ACTIVITY-BASED MODELS: APPROACHES USED TO ACHIEVE INTEGRATION AMONG TRIPS AND TOURS THROUGHOUT THE DAY

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

John L. Bowman, Ph. D.
email: John_L_Bowman@alum.mit.edu, website: http://JBowman.net
Mark A Bradley
mark_bradley@cox.net

ABSTRACT SUMMARY

We compare various integration techniques used by four activity-based models that have been used for travel forecasting in the US, providing conceptual understanding and reasoned discussion of their strengths and weaknesses.

ABSTRACT

This paper examines the so-called activity-based models implemented to date in the United States, explaining and comparing the various techniques that have been used to achieve model integration. These models integrate the representation of activities and travel conducted by an individual, and in some cases an entire household, over the course of an entire day. Such integration is what distinguishes these models from earlier trip-based and tour-based models. Three techniques of integration are typically used. First, a model is developed that simultaneously represents outcomes spanning the tours in a day and, in some cases, the persons in a household. Sometimes called an “activity pattern” model, it provides what could be called horizontal integration across all the dimensions of choice. Second, since the outcomes that need to be modeled are more complex than can be represented in a single activity pattern model, additional aspects of choice are modeled by breaking the outcome into a conditional model hierarchy or a chain of models. Models lower in the hierarchy (or later in the chain) take as given the outcomes higher in the hierarchy. This achieves what has been referred to as downward vertical integrity. Done properly, it assures that lower level models adhere to constraints imposed at higher levels, and makes the lower level models indirectly sensitive to all variables that directly affect the upper level outcomes. Just as important as downward integrity is upward vertical integrity. Upward integrity comes from making the upper level models appropriately sensitive to variables that affect the upper level outcome, but can’t be measured directly because they differ among the undetermined lower level model outcomes.

In practice, a variety of techniques has been developed and used to achieve horizontal, downward and upward integrity. This paper examines and compares the techniques employed by four model systems that have been used for travel

forecasting and policy analysis in the United States (in San Francisco; New York City; Columbus, Ohio; and Sacramento.) Among the techniques considered are:

- individual pattern models that explicitly identify the tours in a day and the intermediate stops that occur on the way to and from the tour's primary destination,
- individual pattern models that identify the purposes for which tours and extra stops occur, without associating the purposes and stops to specific tours,
- models implemented without an activity pattern model, using instead a cascade of tour generation models by purpose,
- household pattern models that identify the primary activity of the day for all persons in the household,
- household pattern models implemented as a hierarchy, identifying the presence of joint tours conducted together by household members and tours conducted by individuals to achieve household maintenance activities,
- time window techniques that enforce realistic time constraints: conditional tours and stops are limited to windows of time not occupied by higher priority tours and stops,
- half-tour models that identify the number and purpose of intermediate stops on each half-tour, given the purposes for which stops are to be made in the day and the stop purposes already included on higher priority half-tours,
- half-tour models that simulate stop locations, timing and mode in a temporal sequence emanating from the tour's origin or destination,
- long-term models of usual work and school locations that condition the daily activity pattern, tour and stop locations,
- disaggregate tour mode and mode-destination logsums used to measure accessibility in various higher level models within the model system, relying on simulated or assumed time-of-day choice in the logsum calculation,
- tour mode and mode-destination logsums calculated for aggregate market segments and representative conditional alternatives, requiring less computation time than disaggregate logsums, and
- aggregate trip destination logsums that measure the attractiveness of locations along the path between a tour's origin and its primary destination.

The techniques employed are not always directly comparable, the models are still relatively new, and the situations in which they are used vary considerably. Therefore, a meaningful empirical comparison of the techniques is not feasible. Accordingly, the emphasis in this paper is on providing a comparative conceptual understanding of the techniques themselves, including a reasoned discussion of the potential strengths and weaknesses of the various approaches.

1. INTRODUCTION

This paper examines the so-called activity-based (AB) regional travel forecasting models implemented to date in the United States, explaining and discussing the various techniques that have been used to achieve model integration.

These models simulate 24-hour activity and travel itineraries for each synthetic resident of a region. The resulting trips are aggregated into trip matrices, combined with commercial trips and trips of non-residents, and assigned to transit and road networks (see Figure 1). In simulating the itineraries of one person, many dimensions of choice are modeled, including activity participation, timing and location, as well as the mode of associated travel. It is necessary to address the issue of integrating multiple model components because, on the one hand, the outcomes are related to such an extent that it seems appropriate to treat them as a single complex outcome, modeling all dimensions simultaneously and, on the other hand, the outcome is so complex that it is impractical to capture all the needed detail in a single model. Therefore, the models have been implemented as a large number of carefully integrated components. The objective of these models, and hence the objective that guides the selection of integration techniques, is to realistically model travel behavior that can be affected by changes in activity opportunities and travel conditions, especially those that are affected by public policies and programs.

Figure 1: Typical Activity-Based Regional Travel Forecasting Model System

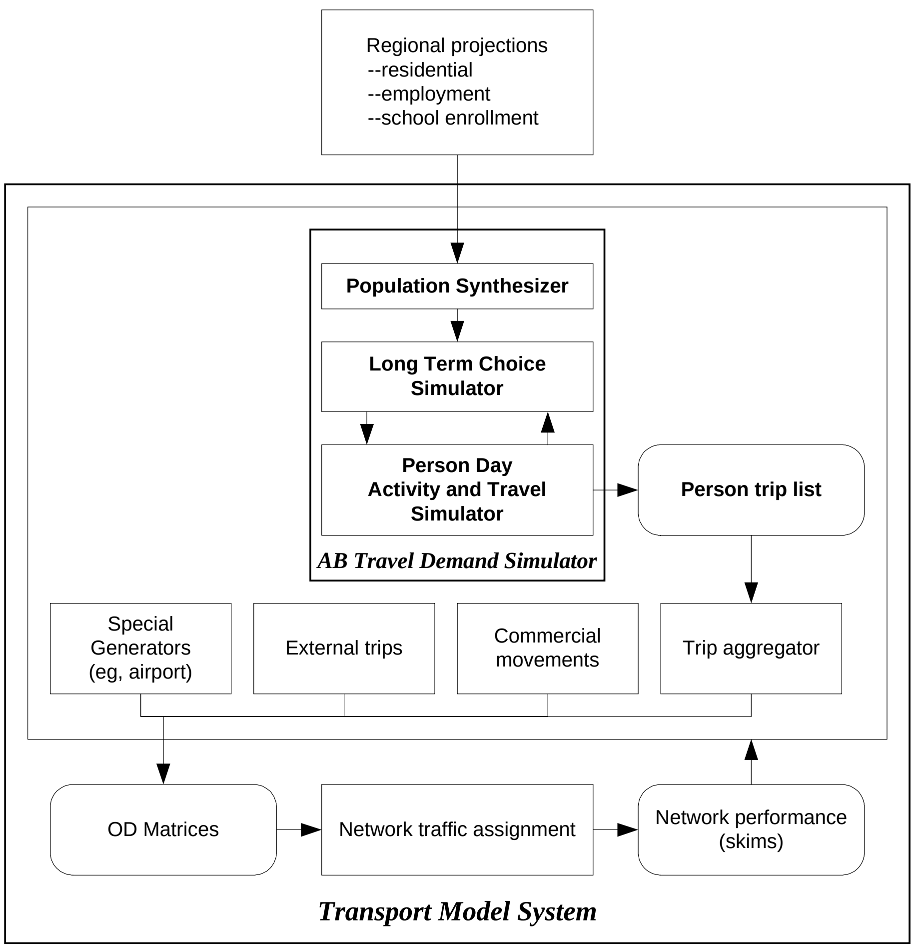

This paper does three things. First, in Section 2, it discusses principles that have guided the development of AB models now in use for travel forecasting in the

United States, focusing on principles related to the integration of systems composed of many models. Second, in Section 3, using the terminology established in the discussion of principles, it describes the specific integration techniques employed by existing AB model systems in the United States. Third, it provides a reasoned (though not empirically supported) discussion of the strengths and weaknesses of the techniques, making some preliminary judgments about their effectiveness. This critique occurs throughout Sections 2 and 3, and is summarized in Section 4 with a prioritized list of integration features. Although some model systems fare better than others in the critique, they all have substantial integration weaknesses. The purpose is not to select winners and losers. Nor is the purpose to give definitive answers regarding the best integration techniques. Rather, it is to increase the awareness of this important topic, focus the issues, and stimulate further thought, discussion and research that may lead to the development of improved integration techniques in AB models.

2. INTEGRATION PRINCIPLES AND TERMINOLOGY

We start by considering two different phenomena that affect the activities and travel a person ends up completing in any given day. The first one is the passage of time. Every new action emerges from the situation in the present created by events of the past. It is tempting to thus organize an AB model in the same way, with outcomes modeled in strict temporal sequence. Indeed, many models of activity and travel have been developed according to this organizing principle. In such a framework, newly modeled activities can take as given all outcomes that occurred earlier, and must treat as unknown all outcomes that may (or may not) occur later. This does not necessarily preclude taking into consideration the various possibilities for the future; but to do so requires some complex techniques to quantify those future possibilities so that they can be considered in the model's prediction of the present activity.

A second phenomenon is purposeful human planning. People often think ahead, schedule important future activities, and arrange other activities around them in order to achieve their objectives more effectively. It is also tempting to organize an AB model in this second way, with outcomes also modeled sequentially, but according to a plan in which more important activities are modeled first, and less important activities fill the remaining time. This does not preclude taking into consideration, at the time that a more important activity is modeled, the various possibilities for the less important activities. But to do so requires complex techniques to quantify the possibilities so that they can be considered in the model's prediction of the more important activity.

In reality, the implemented days of most people are probably the result of both planning and the passage of time. Furthermore, most data that are available for developing these models provide an observed itinerary, but little or no information about how planning and the passage of time combined to cause it.

Because of this, it is tempting to organize an AB model in the same way, representing a person's (or even an entire household's) day as a single complex outcome, simultaneously representing many components that are highly correlated because of the forces of planning and passage of time that together shaped the person's day. Indeed, because of the practical data limitation, and the importance of both planning and the passage of time, it seems better to formulate a simultaneous model than to formulate a sequential model organized only on planning or only on the passage of time.

However, a person's day (and even more so an entire household's day) is so complex that a simultaneous representation is not feasible. It is too complex to understand all at once, put into a mathematical form and carry out the needed computation. Model developers have been forced to break the outcome into pieces that can be implemented sequentially, specify a model of each of the components, specify important relationships among the components, and integrate them in an attempt to preserve the important relationships.

Without explicitly explaining the need to implement the models in a sequence, which is what gives rise to the concern about adequate integration, Vovsha et al (2005) have used the terms "downward integrity" and "upward integrity" to describe effective integration of sequentially applied models:

> "**Downward integrity** means that all lower-level decisions in the choice model hierarchy are properly conditional upon the upper-level decisions and take into account a gradually narrowed scope of lower-level choice alternatives as the upper-level choices progress....Downward integrity is ensured by properly sequencing the models, tracking the important variables from choice to choice that accurately describe the feasible scope left for each subsequent choice, and preventing conflicting choices for the same individual.
>
> "**Upward integrity** means that when modeling upper-level choices the composite measure of quality of the lower-level choices available for each upper-level alternative is properly taken into account."

These terms associate a direction of movement for the model sequence, with the beginning of the sequence at the 'top level' and the end of the sequence at the 'bottom level'. In this paper we adopt this commonly used—and entirely arbitrary—directional reference. Accordingly, we call downward and upward integrity two aspects of **vertical integration**.

However, in some cases, it may be important and feasible to retain the simultaneous modeling approach for a portion of the overall outcome consisting of two or more components. This would be the case where there is important complex correlation among component outcomes that can be correctly represented by a known and practical model structure. We call this **horizontal integration**. Defined in this way, horizontal integration is superior to vertical integration. A sequence of vertically integrated model components is a second best approach, to be used only when horizontal integration is infeasible. The efforts of most academic researchers in this field are often limited to the domain of horizontal integration. In the United States, vertical integration has primarily been the work of consultants carrying out projects with the clear objective of implementing complete model systems that can be used by a public agency for travel forecasting and policy analysis.

Since reality is so complex that full horizontal integration is infeasible, and it is necessary to use a sequential approach with vertical integration, then what organizing principle(s) might be used in choosing the sequence? In the context of a one-day itinerary, where choices are made for only one day, but are often heavily dominated by prior choices with longer term consequences—such as where to live and work, and whether to have a car for every driver in the household—the **time horizon** of decisions can serve as an organizing principle, with longer term choices coming before shorter term choices in the sequence. The principles of **temporal sequence** and **human planning**, described above, are also both good candidates. All three principles have been used in all or most of the existing US AB model systems. Increasingly, modelers have recognized the impact of long-term choices and habits on within-day behavior, specifying more of these models and placing them first in the sequence. For within-day choices, the existing model systems rely primarily on a human planning sequence, with temporal sequence used in some cases for what can be viewed as minor decisions.

Unfortunately, it has not even been feasible for modelers to capture, via vertical integration, all the apparent correlations among the components of a person's one-day itinerary. Nor has it been feasible to test enough variously specified vertically integrated models to state with confidence that the most important correlations have been correctly captured by the selected specifications. As a result, a great deal of modeler judgment undergirds the existing model systems, and will probably continue to do so for the foreseeable future. Judgment guides the modeler's choice of the components to include (and exclude), the components to keep together via horizontal integration, the specific sequence to use for the separate components, and the techniques of vertical integration to employ.

3. EXAMPLES FROM MODEL SYSTEMS NOW IN USE

This section examines specific integration techniques and features of the AB models now in use in the United States. The model systems considered include those of:

- San Francisco County Transportation Authority (SFCTA), (San Francisco County Transportation Authority and Cambridge Systematics, Inc. 2002a);

- New York MetropolitanTransportation Commission (NYMTC), (Parsons Brinckerhoff Quade & Douglas, Inc., 2005);
- Mid-Ohio (Columbus) Regional Planning Commission (MORPC), (PB Consult Inc. and Parsons Brinckerhoff, 2005); and
- Sacramento Area Council of Governments (SACOG), (Bradley, Bowman and Griesenbeck, 2006).

This list reflects the chronological order in which the model systems were developed. The subsequent discussion examines them one at a time, in the same order.

3.1. SFCTA

SFCTA has a day activity pattern model (San Francisco County Transportation Authority and Cambridge Systematics, Inc. 2002a), based on the Bowman and Ben-Akiva prototype (Bowman, 1995; Bowman and Ben-Akiva, 2001), spanning an entire day of activities and travel for one person. As shown in Figure 2, in one (horizontally integrated) model it identifies the most important on-tour activity purpose of the day, whether one or more stops is made before, during and/or after that activity on the same tour, and the presence of one or more additional tours for maintenance and/or discretionary activities during the day. Thus, this single model provides information about (a) the purpose and structure of the primary tour of the day, and (b) participation in additional tours. This enables the model to represent the total amount of tour-making during the day, and to capture trade-offs between trip-chaining on the primary tour and conducting additional tours. However, it provides no information about activity purpose, other than the purpose of the primary on-tour activity of the day.

Figure 2: The SFCTA model includes a horizontally integrated day activity pattern that encompasses tours and trip-chaining. Its upward integration is limited because it uses simple zonal measures instead of logsums that would account for differences among persons and their available travel destinations, modes and times of day.

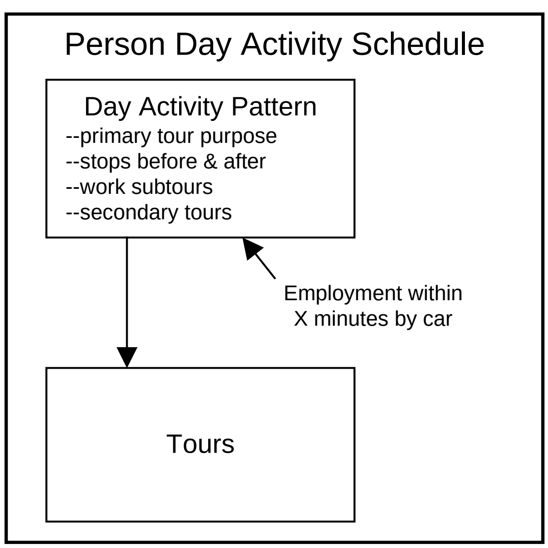

The subsequently simulated tour and stop models are conditioned by the activity pattern, providing downward integration that makes the tour models consistent with the modeled aspects of the day activity pattern. However, the model system distinguishes only five time periods during the day, so the downward integration isn't able to effectively deal with the time dimension.

The horizontal and downward integration are unable to capture the sensitivity of activity participation and tour-making (i.e. the pattern model outcome) to travel conditions such as travel time and cost. However, because the pattern model spans an entire day, it would be easy to implement upward integration that captures the effect simultaneously across all activities and tours in the day. This was a primary emphasis of Bowman and Ben-Akiva, who used logsums from the tour models to do this when they introduced their prototype. However, rather than using logsums from the tour and intermediate stop models, which would provide the best known way of capturing the accessibility effects across multiple travel modes and times, the original implementation of the SFCTA model used ad hoc accessibility measures, primarily retail and service employment within 15 minutes by car at certain times of the day. This renders the pattern model insensitive to changes in transit level of service, and to accessibility changes that affect tours and intermediate stops unequally.

SFCTA includes long-term choices of usual work location for each worker and a downwardly integrated household vehicle ownership model (see Figure 3). These are downwardly integrated with the day activity schedule of all household members. Also, for each person the usual work location model is upwardly integrated via a tour mode choice logsum, and the auto ownership model includes non-logsum zonal accessibility measures that are sensitive to auto and transit travel times.

Figure 3: SFCTA integrates long-term choices of usual work location and household vehicle ownership with the day activity schedule. The usual work location model is upwardly integrated via a tour mode choice logsum, and the auto ownership model includes non-logsum zonal accessibility measures that are sensitive to auto and transit travel times.

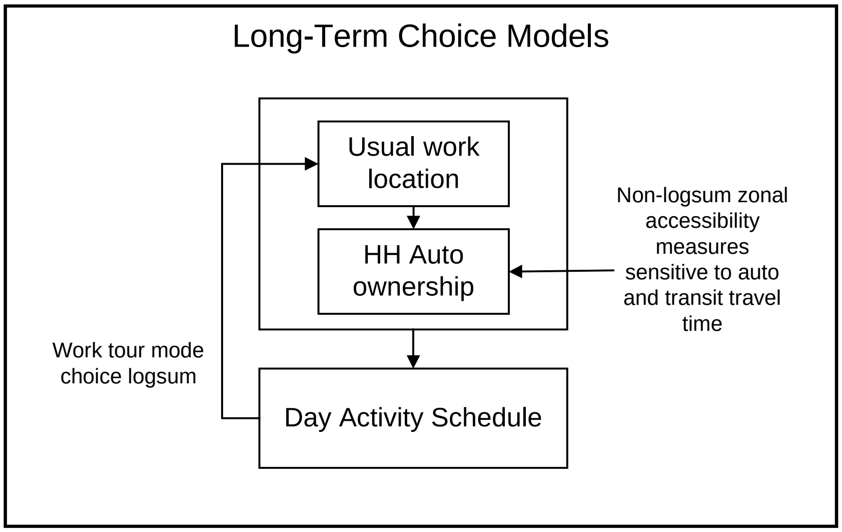

3.2. NYMTC

NYMTC has no day activity pattern that horizontally integrates the representation of a person's day. Instead, as shown in Figure 4, it uses a sequence (or cascade) of thirteen tour generation models, each with a distinct combination of person type (children, non-workers and workers) and purpose (school, university, work, at-work, maintenance, and discretionary). Each subsequent model can take into consideration the outcome of the prior tour generation models. Thus, tour generation for later person-type-purpose combinations is affected by variables indicating generation of tours for earlier-simulated person-type-purpose combinations. Because there is no horizontally integrated activity pattern, there is no good way for earlier person-type-purposes to be affected by the results of those later in the sequence. For the same reason, the overall activity and travel agenda of the day cannot adjust—via upward integrity mechanisms—to changes in travel conditions (Figure 5). Furthermore, none of the thirteen individual NYMTC tour generation models is sensitive to changes in auto travel conditions, and only discretionary tour generation of non-workers is sensitive to the transit travel conditions. Even if the individual models were more sensitive to travel conditions, the cascade approach provides no good way of providing upward integrity that enables the earlier models to compensate for changes that would come indirectly through the lower level tour decisions.

Figure 4: NYMTC uses a sequence of downwardly integrated tour generation models to capture interactions among household members in tour generation. There is no horizontal or upward integration to capture non-hierarchical correlations.

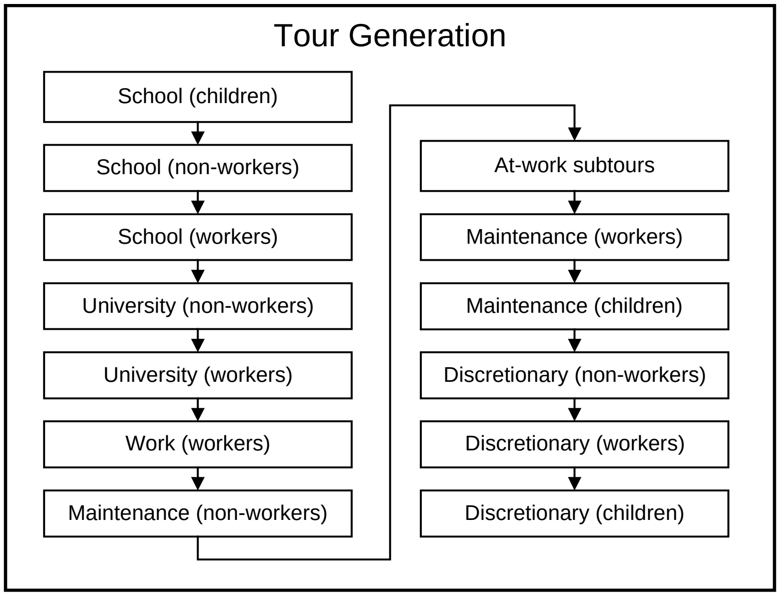

Figure 5: NYMTC lacks integration to make tour generation sensitive to travel conditions

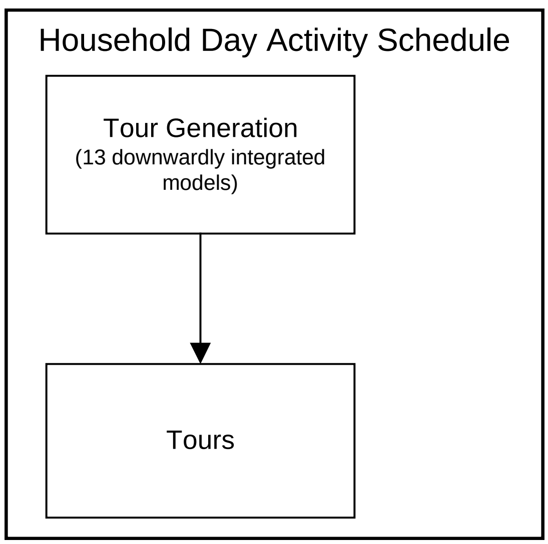

NYMTC includes a long-term model for auto ownership that is downwardly integrated with the rest of the model system, and includes zonal accessibility measures that are sensitive to auto, transit and walk travel times.

The NYMTC model and the SFCTA model employ two contrasting approaches that highlight a major trade-off faced by AB modelers. The trade-off is between emphasizing upward integration and emphasizing integration among household members; it is difficult to achieve both simultaneously. Sacrificing upward integrity weakens the ability of the model to be accurately sensitive to changes in transport conditions, especially at the highest levels of the model system, where tour generation is modeled. Sacrificing intra-household integration weakens the ability of the model to accurately represent the joint behavior of household members. SFCTA favors upward integration, whereas NYMTC favors household integration.

3.3. MORPC

Like NYMTC, MORPC relies heavily on downward integration of a cascade of models to represent the tours and trips of persons in a household. Compared to NYMTC, it has a substantially more complex sequence of models that attempts to more realistically represent interactions among household members. The focus of attention is on downward integration. Horizontal and vertical integration are not emphasized. The model sequence is broken into several subsequences including, in order, those for 'linked activity patterns', 'mandatory' tour details, non-mandatory tour generation (including joint tour, household maintenance tours and individual discretionary tours), non-mandatory tour details, and intermediate stop generation and details (see Figure 6).

Figure 6: MORPC uses a sequence of downwardly integrated tour generation and tour model subsequences that is more complex than NYMTC's in order to capture more interactions among household members in tour generation, including linked activity patterns, joint tours and household maintenance tours. Downward integration among tours prevents conflicting time-of-day results for tours, within and across the schedules of household members. There is no horizontal or upward integration among the subsequences to capture non-hierarchical correlations among them.

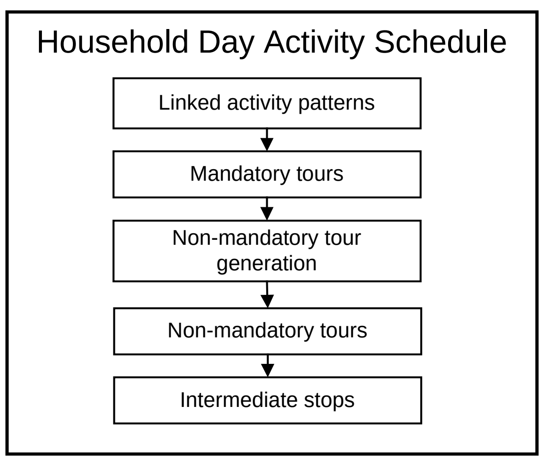

As shown in Figure 7, each model in the 'linked activity pattern' sequence (Figure 7a) predicts one of a few basic pattern types (Figure 7b) for one person in the household: whether they travel for work or school, for only other purposes, or stay at home all day, given the same type of prediction for household members earlier in the sequence (PB Consult Inc. and Parsons Brinckerhoff, 2002). The downward integration among these models allows them to partially capture strong observed correlations in primary activity purpose among household members. In particular, it captures the tendency for multiple people in the household on any given day to together stay at home, or to not go to work or school. The sequential nature of the model makes it infeasible to enable early models in the sequence to compensate for changes that affect the likelihood that lower level persons in the sequence will stay home or not go to work or school. Also, for each person in the household, this isolates the modeling of their primary activity purpose from the modeling of their other (subsequently modeled) activities of the day. This makes it impossible to use a horizontally integrated model to (simultaneously) capture the effect of travel conditions on the content and structure of their entire day. It would not preclude using upward integration—via logsums—to capture the effect of travel conditions on each individual's choice of primary activity purpose. However, the MORPC model relies instead on zonal accessibility indices that are mode- and time-of-day specific. This makes them insensitive to significant differences across households, such as income and car ownership. It also makes each accessibility index sensitive only to policies that affect a single mode and time-of-day. Because of high correlation across modes and times-of-day, it is usually possible to include only one index in a model. As a result, the model is sensitive only to one mode and time of day. In this sequence of the MORPC models, the upward integration makes the models sensitive only to walk access to jobs (for work-tour patterns) and walk access to retail (for non-work patterns). Since walk access is a function of distance, these measures (and hence the linked activity pattern) are sensitive to the distribution of employment, but insensitive to transport conditions.

Figure 7: MORPC's linked activity patterns consist of a sequence of downwardly integrated activity pattern type models, one per person in the household, that capture the tendency for multiple people to stay home together on the same day (Figure 7a). Each person's activity pattern (Figure 7b) represents the purpose of the main tour(s) of the day, but not all tour purposes or any intermediate stops. There is no upward integration from subsequent tour or stop models, nor is there upward integration that makes them sensitive to transit or auto travel conditions.

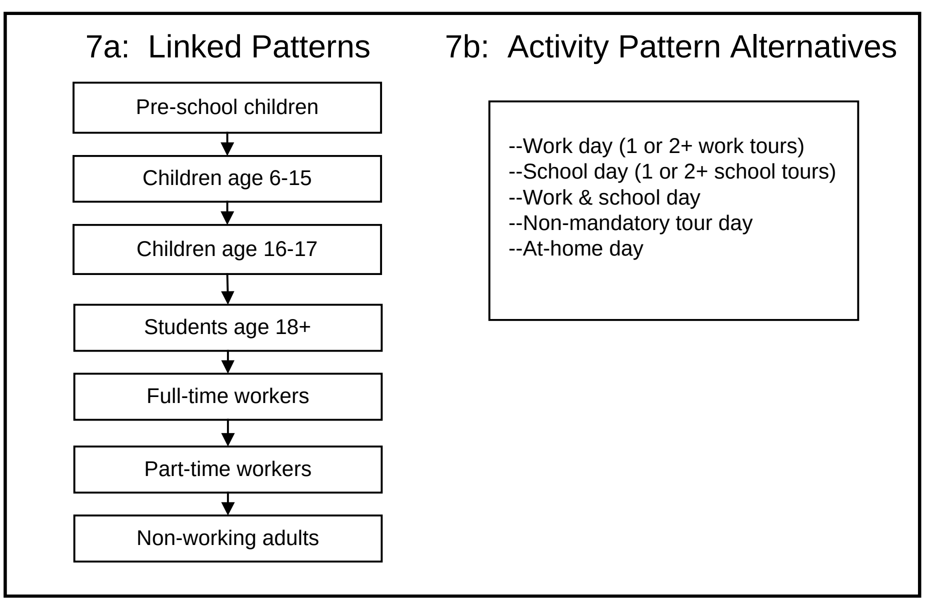

The ‘mandatory’ tour sequence simulates the destination, time of day and travel mode (in that order) for any work, university or school tours (called ‘mandatory’ tours by the developers) prescribed above for household members (see Figure 8). Through downward integration to the subsequent models, this allows the models of tour generation for maintenance and discretionary tours to be conditioned by the amount of time persons devote to their mandatory tours. There is no upward integration from the subsequent models to these detailed models of mandatory tours.

Figure 8: For each tour, MORPC uses a vertically integrated sequence of destination, time of day (TOD), and mode choice. Mode choice logsums upwardly integrate mode choice with TOD and destination choice, but destination choice uses an assumed TOD rather than a TOD logsum. The sensitivity of destination choice is thus limited to effects that occur during the assumed TOD.

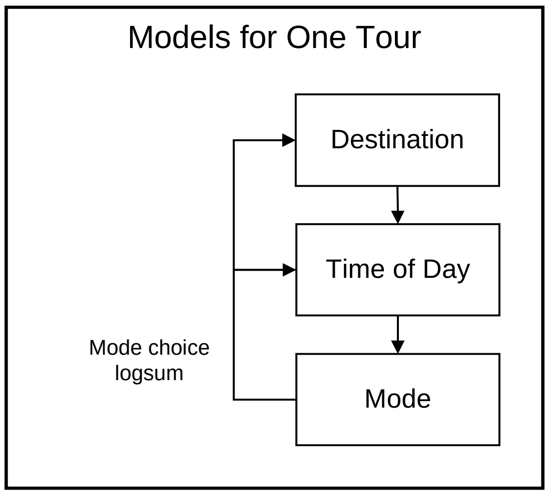

Within the sequence of destination, time and mode models, there is downward integration that prevents unrealistic times to be chosen for a predicted destination, and prevents unrealistic mode for a predicted destination and time. There is also upward integration from mode to time, and from mode to destination. Upward integration is not implemented from time-of-day to destination. Rather, the destination choice model has mode choice logsums that are specific to a particular time period combination (for tour beginning and end); the time period used for the logsum depends on the purpose and other already modeled aspects of the day. This makes the destination choice model sensitive to changes in transport conditions that vary by time of day, but in a biased way, capturing sensitivity for the typical time of day, but not for the other times of day.

The model sequence that generates joint, maintenance and discretionary tours is depicted in Figure 9. The joint tour generation subsequence generates tours taken jointly by two or more household members, and then identifies the household members who participate in each joint tour. The horizontal integration of this choice among household members is consistent with the fact that some tours are indeed joint tours reflecting a mutual decision to conduct a tour together for the same purpose. This group of models has downward integration from the mandatory models, and internally, taking into consideration the time required of the various household members for their mandatory tours. There is no upward integration within this model subsequence, from subsequent models, or from the assignment models; therefore, generation of joint tours is insensitive to transport level of service.

Figure 9: MORPC generates joint tours using a model that horizontally integrates this outcome for all members of the household and subsequently assigns household members to the tours. A similar approach is used for household maintenance tours. Then discretionary tours are generated sequentially for each household member. There is no upward integration to joint tour generation, so it is not sensitive to transport conditions. For maintenance and discretionary tour generation, the upward integration comes from zonal accessibility indices that are mode and time-specific; most models are sensitive to at most one mode (auto, transit or walk) and time of day, which biases their policy sensitivity.

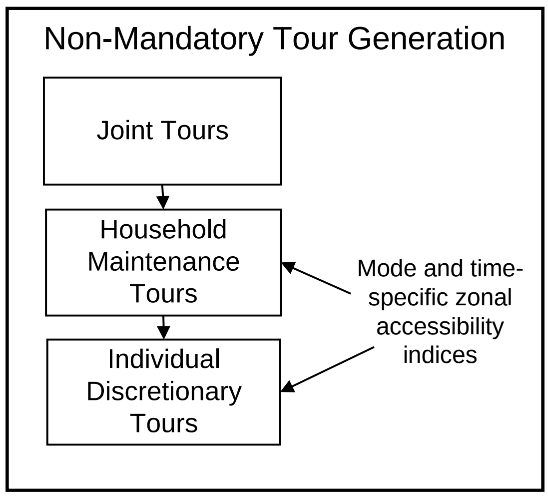

The maintenance tour generation and allocation subsequence has a tour generation model that horizontally integrates the generation of joint tours among all members of the household. The horizontal integration of maintenance tour generation among household members is consistent with the hypothesis that the number of maintenance tours is primarily a household decision for the purpose of achieving household objectives. This is followed by downwardly integrated models that assign the maintenance tours to individual household members. The generation and assignment of maintenance tours to household members is conditioned by the work, university and school tour obligations modeled above. There is no upward integration to capture the effect, on maintenance tour generation and assignment, of an individual’s propensity for discretionary tours. There is upward integration to this generation model using zonal accessibility indices that are mode- and time-of-day specific as described above for linked activity patterns. Whether auto, transit or walk index is used varies from purpose to purpose.

Next comes a subsequence of discretionary tour generation models for each person in the household. For each person, this is conditioned by their prior modeled obligations for work, school, university, joint and maintenance tours, as well as the at-home-all-day status of other household members. Upward integration is handled like it is for maintenance tour generation.

After the generation of all joint, maintenance and discretionary tours, a sequence of models simulates their destinations, times-of-day and modes, one tour at a time, just like the sequence for mandatory tours. Downward integration limits destination and time-of-day choices according to the time already taken for previously simulated tours. This is accomplished by maintaining a set of 19 one-hour time blocks for each person, marking as unavailable those that are occupied by a newly modeled tour, and checking time slot availability to determine the choice set whenever a tour time-of-day is modeled.

The modeling of participation, location and trip mode of intermediate stops occurs on a tour-by-tour basis, but only after the generation, destination, mode and timing of all tours by all household members. The stop models are restricted by time commitments for all modeled tours, and for prior-modeled stops. There is no upward integration from the stop models to the tour models. Therefore, tour models are unable to capture trade-offs between conducting additional tours for maintenance and discretionary activities, on the one hand, and conducting those activities as intermediate stops, on the other hand.

MORPC includes a long-term model for auto ownership that is downwardly integrated with the rest of the model system. It includes zonal accessibility logsums for walk and transit, making the model sensitive to transit and walk travel times, but insensitive to auto travel times.

In summary, MORPC emphasizes aspects of integration that achieve two major improvements over SFCTA and MORPC. The first emphasis is horizontal and downward integration that improves the realism of joint household outcomes and the consistency of individual schedules among household members. The second emphasis is on more detailed time of day modeling, with accompanying downward integration, that improves the consistency of time-of-day outcomes within an individual's day and across household members. This comes at the expense of upward integration which makes MORPC weaker in modeling sensitivity of the upper level models, especially tour generation and trip-chaining, to travel conditions.

3.4. SACOG

Like SFCTA, SACOG has a day activity pattern model spanning an entire day (Bowman and Bradley, 2006). It identifies the participation in one or more tours for each of seven purposes, and the participation in one or more additional stops (throughout the day) for each of the same seven purposes. This single model provides a complete inventory of the activity purposes in a day and, for each purpose, whether it is the primary objective of a tour and/or a supplemental objective on a tour. This enables the model to realistically capture the mix of tours and stops, by purpose, throughout the day. The utility of a tour or stop for one purpose directly affects its probability, as well as the probability of tours and stops for all other purposes. However, the model does not provide information on the exact number and purpose of the stops on a specific tour, or the positions of the stops on the tour; this is left for models later in the sequence.

Figure 10: The SACOG model includes a horizontally integrated day activity pattern that encompasses tours and trip-chaining for seven purposes. Its upward integration uses three kinds of logsums that account for differences among persons and their available travel destinations and modes, but not times of day.

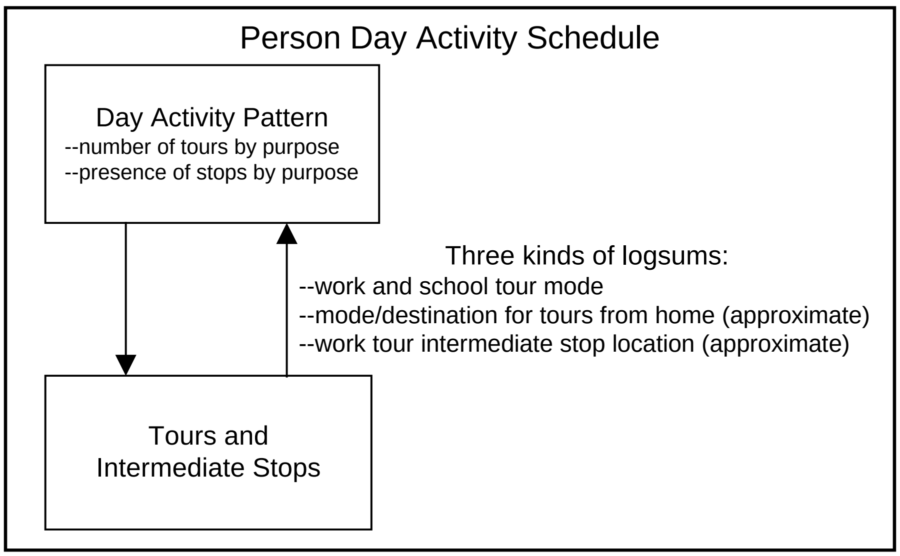

As with the Bowman and Ben-Akiva prototype, the subsequently simulated tour and stop models are conditioned by the activity pattern, providing downward integration that makes the tour models consistent with the modeled aspects of the day activity pattern.

Also as with the Bowman and Ben-Akiva prototype, it is easy to implement upward integration that captures the effect of tour accessibility simultaneously on the overall pattern of activities and tours in the day. In contrast to the SFCTA model, the SACOG model does this with logsum accessibility variables for tours and intermediate stops, capturing several types of composite accessibility effects on the pattern, including mode choice logsums for tours to the usual work and school locations, approximate mode/destination logsums for tours from home, and approximate location choice logsums for intermediate stops on work tours.

The approximate, or aggregate, logsum is calculated in the same basic way as a true logsum, by calculating the utility of multiple alternatives, and then taking expectation across the alternatives by calculating the log of the sum of the exponentiated utilities. However, the amount of computation is reduced, either by ignoring some differences among decisionmakers, or by calculating utility for a carefully chosen subset or aggregation of the available alternatives. The approximate logsum is pre-calculated and used by several of the model components, and can be re-used for many persons. This makes it computationally feasible to use logsums at the upper levels of the model system. The categories of decisionmakers and the aggregation of alternatives are chosen so that in all choice cases an approximate logsum is available that closely approximates the true logsum. In essence, this is a sophisticated ad hoc measure that is intended to achieve most of the realism of the true logsum at a small fraction of the cost.

The approximate tour mode-destination choice logsum is used in situations where information is needed about accessibility to activity opportunities in all surrounding locations by all available transport modes at all times of day. Because of the large amount of computation required for calculating a true logsum for all feasible combinations in these three dimensions, an approximate logsum is used with several simplifications. First, it ignores socio-demographic characteristics, except for car availability. Second, it uses aggregate distance bands for transit walk access. Third, sometimes it uses a logsum for a composite or most likely purpose instead of calculating it across a full set of specific purposes. Finally, instead of basing the logsum on the exact available time window of the choice situation, and calculating it across all of the available time period combinations within the window, it uses a particular available time window size and time period combination. With these simplifications, it is possible to pre-calculate a relatively small number of logsums for each zone, and use them when needed at any point in the simulation of any person's day activity schedule.

The approximate intermediate stop location choice logsum is used in the activity pattern models, where accessibility for making intermediate stops affects whether the pattern will include intermediate stops on tours, and how many. Four logsums are calculated for each OD zone pair, distinguished by tour mode (transit or auto) and time of day (peak or offpeak). Each logsum is calculated across all possible intermediate stop zones, each stop's utility is a function of travel time and zonal attractiveness, and zonal attractiveness is a function of employment and school enrollment, taken from an estimated purpose-non-specific location choice model.

Although the upward integration of the activity pattern is better than SFCTA's, it is still limited. The approximate logsums are limited as described above. In addition, most of the non-work purposes do not have mode/dest logsums, so some of the benefit of the purpose-specific specification is lost in the upward integration.

For each tour there is a vertically integrated sequence of destination, mode and time-of-day models (see Figure 11). Downward integration prevents unrealistic destination, mode and time-of-day combinations. There is also upward integration from mode to destination, but it does not use time-of-day logsums. Rather, a simulation technique was implemented to make the mode and destination choice models sensitive to policy effects that may vary by time of day. The basic idea is to avoid the use of a logsum (and its associated computational costs) when applying an upper level model by treating as given a conditional outcome that is not known, and would otherwise require the calculation of a logsum from all possible conditional outcomes. In this case the assumed conditional outcome is the tour time-of-day. It is selected by a Monte Carlo draw using approximate probabilities for the conditional outcome. Rather than making every simulated outcome sensitive to variability in the conditional outcome, sensitivity is achieved across the population through the variability of outcome in the Monte Carlo draws. In this way, the mode and destination choice models are sensitive to variations in transport level of service and spatial attributes across all possible combinations of time-of-day, with the affects approximately weighted by the time-of-day choice probabilities.

Figure 11: For each tour, SACOG uses a vertically integrated sequence of destination, mode choice, and time of day (TOD). Mode choice upwardly integrates with destination choice; it uses a simulated TOD that makes the destination and mode choice models sensitive to changes in transport conditions that vary by time of day.

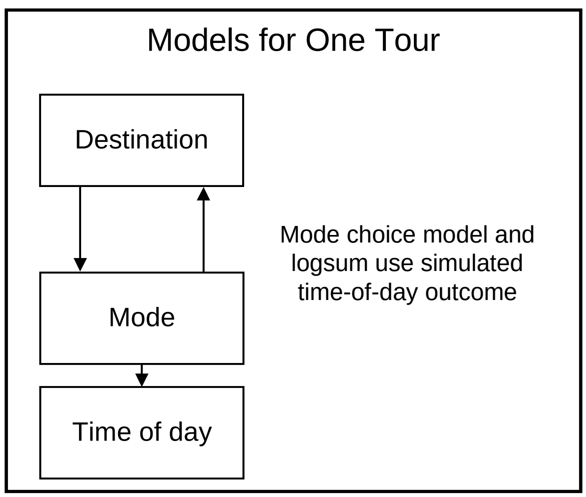

SACOG conditions half-tour stop participation and purpose upon the modeled aspects of the day activity pattern, prior modeled tours, current tour, and—in the case of the second half-tour of the tour—the first half-tour (Figure 12). This enables consistency of stop participation by purpose among models at all levels, and complements the horizontal and upward integrity of stop purpose enabled by the SACOG specification of the activity pattern model.

The SACOG intermediate stop models of location, mode and timing are conditioned by the same models as the half-tour model. In addition, each one is conditioned by the half-tour model that predicts the number and purpose of stops on the half-tour, as well as all prior-modeled stop outcomes on the half-tour. The intermediate stops are simulated in sequence emanating from the tour's primary destination, in reverse chronological order for stops before the tour destination, and in chronological order for stops after the tour destination. This is based on the assumed importance of arriving at and departing from the primary destination at a previously modeled time. Accordingly the timing of each intermediate stop is conditioned by the timing of all prior-modeled stops. Similarly, the intermediate stop location choice is conditioned by all previously modeled locations, and the trip mode is conditioned by the mode set allowed by the modeled tour mode and the modes used on all prior modeled trips on the tour. This downward integration allows a consistent and feasible representation of each tour's entire travel itinerary, with regard to timing, mode and location. This is complemented by upward integrity provided by the use of intermediate stop location choice logsums in the half-tour and activity pattern models. These capture the effect of accessibility on the stop participation choices modeled in those two models.

Figure 12: Downward integration of half-tour and intermediate stop models conditioned by the pattern, tour, half-tour and prior-modeled stop models, with accompanying upward integration via approximate stop location logsums.

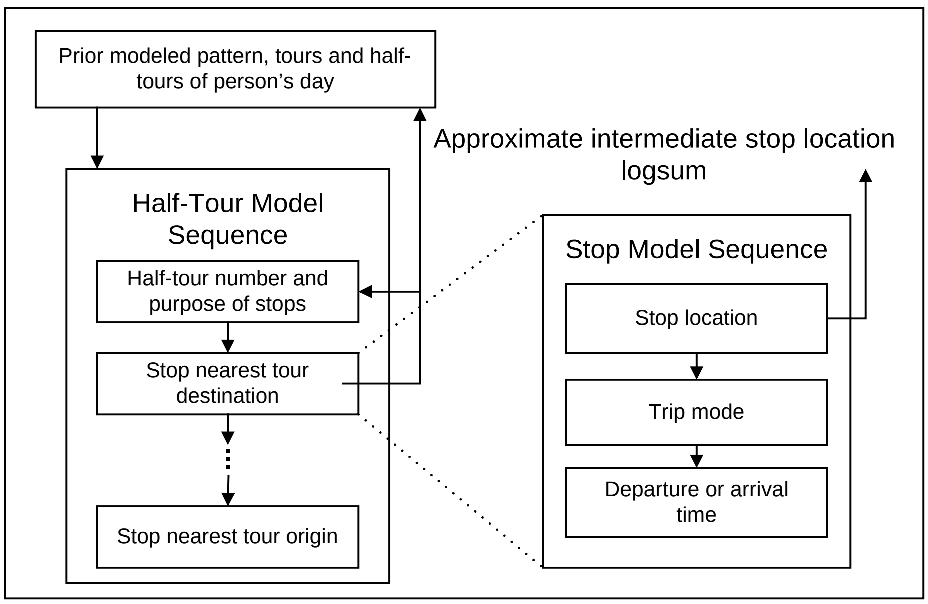

SACOG includes long-term choices of usual work location for each worker and usual school location for each student. For young student workers, the usual school location model conditions the usual work location model, and for old worker students the sequence is reversed. These condition a downwardly integrated household vehicle ownership model, and all of them are downwardly integrated with the day activity schedule of all household members. All three long-term models are upwardly integrated via work and/or school tour mode choice logsums, as well as approximate mode-destination logsums. For the usual location models the mode/dest logsums measure accessibility for tours from the usual location. For auto ownership, they measure accessibility for tours from home.

In summary, SACOG emphasizes and implements techniques for achieving upward integration that improves the model system's ability to accurately capture sensitivity to travel conditions at all levels of the model system, and especially at the upper levels. It also implements more complete downward integrity of the model system with regard to the modeling and accounting of participation, time-of-day, mode and location of intermediate (chained) stops on all tours of an individual in a day. The big weakness of the SACOG model system, relative to MORPC and NYMTC, is that it excludes explicit integration of the day activity schedule model components across household members.

4. DISCUSSION

In this section we shift from examining the integration features of particular model systems to identify a set of integration features that spans the features seen in the reviewed models and also includes a few features that none of the models has yet incorporated. Table 1 lists detailed features within seven major categories, identifying for each feature the model systems that include it. Although all the categories are important, some seem more crucial than others, so we list them in a suggested priority order.

1 **Integration among destination, mode and time of a tour**. Effectively integrating the models of mode and destination choice has long been an important objective in trip-based models. One of the great desired strengths of the activity-based framework is the realistic incorporation of time-of-day modeling into the model system. This expands the focus of tour model integration to include time-of-day choice with mode and destination. It is important to move beyond the use of fixed assumed times in the calculation of mode and destination choice models, and in the calculation of logsums that are used as accessibility measures in higher level models. Doing so will enable the models to more realistically capture how time-specific policies affect choices other than time-of-day, such as mode, destination, and tour generation. In summary, this **is important because it helps the model capture the effects of time-specific policies on all dimensions of choice**.

2 **A horizontally integrated person-day activity pattern model with purpose-specific information about the tours and intermediate stops in the day.** This horizontal integration helps the model system to realistically represent the total amount and mix of activity and travel carried out by a person in one day. Importantly, it is a pre-requisite for additional features that improve the realism of these model predictions. In particular, it allows effective upward integration from purpose-specific tour and stop models. This **allows changes in transport conditions to affect overall tour and stop generation, as well as trade-offs between tour and stop generation**. Including activity purpose information in the pattern model enables it to capture the impact of purpose-specific changes in accessibility on pattern choice. For example, improved work accessibility can have different impact than improved accessibility for shopping.

Table 1: Integration Features of Existing Activity-Based Model Systems
(a 'y' indicates that a model system has a particular feature)

Integration Feature	Priority	SFCTA	NYMTC	MORPC	SACOG
Integration among destination, mode and time of a tour					
—downward	1	y	y	y	y
—upward using assumed times		y	y	y	y
—upward accounting for available times	1				y
Horizontal integration of tour and stop generation in a day **for a person**					
—downward integration among tours			y	y	y
—downward among tours and stops					y
—horizontal among tours	2	y			y
—horizontal among tours and stops simultaneously	2	y			y
—horizontal purpose-specific among tours and stops	2				y
Upward integration: Transport conditions influence **tour generation**	3				
—upward integration in very few cases				y	
—upward integration in most or all cases	3	y			y
—accounting for differences among persons, available destinations and available modes	3				y
—accounting for available times	3				
Upward integration: Transport conditions influence **generation of trip chains**	3	y			y
—via intermediate stop logsums	3				y
—accounting for differences among persons and available destinations	3				y
—accounting for available modes and times	3				
Downward integration of tour and stop details in a day	4	y	y	y	y
—accounting for time used on mandatory tours	4			y	y
—accounting for time used on all trips and tours	4				y
—accounting for stop purposes	4				y
Integration between long-term and within-day models	5				
—downward from long-term to within-day models	5	y	y	y	y
—included auto ownership model	5	y	y	y	y
—included usual work location model	5	y			y
—upward to long-term models via mode choice logsums	5	y			y
—included usual school location model	5				y
Downward integration related to the use of vehicles from auto ownership model through tour and stop models	6				
—accounting for the use of each household vehicle	6				
—with vehicle type in long-term and mode choice models	6				
Integration of tours and stops in a day **for a household**	7				
—upward integration: transport conditions influence all modeled tour and stop generation	7				
—horizontal among persons for joint tours	7			y	
—horizontal among persons for day pattern	7				
—horizontal among persons for maintenance stops	7			y	
—among persons for escorting and shared trips	7				
—downward among persons for staying at home all day	7		y	y	
—downward among persons for additional aspects	7				

3 **Upward integration to the activity pattern model from both tour and intermediate stop models via logsums that account for important differences among persons (especially income, car availability and driving age) and that also account for available modes, destinations and times of day.** Accounting in logsums for available modes, destinations and times of day makes the pattern model sensitive to policies that apply differently across these three dimensions. For example, with this type of upward integration, a peak period toll on a major commute corridor will have an effect on pattern choice, as will a peak period transit improvement or a highway improvement that affects all time periods; the nature and magnitude of the effects will be governed by the probability proportions of the models used to calculate the logsums. Accounting for important differences among persons enables the model to capture changing aggregate effects when population demographics and long term choices change, and helps the model to more accurately capture differences in how policies affect various population segments. For example, highway improvements help households with cars more than they help households without cars. Accounting for tour and intermediate stop accessibility enables the pattern model to better capture trade-offs between tour and intermediate stop generation. For example, improvements along a person's commute route should increase the likelihood that they will make intermediate stops on their way to or from work, whereas improvements away from their commute route should instead increase their likelihood of making additional tours. In summary, this **makes the sensitivity of pattern level changes to transport conditions and demographic changes more realistic**.

4 **Extensive and consistent downward integration of the within-day models: from activity pattern to tour models and then to all intermediate stops within each tour, including participation, purpose, location, travel mode and timing of all travel.** We think that a model sequence proceeding from day to tour to intermediate stop provides a reasonably realistic modeling sequence based on purposeful human planning. For example, it seems realistic to assume that the exact number and location of intermediate stops on a work tour should depend directly on the work location and the main mode used to get to work, even if the stops occur on the way to work. Careful downward integration helps assure that all modeled aspects of the entire modeled outcome are mutually consistent. For example, the use of time window accounting for each person can prevent the model from scheduling two of their tours to occur at the same time. Similarly, time window accounting for each household vehicle could be used to prevent the model from having two different household members use auto drive mode for separate tours if the household has only one vehicle. Extending the downward integration to include all intermediate stops on all tours would also yield travel itineraries that might eventually integrate effectively with traffic simulation models because of their realism. Finally, downward integration of purpose-specific choices can enable the model system to capture the correlation between activity choice and travel

conditions, when combined with purpose-specific upper model levels and upward integration of purpose-specific tour and trip effects into the upper levels of the model. Such features can move the model system closer to being truly activity-based. This type of accounting does not directly improve the policy responsiveness of the model system. However, it **imposes constraints that should improve the accuracy of the policy responsive aspects of the model system, in particular sensitivity to transport conditions** that cause people to travel more or less, or to change travel modes, destinations and/or times of day.

5 **Integration that conditions within-day choices upon long-term choices (downward), and incorporates the effects of short-term opportunities and conditions on the long-term choices (upward).** The long-term choices include residential location, work location, school location, auto ownership, and possibly also vehicle type, usual mode to work, usual mode to school, and transit pass ownership. An important benefit of conditioning within-day choices upon long-term choices is that the day activity pattern and other within-day models can directly use information related to the long-term outcomes. For example, the person activity pattern model can be influenced by the mode choice logsum associated with a tour to the usual work place. As another example, the location choice model of a household adult with an escort tour or stop can assign higher probability to the usual work and school locations of other household members. Thus, the use of long-term models with downward integration to the within-day models can make the within-day models more realistically policy responsive. In other words, the short-term elasticities of the model should be more realistic. There is a danger in this, however. Conditioning the day activity schedule on long-term outcomes without making the long-term outcomes correctly responsive to policy changes could severely bias the model predictions. This is because some long-term decisions, such as usual mode to work, can be much more elastic to travel conditions than their within-day choice of how to get to work today. Thus, any modeling of long-term choices should be accompanied not only by downward integration, but also by upward integration that captures the effect of travel conditions on the long-term choice itself, yielding accurate long-term elasticities. Doing this **should enable the model system to more realistically distinguish between long-term and short-term responses to policies**. To represent short-term responses to a policy scenario, the long-term model outcomes can be held fixed, and to represent long-term responses, they can be allowed to change in response to the policy.

6 **Downward integration related to the use of vehicles, from auto ownership model down through tour and stop models.** This feature, which is not included in any of the reviewed model systems, **should enhance the ability to forecast air quality impacts and fuel consumption of alternative future scenarios**. To achieve this benefit, the vehicle type would need to be modeled along with vehicle ownership. Vehicle types would need to be defined so as to be useful for policy analysis,

while at the same time distinguishing types that represent realistic differences that matter when households acquire vehicles. The model system would also need to include the choice of tour vehicle for each auto driver tour and a full accounting of household vehicle use by time of day would need to occur, enforcing time-space constraints on every vehicle in the fleet. Doing this would enable a specific vehicle to be assigned to each vehicle trip, significantly improving the ability to provide information for air quality analysis.

7 **Integration of activities, tours and stops in a day for a household.** This category includes horizontal and downward integration for joint and correlated outcomes, such as coordinated day patterns, joint tours, household maintenance tours, and shared trips such as escort trips. Importantly, it also includes vertical integration, so that the benefits of upward vertical integration identified in features one through five above are preserved.

A horizontally integrated model of household joint tour generation would capture the tendency of persons in a household to conduct activities and associated travel together. A potential advantage is that the conditional models of destination, mode and timing for joint tours might differ from those for individual tours, and the generation of joint tours might be more or less sensitive to transport conditions than the generation of individual tours. In order to achieve these benefits, the joint tour generation model would need to be effectively downwardly integrated with the person models, so that each person's day activity pattern would include the joint tour and their individual tours would not conflict with the joint tours. Just as importantly, the joint tour generation model would need to be upwardly integrated via logsums from the tour mode, destination and timing models, so that it would be realistically sensitive to transport conditions.

A horizontally integrated household day pattern would simultaneously represent the major activity and travel choices of the day for all members of a household, especially whether they traveled at all during the day, whether they traveled to work or school, and perhaps whether they worked at home. It could thus naturally extend the day pattern approach to encompass the entire household. Starting at the top, the basic within-day hierarchy would become household-day > person-day > tour > trip. The most tangible advantage of integrating the person-day models in this way is that it would yield more realistic household day patterns, capturing tendencies for persons in a household to coordinate their schedules. For example, in two-worker households without children, workers might be inclined to work on the same days, whereas in two-worker households with children they might be inclined to not work on the same days, or to stay home when a child stays home. However, to our knowledge, there is no current evidence that modeling the household day pattern would make the model system more accurately sensitive to transport prices and policies. Furthermore, using the household day pattern prevents all the major aspects of pattern choice for a given person from being horizontally integrated in the person day activity pattern. For example, the choice of whether to conduct a work tour is separated from the choices related to tours for other purposes. This has the potential of reducing the realism of the activity pattern's response to changes in accessibility. Therefore it is not clear whether the household day pattern would improve the overall model performance. In order to maximize the benefits of the household-day horizontal integration and minimize the problems caused by breaking the person-day horizontal integration it would be important to include upward integration to the household day pattern model from tour models via logsums that account for important differences among persons (especially income, car availability and driving age) and that also account for available modes, destinations and times of day (see discussion above related to upward integration to the person-day activity pattern model). This would help the household day pattern model more realistically respond to transport conditions.

Downward integration that captures correlations in activity and travel decisions among household members would involve conditioning day activity pattern, tour and trip choices of a person in a household upon earlier modeled outcomes of other persons in the household. It would depend heavily on the choice of run order of the model components among household members. The benefit is that it would increase the consistency of the modeled outcomes among the household members, capturing natural intra-household correlations in those outcomes. However, it is not clear that this would improve the accuracy of the model system's policy responsiveness, because the coordination of activity and travel choices (such as staying home with a sick child) may not be sensitive to transport conditions. Furthermore, vertically integrating among household members may make it more difficult to accurately model policy responsiveness of each person's activity and travel choices.

Integrating households in the models has been strongly advocated by academics because of the significant influence of the household on individual behavior. A clear benefit of this type of integration is that the predicted itineraries of household members would be much more realistic when viewed from the household perspective. For example, without this type of integration, a household with young children could easily be predicted to send all the adults to work and leave a small child at home alone. Or in a household with two adults and one car, one of them might be predicted to drive alone to work, and the other to ride as a passenger to work. Other potential benefits of explicit household integration in the models would occur if the integration caused the model to behave differently, and more accurately, to changes in demographics, land use or travel conditions. This would be the case if, for example, increasing real estate prices and declining incomes lead to more large non-family households, and modeling the activity and travel of such a household as a unit yields more or less travel than modeling the activities of its members separately.

However, we place the explicit integration of household behavior at the bottom of the priority list for several reasons. First, we think that the benefit of more realistic sensitivity to travel conditions brought by effective upward integration is of utmost importance. Second, introducing explicit household interactions makes it more difficult to implement the needed upward integration. Third, even if such upward integration could be implemented in a model with extensive household integration, it is not clear that doing so would substantially improve the quality over a well-integrated individual model that includes household effects indirectly through the use of household characteristics in the model equations. In summary, we think that the implementation of household integration **is important because of its ability to improve the realism of predicted household schedules and its potential to improve the model's responsiveness to changes in demographics, land use or travel conditions, but it should only be implemented without sacrificing the effective integration of features we identify as higher priorities.**

In this paper we have discussed integration principles that have guided the development of AB models now in use in the United States, described and critiqued the integration techniques employed by these models, and listed the techniques that we think are important, in order of importance, based on our experience and judgment about how the techniques would work in real-world model systems.

We reiterate that our purpose is not to select winners and losers among the reviewed models. Each one has made important innovative contributions to the state of the knowledge and practice in travel demand modeling. Rather, we hope that this discussion increases the awareness of the important topic of model system integration, focuses the issues, and stimulates further thought, discussion and research that may lead to the development of improved integration techniques in AB models.

Bibliography

Bowman, John. L. (1995) Activity based travel demand model system with daily activity schedules, Master of Science Thesis in Transportation, Massachusetts Institute of Technology, 92 pages. Available at http://jbowman.net.

Bowman, J. L., Ben-Akiva, M. (2001) Activity-based disaggregate travel demand model system with activity schedules, *Transportation Research Part A*, **35** 1-28. Available at http://jbowman.net.

Bowman, John L. and Mark Bradley (2006), Activity-Based Travel Forecasting Model for SACOG: Technical Memo Number 6: Day Pattern Activity Generation Models, July 31, 2006 – Draft 2, 21 pages. Available at http://jbowman.net.

Bradley, Mark A., John L. Bowman and Bruce Griesenbeck (2006) Development and application of the SACSIM activity-based model system, presented at the 11th World Conference on Transport Research, June, 2007, Berkeley, California, USA. Available at http://jbowman.net.

Parsons Brinckerhoff Quade & Douglas, Inc. (2005) New York Metropolitan Transportation Council, The Metropolitan Planning Organization, Transportation Models and Data Initiative, General Final Report: New York Best Practice Model (NYBPM), 541 pages.

PB Consult Inc. and Parsons Brinckerhoff (2002) Task 3.1A Modeling Individual Daily Activity Pattern: Statistical Evidence and Approach for MORPC (project task deliverable report), 42 pages.

PB Consult Inc. and Parsons Brinckerhoff (2005) The MORPC Travel Demand Model: Validation and Final Report, 237 pages.

San Francisco County Transportation Authority and Cambridge Systematics, Inc. (2002a) San Francisco Travel Demand Forecasting Model Development Final Report: Executive Summary, 39 pages. Available at http://www.sfcta.org.

San Francisco County Transportation Authority and Cambridge Systematics, Inc. (2002b) San Francisco Travel Demand Forecasting Model Development Final Report: Tour Generation & Time of Day Models, 44 pages. Available at http://www.sfcta.org.

Vovsha, P., Bradley, M. A., Bowman, J. L. (2005) Activity-based travel forecasting models in the United States: Progress since 1995 and Prospects for the Future, in Progress in Activity-Based Analysis, H Timmermans, ed., Elsevier, pp 389-414. Pre-publication draft available at http://jbowman.net.