessibility for subsequently modeled tour purposes

--to accessibility for intermediate stops

--and to transport accessibility for more than one mode and time period

The generation of joint tours is not sensitive to transport conditions at all.

(page 45) Now let’s turn attention to the SACOG model system.

SACOG uses a day pattern model that identifies for each of seven purposes

--whether one or more tours are made for that purpose

--and whether an intermediate stop is made for that purpose on at least one tour

This horizontal integration enables the utility for any one purpose to directly affect the probability of it and all other purposes.

The model is also vertically integrated through the use of three types of purpose-specific logsum accessibility measures.

(page 46) Here we see the seven purposes modeled in the day activity pattern

(page 47) and the three types of logsums

(page 48) Work and school mode choice logsums are used in the utility components of work and school tours and stops. These are specific to the person’s usual work and school locations. They also use an assumed time of day; it would be better if they took all possible times into consideration, at least through the use of simulated times. (pause)

The next type of logsum is a mode and destination logsum that measures accessibility to all available destinations by all available modes.

(page 49) Because the calculation of mode-destination logsums is very time consuming, SACOG pre-calculates a large number of them for each zone, each time the model system is run. For a given zone, there are 84 distinct logsums defined in ways that most affect the logsum values:

--activity purpose,

--car availability,

--and proximity of origin parcel to transit.

Then, in each case where such a logsum is needed, the one that matches the purpose, car availability and transit proximity is used.

(page 50) This ‘approximate’ or ‘aggregate’ logsum affects the attractiveness of tours and stops for all purposes in the pattern model. (pause)

The third type of logsum measures the accessibility for making a stop on the way to or from a tour destination. In the pattern model, it is used to measure intermediate stop accessibility for tours to the usual work location.

(page 51) It is calculated as the logsum of the utilities of the stop location choice model.

The upward arrows added to the half-tour model diagram show that this intermediate logsum is not only used in the pattern model, but it also affects the prediction of the exact number and purpose of intermediate stops on each half-tour.

(page 52) Intermediate stop logsums are especially time consuming to calculate, so these are also aggregate. There are four for each zone pair. Each one is mode-specific and time-of-day specific.

(page 53) In the pattern model, since the mode and time of the modeled tours is not known, the logsums are specific to the auto mode, and for an assumed time of day.

(page 54) This completes my examination of five specific aspects of integration used in the existing US activity-based models.

(page 55) Before we open things up for discussion I’ll briefly summarize my opinions on integration priorities.

A basic distinction can be made between techniques used to integrate tours and stops and those used to integrate choices made among household members.

In my view, both are valuable, but integration among tours and stops for individuals should be higher priority than integration among household members,

because this is what makes the model appropriately sensitive, at all levels, to transport conditions and land use attributes.

(page 56) The integration among household members certainly has the potential to make the models more realistic.

I would like to see horizontally integrated household day pattern models

--and would like to see them incorporate the generation of household maintenance tasks that could be carried out as tours or intermediate stops.

The modeling of joint tours is also valuable,

as would be the modeling of shared partial tours.

However, based on my examination of the working model systems, it is clear that it is difficult to implement household integration without severely compromising the integration of tours and stops.

I think that when model systems include advances in household integration, they should also preserve rigorous tour and stop integration.

(page 57) There should be purpose-specific horizontal integration among tours and stops in a day pattern.

It should be accompanied by upward integration that captures the affect of transport conditions on that pattern.

In doing so, the vertical integration should account for

--differences among persons and purposes,

--as well as available destinations, modes and times.

It should also be accompanied by downward integration that accounts for chosen purposes, locations, modes and times.

And I also see great potential value to enhance the models to account explicitly for vehicle type and usage.