

**San Francisco Dynamic Traffic Assignment Project
“DTA Anyway”**

Future Research Topics Report

Prepared by:

Parsons Brinckerhoff
&
San Francisco County Transportation Authority



November 28, 2012

1. Introduction

Having recently completed the implementation of a citywide dynamic traffic assignment (DTA) model, this document outlines a work plan for future model development. It contains a mix of both short-term and long-term goals, with some discussion of prioritization and ordering provided in the conclusions. The goals are both for the DTA model, and supporting enhancements to SF-CHAMP.

The potential improvements are broadly broken into three categories:

- *Improvements in representation* are areas in which a certain aspect of reality can be better captured in the model system.
- *Improvements in usability* are not expected to change the model results, but simply make the process easier to apply for the user.
- *Improvements in linkages* represent ways in which the flow of information between SF-CHAMP and DTA can be enhanced to take advantage of the DTA structure.

For each potential improvement, a low and high cost estimate is provided. These cost estimates are included to provide a rough estimate of the level-of-effort associated with particular tasks, and do not represent any specific cost proposal. The tasks could be completed by SFCTA staff, consultant staff, or a combination of both. For the purpose of this document, the two are treated as interchangeable--an hour of SFCTA staff time is valued at the same rate as an hour of consultant time. The final cost will vary based both on who is completing the work, and on the complexity of the approach selected.

This document should not be construed as a requirements list that is necessary before the model can be successfully applied. Rather, the model has been validated, and applications work can begin immediately. These goals instead represent improvements that can be made to further improve the results. In addition, the experience gained from model applications may be used to add to or re-prioritize this list.

2. Improvements in Representation

These improvements are ways in which some aspect of reality can be better captured in the DTA model.

2.1 Transit Representation

Why would this be nice?

As a dense urban environment with high transit usage, San Francisco is faced with an array of challenges in reliably operating a bus system. The buses are subject to congestion delays, as well as crowding and delays associated with boarding and alighting. Planning efforts have focused on improving the bus system in a range of ways, including the use of bus-only lanes, and implementing dedicated bus rapid transit. The DTA model offers the potential to better understand how the busses interact with highway traffic and move through the system, and provide a better tool for evaluating such policies. Our experience calibrating the model has shown that specific chokepoints can have a large effect on the overall DTA result, and that bus only lanes can serve as chokepoints. For all of these reasons, it is desirable to further improve the representation of transit in the DTA model.

Data Requirements

The main data requirement is knowledge of the transit system, which is already reflected both in SF-CHAMP and DTA. Optionally, automated vehicle location (AVL) data and automated passenger counter (APC) data from the transit system can be processed to serve as a validation data set. Measures of interest are the runtimes for bus routes by time of day, and the variation in runtimes or arrival times.

Possible Methodology

There are two possible methodologies covered here: better representation of transit only lanes and smoother FAST-TrIPs integration.

Better Transit Only Lane Representations

The transit only lanes in San Francisco the transit only lanes allow busses are generally located on the right side of arterials in the downtown area. Busses are allowed entry, as are autos if they are making right turns. In addition, there is some fraction of vehicles that are violators, particularly as queues build.

The current version of Dynameq does not replicate this condition precisely. Instead, each lane can be given a restriction by vehicle class, but that vehicle class specific restriction is not movement specific. So either autos are allowed into the lane or they are not. There is no accommodation for right turns and no accommodation for violators. Initial testing during calibration showed that completely excluding autos from the bus only lanes resulted in gridlock at several key locations. To approximate this real world condition, the DTA model currently uses a split-link approach, where bus-only links are split in half, with the back half of the link allowing only busses, and the front half of the link allowing all vehicles. In this way, autos can use the front half of the link to make a right turn, but to continue straight through the intersection would need to move over a lane.

This task would involve incorporating a more refined approach to simulating the transit only lanes. An upcoming version of Dynameq will allow for vehicle class restrictions to be movement-specific, which will allow appropriate coding of autos only allowed to use the transit lanes for right turns. When this version of Dynameq is released, changes can be made to take advantage of this feature and the calibration can be updated.

Smoother FAST-TrIPs Integration

Another option to improve the transit representation is to develop a disaggregate, person-based transit accessibility and assignment tool to be used within their SF-CHAMP model that is capable of representing transit reliability and flexible enough to accommodate future plans to develop a route choice model from person-based GPS data in the 2012 California Household Travel Survey. Based on preliminary work undertaken during Summer 2012, the Authority has identified the open-source FAST-TrIPs (**F**lexible **A**ssignment and **S**imulation **T**ool for **T**ransit and **I**ntermodal **P**assengers), developed at University of Arizona, as an appropriate tool. The underlying agent-based simulation in FAST-TrIPs also makes it an ideal tool for simulating the reliability of transit paths based on changes in demand. These simulations, partnered with the San Francisco's Dynamic Traffic Assignment (DTA) model and extensive observed data from Automatic Passenger Counters (APCs) and Automatic Vehicle Location (AVL) systems provide a unique package to examine transit reliability from both demand and network perspectives.

Preliminary work involved successfully setting up a transit network and transit operating schedules using Google's General Transit Feed Specification (GTFS) data from San Francisco Muni. Specifically, using FAST-TrIPs, we have completed a deterministic, capacity-constrained, schedule-based assignment using the Muni transit network and the SF-CHAMP disaggregate trip data. We have completed a stochastic (discrete choice) assignment model, which is also capacity-constrained and schedule-based. Finally, we have successfully achieved several feedback loops between FAST-TrIPs and the San Francisco Citywide DTA model in Dynameq whereby the boardings and alightings predicted by FAST-TrIPs update the transit dwell times in DTA, and the subsequent travel times from specific transit-vehicle-trip departures are fed back to FAST-TrIPs.

This task would involve moving that initial implementation from a test case to be a practical tool for application. Specifically, this would involve:

- Improve integration with San Francisco Citywide DTA model for PM Peak Period
- Assess possible convergence methodologies between FAST-TrIPs and DTA
- Add other transit modes and providers to the SF network (currently only Muni buses are being used)
 - Develop an option to use static traffic assignment results in FAST-TrIPs
 - Develop skimming procedures in Fast TrIPs
 - Conduct sensitivity analysis of FAST-TrIPs skims within SF-CHAMP
 - Develop flexible and open-source APC and AVL data processing techniques in R or Python
 - Develop a methodology for quantifying reliability in transit as a function of demand
 - Develop a methodology for quantifying reliability in transit as a function of network and service characteristics
 - Develop methodology in FAST-TrIPs to quantify various aspects of reliability from Task 4
 - Examine how demand plays a role in travel time reliability using the FAST-TrIPs passenger simulation
 - Improve representation of preferred arrival or departure time in FAST-TrIPs based on research on values of time for arriving earlier or later.

The resulting product would be both a usable FAST-TrIPs model and a series of white papers documenting the findings.

What would it take?

Table 1 shows the estimated level of effort for better transit only lane representation. Table 2 shows the estimated level of effort if the task focuses on smoother FAST-TrIPs integration.

Table 1. Estimated Level of Effort for Better Transit Only Lane Representation

	Low Estimate			High Estimate		
	Senior Modeler	Modeler	Total Cost	Senior Modeler	Modeler	Total Cost
Loaded Rate	\$160	\$100		\$160	\$100	
1. Update Software	8	20	\$3,280	8	40	\$5,280
2. Recalibrate Model	16	40	\$6,560	20	80	\$11,200
Total Cost	\$3,840	\$6,000	\$9,840	\$4,480	\$12,000	\$16,480

Table 2. Estimated Level of Effort for Smoother FAST-TrIPs Integration

	Low Estimate			High Estimate		
	Senior Modeler	Modeler	Total Cost	Senior Modeler	Modeler	Total Cost
Loaded Rate	\$160	\$100		\$160	\$100	
0. Project Management	40		\$6,400	50	0	\$8,000
1. DTA Integration	100	150	\$31,000	120	200	\$39,200
2. Network Development	40	250	\$31,400	50	300	\$38,000
3. SF-CHAMP Integration	100	200	\$36,000	120	250	\$44,200
4. Examine Reliability in Data	50	160	\$24,000	60	200	\$29,600
5. Develop Reliability Metrics	100	150	\$31,000	120	200	\$39,200
Total Cost	\$68,800	\$91,000	\$159,800	\$83,200	\$115,000	\$198,200

2.2 Improved Non-Motorized Mode Representation

Why would this be nice?

During model calibration, it became clear that in a dense urban environment, congestion is driven not only by auto traffic on the roadway, but also by conflicts with pedestrians, and to a lesser degree, bicycles. These conflicts have the biggest influence on the capacity for right and left turns, where vehicles attempted to complete turning movements must first wait for a herd of pedestrians to clear the intersection. As a result, not only is that particular vehicle delayed, but a queue builds up behind it. The delay is a function of the number of pedestrians crossing. Travel patterns and congestion levels in the CBD area cannot be properly modeled without some representation of the pedestrian conflicts.

Data Requirements

Two data elements would be of particular interest to this work. The first is counts of pedestrian and bicycle demand at intersections, which can be used to validate what SF-CHAMP predicts. The second is observations of the relationship between pedestrian and bicycle demand and vehicular traffic flow. This can be observed by counting the pedestrian volumes at intersections and the number of vehicles that can pass through or are waiting for those pedestrians to pass. In some situations, pedestrian queues may need to dissipate before any vehicle can execute a movement that conflicts with pedestrian crossings. In these situations the relationship between pedestrian volume and vehicular saturation flow may be more nuanced than a simple pedestrian friction factor. Observations of conflicting pedestrian, bicycle, and vehicular movements would ideally capture time information relative to signal phase start and end times. Furthermore, the pedestrian and vehicle flow relationship may also be impacted by signal phasing, street widths, or other factors. These should be recorded as well.

Possible Methodology

There are two possible approaches that could be used to model this: a demand-specific approach and a demand generic approach. The current DTA implementation uses a demand generic approach, where the capacity of right turns is reduced for denser area types by increasing the follow up time. This approach assumes that pedestrian demand is highest in the CBD and lowest in the outlying areas and introduces friction accordingly. The weakness of this approach is that pedestrian volumes and the impact of pedestrian activity on vehicle throughput vary significantly over relatively short distances. Whereas an intersection

adjacent to a rapid transit station may have pedestrian volumes that severely limit vehicular turning movements, an intersection a couple blocks away may have pedestrian volumes low enough that vehicular turning movement capacity is not significantly impacted by pedestrians.

An enhancement to the current approach would be to use a demand specific approach where the right turn capacity given the DTA model would be coded as a function of the pedestrian demand predicted by SF-CHAMP. In this way the model would be responsive both to the specific locations where pedestrian demand is high as well as to changes in pedestrian demand in the future or in alternate scenarios.

Friction with bicycles could be coded in an equivalent way. If a demand-specific approach is used for bicycles, the bicycle demand would be taken from the SFCTA bicycle assignment model, and the traffic flow parameters would be a function of those bicycle flows.

What would it take?

Table 3 shows the estimated level of effort for this task.

Table 3. Estimated Level of Effort for Improved Non-Motorized Mode Representation

	Low Estimate			High Estimate		
	Senior Modeler	Modeler	Total Cost	Senior Modeler	Modeler	Total Cost
Loaded Rate	\$160	\$100		\$160	\$100	
1. Analyze Data	6	16	\$2,560	16	40	\$6,560
2. Implement Changes	4	16	\$2,240	8	30	\$4,280
3. Recalibrate Model	8	32	\$4,480	16	80	\$10,560
Total Cost	\$2,880	\$6,400	\$9,280	\$6,400	\$15,000	\$21,400

2.3 Robust Parking Model

Why would this be nice?

One of the weaknesses in the SF-CHAMP roadway assignment model (both the static model and the DTA model) is the current lack of representation of parking supply. Because of this shortcoming, TAZs with a shortage of parking may have unrealistically high levels of roadway volumes in their vicinities, and traffic congestion caused by large concentrated parking supply (large garages) may be under-represented. Thus, this lack of detail may be causing inaccuracy on a local level, and enhancing model sensitivity to parking supply locations would improve the usefulness of the model for location-sensitive studies such as cordon pricing analysis and parking pricing strategies.

Data Requirements

For the last several years, SFMTA has been collecting information on San Francisco parking supply, including off-street publicly available parking and the price of that supply, as well as on-street parking (both metered and unmetered). This supply information is available in GIS format. It is important to note that parking data is fairly complex due to the constantly changing nature of the parking pricing, the complexity of the parking pricing structures, and the lack of information about how much parking supply is reserved (for monthly parkers who are paying, or for monthly parkers with parking subsidized by employers).

Possible Methodology

In 2009, the SFCTA modeling team implemented a preliminary enhancement to SF-CHAMP's static model to include parking locations, capacity and the tradeoffs between parking cost, search time and walking distance (Zorn, et al, 2009). In this study, a discrete choice model was estimated using a 2006 stated preference survey on Parking where San Francisco parkers were asked to trade off parking cost, walk distance and parking search time. Results from this model were used to formulate a generalized cost function encapsulating this trade-off. Parking lots, garages and on-street parking were represented as nodes in the street network, which were connected to the road links via "parking ramps", and connected to the TAZ via "walk links." Drivers without reserved parking spaces at their destination would incur the cost of parking and search time on the parking ramp links (depending the vacancy of the relevant parking capacity), and they would incur the cost of the walk on the walk links.

Although the initial implementation was shelved due to missing parking data (reserved parking vs unreserved parking, as well as on-street un-metered parking inventory), some of this data has become available and the implementation could be explored in the DTA model. There are several advantages and disadvantages to a DTA-based parking model as described.

One potential advantage is that the DTA model is more flexible in terms of enabling vehicle trip tables to be split into multiple user classes. In the initial implementation, the static assignment software limited the number of user classes in the roadway assignment, so driver attributes such as value of time and parking duration (affecting parking cost) were not representable; this restriction could be alleviated with the flexibility available in the DTA software package, although additional user classes must be carefully weighed against computational requirements.

An alternative to further segmenting the trip tables would be to push the parking location choice into SF-CHAMP and out of the assignment. The advantage to doing this is that because SF-CHAMP uses a synthetic population the amount of market segmentation that can be used is essentially unlimited, with a continuously distributed value of time of particular relevance here. To accomplish this, the parking location choice model would need to be run iteratively to determine the appropriate shadow prices to avoid exceeding the parking capacity. Examples of similar models can be found in the University of Arizona sub-model, the LA Metro park-and-ride model, and the Atlanta Regional Commission model, among others.

A second alternative would be to leverage the dynamic nature of the DTA model to account for the dynamic nature of parking. In the static assignment parking model, parking capacity is imposed through a shadow price on the link associated with the parking lot. The effect of this is that everyone who wants to use that link has to endure that perceived cost. In reality, until a parking lot is nearly full, that cost is near zero, and then once it fills all additional users are denied entry. Therefore, the parking options can be dramatically different at 9 am than at 7 am. Because the DTA simulates vehicle arrivals in detailed time steps, it should be possible to add the accounting to reflect the dynamic of lots filling.

Finally, modeling parking garages as nodes in the DTA network is a useful exercise even if the tradeoffs between parking choices are not included, because this level of detail enhances the accuracy of roadway volumes affected by these concentrations of parking supply. Our experience calibrating the DTA model has shown that specific movements, down to the level of driveways, can play an important role in the DTA

results. One limitation to this approach is the awkwardness of representing the walk links within the DTA network.

What would it take? (Range of Level of Effort)

Table 4 shows the estimated level of effort for this task. Note that the cost for “3. Additional User Classes” and “4. Account for Dynamic Lot Filling” are zero for the low estimate because estimate assumes that these features would only be included in the high estimate.

Table 4. Estimated Level of Effort for Robust Parking Model

	Low Estimate			High Estimate		
	Senior Modeler	Modeler	Total Cost	Senior Modeler	Modeler	Total Cost
Loaded Rate	\$160	\$100		\$160	\$100	
1. Code Lots, Driveways, Capacity	24	200	\$23,840	24	300	\$33,840
2. Extend Current Model to DTA	32	100	\$15,120	40	150	\$21,400
3. Additional User Classes			\$0	30	80	\$12,800
4. Account for Dynamic Lot Filling			\$0	150	200	\$44,000
5. Calibration	32	100	\$15,120	40	160	\$22,400
Total Cost	\$14,080	\$40,000	\$54,080	\$45,440	\$89,000	\$134,440

2.4 Add External Geographic Representation

Why would this be nice?

The SF-CHAMP model covers the 9-county San Francisco Bay Area, but the SF-DTA model is run only for San Francisco County. To accommodate this difference, a subarea extraction is run on the static assignment model to create a county trip table for use in the DTA model. Unfortunately, during the subarea extraction process, there is a loss of information such that the TAZ of the external trip end is not maintained. This information loss makes it impossible directly relate the DTA trip tables to the trip lists that are output by the SF-CHAMP micro-simulation.

A more direct linkage between the DTA model and the trip lists would be valuable because it would allow information from the trip lists to be passed to the DTA model, particularly as it relates to time-of-day. To provide a more detailed temporal profile to the DTA model, it is valuable to understand the characteristics of the trips leaving each zone to each zone or external station. This is of particular relevance as the time periods being simulated by DTA get longer. For example, consider the current warm-up period that starts at 2:30 pm, drawing from the SF-CHAMP mid-day period for the 2:30-3:30 hour. The best the model can currently assume is that a uniform share of all trips in the mid-day period occur between 2:30 and 3:30 pm. In reality, there may be a tendency for short work-based sub-tours to occur during the lunch hour, and that the end of the mid-day period is populated by longer commuting trips, or by school trips. Maintaining a linkage to the trip lists would allow the purpose of the trip to be considered when building a temporal profile, and allow that profile to vary based on the origin and destination of the trip.

In addition, relating the DTA information more directly to the trip lists opens additional possibilities for integration between the two models.

Possible Methodology

The proposed methodology to maintain this linkage follows that of the Path Analysis (PA) module of the Oregon Statewide Integrated Model (SWIM). SWIM is an integrated economic-land use-transportation model, where the transportation model is an activity-based model that operates in a micro-simulation framework similar to SF-CHAMP. An enhanced subarea extraction method was developed for SWIM to generate trip tables and trip lists for urban areas from the statewide model results.

The approach works by combining a select link analysis at each external station with the subarea extraction. A select link matrix is created for each external station, both going both inbound and outbound. These matrices indicate which zone pairs in SF-CHAMP use each external station. Each matrix is then converted to a flat file format that lists the origin TAZ, the destination TAZ and the external station, if any, that is used to travel between them. The files are merged for all external stations. It is possible that for a single zone pair, multiple external stations are used. In such cases, the fraction of trips between that zone pair using each external station is calculated. The format of the resulting file is similar to that shown in Table 5.

Table 5. Sample Select Link Results

Assignment Class	Select Link A-B Nodes	Direction	From TAZ	To TAZ	Share Using this External Station	External Station Number
Auto	17813 17811	OUT	1501	1	1	10012
Auto	17813 17811	OUT	1501	3	1	10012
Auto	17813 17811	OUT	1501	9	1	10012
Auto	17813 17811	OUT	1501	52	1	10012
Auto	17813 17811	OUT	1501	53	1	10012
Auto	17813 17811	OUT	1501	54	1	10012
Auto	17813 17811	OUT	1501	64	1	10012
Auto	17813 17811	OUT	1501	73	1	10012

This select link results file is then joined to the trip list with the inbound and outbound external station appended to the trip list. If multiple external stations are used between a zone pair, then a probability is assigned for each. These probabilities can be used fractionally, or Monte Carlo draws can be conducted to assign each trip to a single external station. The resulting trip list would look similar to that shown in Table 6.

Table 6. Trip List with External Stations Appended

Trip ID	Tour Purpose	Orig	Dest	Trip Start Time	Trip Mode	Income	Station 10012	Station 10013	Station 10012	Station 10013
							IN	IN	OUT	OUT
15538	SHOP	2327	2363	1400	SR2	\$3,400	80%	20%	0%	0%
1903376	SHOP	2328	2372	900	DA	\$,9500	100%	0%	0%	0%

Once the external station numbers have been appended to the trip list, the trip list contains all the necessary information to build a subarea trip table directly from the list, bypassing the aggregation necessary for the subarea assignment. In this way, subarea trip tables can be built with a much more refined temporal

resolution, defined specific to the characteristics of each individual trip. In addition, the process would make it easy to create trip tables segmented by any trip characteristic that is desired, such as income or purpose.

It is worth noting that the accounting used in this process becomes quite complex when analyzing weaving paths (paths that enter and exit the subarea multiple times). Therefore, it is desirable that the subarea be defined at physical breakpoints that avoid the potential for weaving. Fortunately, this is readily accommodated by San Francisco’s location on a peninsula.

What would it take? (Range of Level of Effort)

Table 7 shows the estimated level of effort for this task.

Table 7. Estimated Level of Effort for Robust Parking Model

	Low Estimate			High Estimate		
	Senior Modeler	Modeler	Total Cost	Senior Modeler	Modeler	Total Cost
Loaded Rate	\$160	\$100		\$160	\$100	
1. Setup External Station Matching	60	80	\$17,600	80	120	\$24,800
2. Handle Special Cases	20	80	\$11,200	40	150	\$21,400
3. Scripts to Create Demand Matrices	12	20	\$3,920	20	40	\$7,200
Total Cost	\$14,720	\$18,000	\$32,720	\$22,400	\$31,000	\$53,400

2.5 Improved Truck/Commercial Vehicle Model

Why would this be nice?

SF-CHAMP currently includes a truck and commercial vehicle model that is adapted from the Metropolitan Transportation Commission (MTC) truck model. The MTC truck model has two components, one for “very small trucks” and one for all other trucks. “Very small trucks” are defined as commercial vehicles with four tires, and includes taxis, vans, police cars, pizza delivery vehicles, and a wide range of other vehicles. The remaining trucks, which are referred to as the “truck model” and include any trucks with 6 or more tires, including delivery trucks, construction vehicles, and tractor-trailers.

The truck model is based on trip rates taken from the Quick Response Freight Model (QRFM) and calibrated to local conditions. The QRFM is based on rates of truck trips per employee derived from national level data, and they clearly show a propensity to over-estimate truck travel in dense downtown areas. This occurs for two reasons. First, the nature of employment found in central business districts (CBDs) is different. Whereas manufacturing employment in many areas may be related to actual production, manufacturing industry employment (think of firms such as Apple, Intel, etc.) in CBDs is more likely to be related to management or design, and be just typical office employment. Second, multi-story buildings allow trucks to consolidate trips not only for a single firm, but for multiple firms in the same building. Therefore, the rate of truck travel in CBDs should be lower than in other areas, but how much lower is unknown.

The very small truck trip table presents its own set of challenges. Because traffic counts do not distinguish four-tire commercial vehicles from four-tire personal vehicles, we do not have a good observation of the share of traffic attributable to very small trucks. For this reason, the very small truck trip table can only be calibrated by seeking to match total VMT, an approach that is challenging because if the model under-predicts VMT relative to counts, it is impossible to know whether that difference is due to an underestimate of commercial travel or underreporting of travel in the household survey.

The truck model is of particular importance to the DTA results because trucks have a disproportionate effect on congestion, and because there is a higher degree of uncertainty in the truck and commercial vehicle flows. While calibrating the DTA model, it was observed that trucks can make a big difference to the results, depending on how many trucks are included in the peak period, and the vehicle length used in the DTA model for trucks.

Because truck models are typically the domain of larger areas—either statewide models or metropolitan area models—this topic presents a strong opportunity for teaming with Bay Area partners such as MTC to jointly collect data and develop a refined truck and commercial vehicle model.

Data Requirements

One of the biggest challenges in developing a truck and commercial vehicle model is obtaining a good observation of current conditions. For trucks, an important part of this is a database of truck counts, some of which already exist. Long-haul trucks can be modeled effectively from commodity flow data, such as the Freight Analysis Framework, version 3 (FAF3), but commodity flow data is known to under-predict short trips, so is of limited use in San Francisco. For “very small trucks”, obtaining this information can be more challenging because they are not distinguished in traffic counts. Typically, this is done via establishment surveys, but establishment surveys are notoriously difficult, both because of the difficulty of obtaining participation in the survey, and because it is difficult to know how to expand the survey. Therefore, the data recommendations presented here are based on options for lower cost or alternative data sources.

The data that could be considered for use in a freight and commercial vehicle model includes:

- A database of truck counts, assembled both from the PeMS and from classification counts on surface streets in San Francisco. Ideally, these counts would provide a 24-hour profile of truck volumes, allowing for a basis to understand what fraction of total truck traffic occurs in the peak periods.
- While automated counts cannot observe four-tire commercial vehicles, it may be possible to conduct some manual spot observations to determine the approximate share. This can be done by manually watching the traffic and observing the number of taxis, police cars, and vehicles with a commercial decal. While this may be a useful spot check, it is inherently an imperfect observation because not all commercial vehicles can be clearly identified.
- GPS databases of commercial vehicles could be obtained and used as the basis for building a commercial trip table. These may include the database of taxis in San Francisco, a trucking industry database of long-haul truck movements, and databases of movements from other commercial fleets such as UPS. Research would need to be conducted to determine exactly which databases are available and how representative they would be, but this approach offers significant promise over a traditional establishment survey.

- An important supporting piece of information is an estimate of the size of a TAZ that goes beyond number of employees. This could be either number of firms or number of buildings. Because deliveries can be consolidated within firms and buildings, these may prove to be a better or complementary driver of truck and commercial vehicle travel than number of employees. It is possible that these data can be derived from the city’s UrbanSIM model.

Possible Methodology

The possible methodology will depend both on the available data and the number of resources that are desired to be dedicated to the task. Two options are presented here.

3-Step Truck and Commercial Vehicle Model. In this approach, the structure of the truck and commercial vehicle model would remain as it is—a 3-step model with trip generation rates, a gravity model for trip distribution, and trips assigned in the normal assignment process. The difference is that the model would be either estimated or calibrated to match whatever newly obtained data can be obtained. In this approach, the bulk of the effort will be in assembling and processing the data to provide the best possible base-year estimate of truck and commercial vehicle travel. This may include aggregating GPS records to generate link flows or trip ends by district, testing trip generation rates as a function of different employment or firm groupings, or using matrix estimation from counts to create a truck trip table that serves as a “target” for calibrating trip generation and distribution models.

Tour-Based Truck and Commercial Vehicle Model. Several regions have sought to develop more behavioral models of truck and commercial vehicle flows. These include the Calgary commercial vehicle model, the Ohio disaggregate commercial vehicle model (DCOM), an activity-based freight model developed in Australia, and one being developed for London. These models recognize that trip chaining behavior is of particular importance to commercial movements, because vehicles are routed specifically to take advantage of trip chaining, rather than to make individual trips back and forth to their establishment as they would in a hub-and-spoke system. The specific structures of these models vary. In this approach, a full literature review will be conducted, and an appropriate model structure will be selected for San Francisco based on the available data.

What would it take? (Range of Level of Effort)

Table 8 shows the estimated level of effort for this task if it focuses on a 3-step truck and commercial vehicle model. Table 9 shows the estimated effort if it focuses on building a tour-based truck and commercial vehicle model.

Table 8. Estimated Level of Effort for 3-Step Truck and Commercial Vehicle Model

	Low Estimate			High Estimate		
	Senior Modeler	Modeler	Total Cost	Senior Modeler	Modeler	Total Cost
Loaded Rate	\$160	\$100		\$160	\$100	
1. Data Analysis	20	40	\$7,200	60	160	\$25,600
2. Model Update	16	32	\$5,760	20	40	\$7,200
3. Calibration	16	32	\$5,760	30	80	\$12,800
Total Cost	\$8,320	\$10,400	\$18,720	\$17,600	\$28,000	\$45,600

Table 9. Estimated Level of Effort for Tour-Based Truck and Commercial Vehicle Model

	Low Estimate			High Estimate		
	Senior Modeler	Modeler	Total Cost	Senior Modeler	Modeler	Total Cost
Loaded Rate	\$160	\$100		\$160	\$100	
1. Data Analysis	80	200	\$32,800	100	200	\$36,000
2. Model Update	100	250	\$41,000	200	500	\$82,000
3. Calibration	60	150	\$24,600	80	200	\$32,800
Total Cost	\$38,400	\$60,000	\$98,400	\$60,800	\$90,000	\$150,800

3. Improvements in Usability

These improvements relate to the user experience of the software and the application process.

3.1 Investigate Stability Across Scenarios

Why would this be nice?

Sensitivity testing with the DTA model included two tests to evaluate the stability of the model results across scenarios.

The first of these tested running an identical scenario, but with a different random number seed. The random number seed affects the bucket rounding of non-integer matrices, as well as the exact departure time of trips within a 15-minute window. This test revealed some modest differences in traffic flows and travel times between the two runs, with the larger differences focused in the more congested portions of the network. As expected, this result indicates that there is some observable stochasticity associated with the DTA model.

The second test involved performing two model runs that were identical except for a minor change on an uncongested link. This change was not expected to have large effects on the traffic flow pattern or the travel times throughout the city. The model results showed small changes elsewhere in the network that cannot be attributed to the coded project. This result is thought to be due either to stochasticity in the bucket rounding and departure time of trips, or due to a level of convergence that while good in DTA terms is less than would typically be performed with a static assignment.

It is also possible that a contributing factor to either or both of these results is dual equilibria on certain network links due to the shape of the traffic flow curve.

Because the observed differences are relatively modest, it may be that the magnitude of difference resulting from most project applications is much greater than the differences observed here. It is not clear whether that is the case or not for smaller applications, so the topic warrants further investigation. The goal is to better understand the limits of precision of the model, and to identify steps to improve that precision.

Possible Methodology

This analysis will begin with a series of systematic tests to better isolate the source of the observed differences. These may include:

- Repeat the random number seed test, but with identical integer trip tables input to Dynameq to begin the simulation. This test will give an indication of how much stochasticity is due to the bucket rounding process of getting 15-minute integer trip tables, versus stochasticity in the exact departure time of vehicles within those 15-minute periods.
- Run each of the sensitivity tests many more iterations in an effort to achieve a smaller relative gap. Evaluate whether the tighter convergence results in smaller differences and discuss the runtime versus convergence tradeoff.
- Run identical scenarios, but with different random number seeds given to SF-CHAMP. This will help evaluate the level of stochasticity in SF-DTA versus SF-CHAMP.
- Trace paths where the differences occur to identify whether path switching appears to occur for those particular links.
- Evaluate the traffic flow conditions on the links with large travel time changes to identify whether a dual equilibria problem may be occurring.

Following this investigation, strategies for mitigating the level of noise between scenarios will be recommended. These may include: do nothing, run the DTA multiple times and average the results, run additional iterations or implement a method to integerize the trip tables prior to inputting them to Dynameq. The recommended strategies will be implemented and tested.

What would it take? (Range of Level of Effort)

Table 10 shows the estimated level of effort for this task.

Table 10. Estimated Level of Effort for Stability Testing

	Low Estimate			High Estimate		
	Senior Modeler	Modeler	Total Cost	Senior Modeler	Modeler	Total Cost
Loaded Rate	\$160	\$100		\$160	\$100	
1. Test Sources of Noise	16	40	\$6,560	30	80	\$12,800
2. Implement/Test Mitigation	24	60	\$9,840	40	80	\$14,400
Total Cost	\$6,400	\$10,000	\$16,400	\$11,200	\$16,000	\$27,200

3.2 Computing Efficiency

Why would this be nice?

Currently, the DTA model is run for the PM peak period, from 3:30-6:30 pm. The simulation includes one hour of warm-up time, and one hour of cool down time, extending the demand-loading period to 2:30-7:30 pm, as well as an additional 5 hours to allow all traffic to clear without loading any additional demand. The additional simulation time is a standard Dynameq setting that ensures that the amount of time that demand is loaded is added to the end of the simulation as time to clear the network after the demand period ends. The typical runtime is about 50-60 hours for 80 iterations, running two models simultaneously on a machine with 2 3.30GHz processors and 48.0 GB of RAM running a 64-bit operating system. The required time is significantly reduced when running only one model. These current runtimes are tolerable for the PM peak, but could quickly become impractical if we seek to further expand the complexity of the model, either by shifting towards a 24-hour DTA or by adding vehicle classes. Therefore, it is desirable to find strategies to minimize runtime to allow for the full set of enhancements to be considered.

Possible Methodology

There are a limited number of strategies that could be employed to improve the DTA model runtimes. These are listed below.

Buy Faster Hardware. Model runs for this project have been conducted both on a PB machine with two 3.30GHz processors and 48.0 GB of RAM running a 64-bit operating system and on an SFCTA machine with one 2.93GHz processor and 31.9 GB of RAM running a 64-bit operating system. Testing has shown that for an equivalent model run, the runtime on the SFCTA machine is approximately 1.5 times the runtime on the PB machine. This clearly reveals that there are gains to be made by buying faster hardware. It is worth noting however, that because Dynameq does not take advantage of parallel processing, the gains are related to the processor speed rather than the number of processors. This imposes some limit on the gains that can be made through hardware.

Integerize the Trip Tables Before Inputting to Dynameq. Dynameq only builds shortest paths for zone pairs that have a trip going between them, making it inefficient to input aggregate trip tables with a small number of fractional trips between a large number of zone pairs. As part of the simulation, Dynameq uses fractional trips as a probability and simulates whether or not a trip leaves between that zone pair in each time step. Therefore, it is whole vehicles that are simulated traveling across the network.

During the course of this project, an option was explored to use Dynameq to apply a bucket-rounding approach to integerize the trip table. To accomplish this, a few iterations of the assignment are run, then the assignment is stopped, and a set of trip tables exported based on the vehicles actually assigned by Dynameq. Then the integer trip tables are then used as the input to a full assignment. The integer trip tables offer the potential for improved runtimes. This process was explored, but never institutionalized, due to a focus on other priorities. One option is to formalize this process as part of the standard DTA model run approach.

A related approach is possible that takes advantage of SF-CHAMP's characteristic as a demand microsimulation model, where integer trips are simulated through the demand model. Rather than using Dynameq to implement a bucket rounding approach, it may be beneficial to integerize trip tables before inputting them into Dynameq. This would offer an advantage both of improved consistency between the two models, and of a more seamless application process. There are three reasons that the current trip tables are not integer trip tables, so to accomplish this task each of these three reasons would need to be addressed:

1. Drive alone (DA) trips become a single vehicle trip, but shared ride 2 (SR2) trips become half a vehicle trip, and shared ride 3+ (SR3+) trips become a fractional vehicle trip as well based on the average occupancy. The shared ride trips could be integerized either by random draws to determine if the trip is made by a driver, or by building an explicit driver-passenger model.
2. To minimize simulation noise, SF-CHAMP is run five times in parallel, and the trip tables are averaged before assignment. A bucket rounding approach, however, introduces its own simulation noise, partially defeating the purpose of the multiple SF-CHAMP runs. To integerize the trip tables, a strategy could be employed either to run the model just a single time, or to evaluate the five simulations and pick the one closest to the mean.
3. The truck and commercial vehicle models are aggregate models that result in fractional trips across a large number of zone pairs. There is little alternative to employing a bucket rounding approach to these trip tables, although that approach could be implemented prior to assignment.

Lean on the Software Vendor to Improve Computational Efficiency. The final efficiency approach recognizes that in working with commercial software, the software vendor has a great deal more control than the end user. Therefore, this strategy involves leaning on the software vendor to prioritize improvements related to computation efficiency. The most obvious of these would be to take advantage of parallel processing, although we do recognize that for an event-based simulation like Dynameq that this is a non-trivial task. Ultimately, the best we can do is express our desires as an end user, and rely on the expertise of the software developers to find the most appropriate solutions.

What would it take? (Range of Level of Effort)

Table 11 shows the estimated level of effort for this task if it focuses on buying faster hardware. Table 12 shows the estimated effort if it focuses on integerizing the trip tables before they are input to Dynameq. No cost estimate is provided for leaning on the software vendor, because that cost is both unknown, and likely to be external to the Authority.

Table 11. Estimated Level of Effort for Buy Faster Hardware

	Low Estimate			High Estimate		
	Senior Modeler	Modeler	Total Cost	Senior Modeler	Modeler	Total Cost
Loaded Rate	\$160	\$100		\$160	\$100	
1. Setup	8		\$1,280	8		\$1,280
Direct Expenses			\$10,000			\$25,000
Total Cost	\$1,280	\$0	\$11,280	\$1,280	\$0	\$26,280

Table 12. Estimated Level of Effort for Integerize the Trip Tables Before Inputting to Dynameq

	Low Estimate			High Estimate		
	Senior Modeler	Modeler	Total Cost	Senior Modeler	Modeler	Total Cost
Loaded Rate	\$160	\$100		\$160	\$100	
1. Implement Bucket Rounding Strategies	8	24	\$3,680	8	40	\$5,280
2. Explicit Driver-Passenger Model			\$0	40	160	\$22,400
3. Testing and Debugging	8	20	\$3,280	8	32	\$4,480
Total Cost	\$2,560	\$4,400	\$6,960	\$8,960	\$23,200	\$32,160

4. Improvements in Linkages

These improvements relate to the flow of information between SF-CHAMP and DTA.

4.1 Develop Reliability Variables from DTA

Why would this be nice?

Travel time reliability can be an important measure both as an output of travel supply models and as an input to travel demand models. Also related to travel time reliability is the variability in traffic volumes, which in some cases can be observed more directly than travel times themselves.

If it is available as an output of a supply model (such as DTA), reliability can be used as a performance measure for proposed projects or plans. For example, consider a choice between improving a single primary roadway in a corridor versus a set of parallel roadways. The single roadway project may prove to have a better cost-benefit if only travel time savings are considered, but the set of parallel improvements may prove to be a more robust approach. This choice can only be fully evaluated if reliability is available as a measurable criterion. Tolling research¹ indicates that the value of reliability may be as high as the value of time.

Similarly, reliability can be an important measure to be fed back into travel demand models. Consider a traveler who must be at work or at an appointment on-time. This traveler is likely to leave early to allow for a buffer to ensure that they arrive on time. This buffer is not necessarily travel time, but is still wasted time. Accounting for such behavior may have an important effect on travel choices generally, and the outputs of travel demand models specifically.

With these enhancements, the model would become sensitive to policy changes that improve the reliability of the transportation system. For example, congested conditions tend to have high variance of travel times due to the unstable nature of their flow. Therefore, congestion relief project should have some benefit beyond the improvement in the average travel time that is currently measured. The model would also better reflect the benefits of redundancy in the transportation system in the form of grid networks that are better able to absorb changes in traffic.

Data Requirements

While reliability can be simulated without data, an observation of travel time reliability is important in that it would provide a basis for calibrating the model. There are two promising approaches to measuring travel time reliability, both of which are technological solutions.

The first approach is to leverage automated traffic recorders, such as those used by the Performance Measurement System (PeMS). These data are readily available on Bay Area freeways and can continuously record the speeds and volumes on specific links. These data will provide a basis for calibrating link-level reliability measures for a small subset of links in San Francisco.

The second approach is to use high-fidelity location data associated with vehicles. Possible data sets include GPS data from fleets of commercial vehicles, such as taxis, or for transit, from the busses. These data are expected to provide broader coverage of the city, and could provide origin-destination based measures of reliability.

Possible Methodology

The choice of how to measure reliability can be broken into two important dimensions. In the first dimension, there is a choice of whether to measure reliability within the confines of a daily simulation, or whether to measure it across multiple days. In the second dimension, there is a choice of whether to incorporate an analytical proxy that requires a single model run, or to perform multiple dynamic simulations.

Measuring reliability within the confines of a day reflects the fact that the travel time in a dynamic simulation (and in the real world) can vary substantially for different departure times, or even depending on the precise

¹ Concas and Kolpakov, *Synthesis of Research on Value of Time and Value of Reliability*, 2009

timing with respect to signals and queues. Routes where the travel time is relatively stable throughout the simulation period can be considered more reliable. The advantage to measuring reliability within the simulation period is that it is better suited to a single simulation run. However, representing within-period reliability measures using a DTA model offers some interesting challenges. The theory on which dynamic user-optimal equilibrium is based assumes that travelers have perfect network information and make rational choices to minimize their travel disutility. At equilibrium, there would be no variation in travel disutility between origins and destinations for travelers departing in the same specified time intervals. Under these assumptions, reliability must be determined with respect to variation in disutility over time intervals. Given relatively short interval link attributes, it would be possible to measure variability of route travel times assuming travelers left 5 or 10 minutes earlier or later than their equilibrium departure interval.

Measuring reliability across multiple days offers the potential to capture a much broader range of differences, resulting from weather, incidents, special events, or other variability in traffic patterns. The main challenge in this is that it breaks away from the approach of modeling an average weekday.

A related question is whether to consider reliability with an analytical approximation from a single simulation, or whether to run multiple simulations. The multiple simulation approach is more realistic, but less practical because it requires a substantial amount of computing resources. If the multiple simulation approach is taken, it may be best suited to a one-time study, rather than a regular part of the model run approach. A baseline simulation would be run first, and then a series of perturbations would be introduced, by varying one or more of the following:

- The random number seed given to the DTA;
- The random number seed given to SF-CHAMP;
- The overall quantity of demand;
- Randomly introduced special events; or
- Randomly introduced capacity reductions to represent incidents.

Upon completion of these test runs, measures from the DTA model would be assembled and validated against observed data.

A single-run approach may be better suited to incorporate into a regular model run system. The recent *SHRP 2 Project L04 on Incorporating Reliability Performance Measures in Operations and Planning Modeling Tools* explores options for reliability measures, specifically in the context of DTA. The primary options are to:

- Incorporate perceived highway time by congestion levels; or
- Include a time variability measure at the link level that is calibrated as a function of the link type and congestion level.

If a single-run approach selected, it is recommended that the approach build upon the work of SHRP 2 L04.

What would it take? (Range of Level of Effort)

Table 13 shows the estimated level of effort for developing reliability measures from DTA. Note that the multi-run testing is assumed to only occur in the high estimate.

Table 13. Estimated Level of Effort for Developing Reliability Variables

	Low Estimate			High Estimate		
	Senior Modeler	Modeler	Total Cost	Senior Modeler	Modeler	Total Cost
Loaded Rate	\$160	\$100		\$160	\$100	
1. Design	40		\$6,400	80		\$12,800
2. Data Processing and Analysis	40	60	\$12,400	80	200	\$32,800
3. Implement Single-Run Proxies	40	100	\$16,400	40	120	\$18,400
4. Perform Multi-Run Testing			\$0	120	400	\$59,200
5. Evaluation	40	20	\$8,400	80	40	\$16,800
Total Cost	\$25,600	\$18,000	\$43,600	\$64,000	\$76,000	\$140,000

4.2 24-Hour DTA

Why would this be nice?

Modeling an area with DTA for only a peak period, or multiple separate periods raises temporal boundary issues. There is always a question about how to pre-load the network such that as the demand period of interest begins loading, the network conditions are being represented realistically. It is practically difficult to get the initial conditions set appropriately for a large network where demand is presumably loading at different rates, at different times, depending on the area.

One potential benefit of such a model, beyond the elimination of temporal boundary issues, is the ability to develop more refined level of service information for the 24-hour ABM. In addition, a 24-hour DTA model provides an ideal base model from which to draw subarea models for multi-resolution modeling purposes. Different subareas from the DTA model may have time intervals of interest that begin and end at different times. A single 24-hour DTA model allows analysis different subareas, at the specific time relevant to that area, with minimal subarea configuration.

A 24-hour DTA model is also needed if it is to serve as the basis for level of service information that is fed back to SF-CHAMP.

Data Requirements

The data requirements for a 24-hour DTA are simply an expansion of the data that were used to develop the PM model. This includes signal timing plans that vary by time of day, an accounting of bus lanes, parking lanes or reversible travel lanes that vary by time of day, and a more complete set of traffic counts that go beyond the PM peak period.

Possible Methodology

Implementing a 24-hour DTA model would first require some mechanical changes to the DTA Anyway code to read in the appropriate network and control features for that particular time of day, and to do subarea extractions from SF-CHAMP for each of the five time periods. Mechanically, these changes are not difficult.

The bigger issue is the runtime requirements for such a model, so implementing this strategy is dependent upon what efficiency gains can be made.

In addition, it is reasonable to expect that the boundaries between the five time periods used in SF-CHAMP will be a source of difficulty. Therefore, it would be highly valuable to pair this effort with a stronger external geographic representation, as discussed above. This would allow the temporal profiles to be better smoothed at the boundaries between periods by applying temporal profiles to the trip lists, then generating half-hour trip tables directly.

Finally, a significant portion of the effort is expected to be in calibrating the models. The queuing and choke points will be in different locations in the AM peak than in the PM peak, and will cause a different set of gridlock issues that need to be resolved.

What would it take? (Range of Level of Effort)

Table 14 shows the estimated level of effort for this task.

Table 14. Estimated Level of Effort for 24-Hour DTA

	Low Estimate			High Estimate		
	Senior Modeler	Modeler	Total Cost	Senior Modeler	Modeler	Total Cost
Loaded Rate	\$160	\$100		\$160	\$100	
1. Data Assembly and Checking	8	40	\$5,280	16	80	\$10,560
2. Software Updates	16	40	\$6,560	20	80	\$11,200
3. Calibration	80	240	\$36,800	120	400	\$59,200
Total Cost	\$16,640	\$32,000	\$48,640	\$24,960	\$56,000	\$80,960

4.3 Add Temporal Robustness

Why would this be nice?

The current implementation of SF-CHAMP has a fairly coarse time of day representation, with only 5 time-periods represented (Early AM: 3 AM – 6 AM, AM Peak: 6 AM -9 AM, Midday: 9 AM – 3:30 PM, PM: 3:30 PM - 6:30 PM, and Evening 6:30 PM – 3 AM). This temporal resolution has several obvious drawbacks. Travel conditions may vary widely within these time periods. For example, the midday period includes lunch travel and also end-of-school student travel, which have very different patterns and peaks. The evening period is especially long and poorly represents the variation in early evening versus middle-of-the-night travel times, which are especially different for transit travel. This period is also so long that it presents challenges to coding transit schedules. Some of the advantages of using DTA, with its finer temporal resolution, are lost. Since the San Francisco DTA model can produce auto and transit skims at a much higher temporal resolution, to throw away this information and aggregate these travel conditions into the coarser skims is to waste an opportunity for improving SF-CHAMP.

Adding finer temporal resolution to SF-CHAMP would improve the accuracy of the model in several key ways. One of the key sets of questions often posed in modeling project is: how does the built project affect mode choice, route choice, time-of-day choice, or the decision to travel at all? These decisions are clearly all related; if a traveler can shift their tour or trip 15 or 30 minutes and face much better travel conditions without changing mode or route, then this fact will likely influence all of their travel decisions, especially if the tour activity allows for flexibility. Thus, the potential for improving the performance of SF-CHAMP with better temporal resolution is evident.

Data Requirements

In order to develop a refined Time of Day model for SF-CHAMP, a household travel survey dataset with the requisite time of day granularity would be required. The 2012 California Household Travel survey will have travel and activity times specified to the minute, and should provide enough detail for a variety of implementation options.

Possible Methodology

The most straightforward option for improving the temporal robustness of SF-CHAMP is to simply increase the number of time periods represented so that each time period is representative of a fairly consistent set of travel conditions. These time periods could be defined by analyzing observed transit and auto travel times from various data sources (including CMA data, bus GPS data, and household travel survey GPS data) and designating the desired level of consistency. In addition, the travel times of the DTA model could also be analyzed for consistency with the observed data. Once the new time periods are defined, the tour time of day model would be updated to reflect this new set of options, and re-estimated based on the 2012 California Household Travel Survey dataset. The downside to simply extending the current time-of-day model is that it relies on constants associated with every combination of departure and arrival periods. With only five time periods, there are fifteen alternatives (because you cannot return earlier than you depart), and it is entirely practical to include 15 alternative specific constants in the model. However, if the model included 48 half-hour periods, there would be 1,176 possible combinations of departure and arrival periods, and it would no longer be realistic to estimate or calibrate that many alternative specific constants.

Instead, the model could be reformulated in the style of the Columbus AB model system (Vovsha & Bradley, 2005), with an arbitrarily large number of alternatives (based on the time period granularity), but with a limited number of estimation coefficients and a “continuous shift” interpretation of the duration variables. The “continuous shift” variables allow for coefficients to be estimated based on the variable’s tendency to shift the arrival or departure time earlier or later in the day, or to shift the duration of the tour to be longer or shorter. This approach maintains the more detailed temporal resolution, but reduces the dimensions that must be estimated and calibrated.

A related set of enhancements relates to the way in which time windows are tracked. Persons cannot participate in more than one activity at a time, so it is realistic to first schedule the highest priority tours, reserve that time, and then go on to schedule the remaining tours. These time windows provide constraints that bound how much tours can shift, and prevent the model from being overly sensitive to peak spreading. There is a question as to whether rescheduling should be allowed to occur later in the model stream in response to different than expected travel times or in response impedance differences.

Another important question is the appropriate level of temporal detail to include in the time-of-day model. A separate, but related question is the appropriate level of temporal detail at which to feed travel time skims into SF-CHAMP for this model. It is possible for the time-of-day model to operate on 5, 15, or 30 minute periods, but look up impedance information from more aggregate skims. Much of the motivation of updating the time-of-day model is to include travel time information for more detailed skims, so the time-of-day model should operate on a scale at least as detailed as the skims.

One option for feeding the skims back is to simply go to the 5-minute intervals represented in the DTA model. While this would simplify the skimming process for feeding travel conditions from the DTA model into the mode choice model for SF-CHAMP, the biggest challenge to this approach would be how to resolve

the performance issue: currently the number of variables involved in the trip mode choice model for SF-CHAMP have expanded to such an extent that the travel skims for all time periods have a significant memory footprint. In order to drastically increase the number of time periods, an alternative approach would be needed, such as a database- or network-based solution, or one that involves utilizing multiple machines such that not all the travel skims would be required in memory on all machines. Due to these potential complications, the peer review panel recommended that we use 30 minute periods for the skims in the peak, with the option to aggregate in the off-peak. Because it is the skims that require the large memory footprint, the time-of-day models can still be run in more detail if desired.

What would it take? (Range of Level of Effort)

Table 15 shows the estimated level of effort for this task.

Table 15. Estimated Level of Effort for Add Temporal Robustness

	Low Estimate			High Estimate		
	Senior Modeler	Modeler	Total Cost	Senior Modeler	Modeler	Total Cost
Loaded Rate	\$160	\$100		\$160	\$100	
1. Data Processing (code tours, modes, etc)			\$0	80	160	\$28,800
2. Model Estimation	60	80	\$17,600	60	120	\$21,600
3. Implement Time-of-Day Models	24	80	\$11,840	30	120	\$16,800
4. Update Remainder of SF-CHAMP	24	60	\$9,840	30	80	\$12,800
5. Calibration	20	40	\$7,200	20	80	\$11,200
Total Cost	\$20,480	\$26,000	\$46,480	\$35,200	\$56,000	\$91,200

4.4 Create Skims that use DTA at an Appropriate Level of Aggregation

Why would this be nice?

A related enhancement to adding temporal robustness to SF-CHAMP is to be able to measure the level-of-service at the same resolution. Half-hourly or fifteen-minute time periods could be a reasonable option. The great benefit of this enhancement would be to provide the ability to model the dynamic changes in time and price (such as peak spreading or dynamic pricing) throughout the model system.

Currently the households' simulated choices of departure times are influenced only by average travel times over the five aggregate time periods. If dynamic skims matrices were created for smaller time periods, the mode choice logsums in the time of day choice models would vary within the peak periods. DTA can measure the extent to which projects and policies change the attractiveness of departing at 7:30 AM relative to 8:00 AM, and shift the temporal distribution of trips within the peaks. This capability is in contrast to the current model structure, which can only model much greater and less common time shifts. In addition the models would predict that travelers may be somewhat more likely to use transit during the peak hour or peak half-hour than during the shoulders of the peak due to the level-of-service differences.

Data Requirements

No additional data are required to create DTA skims. However, observed travel time data summarized at the same temporal resolution could be used to validate the skims. While route-level travel times have already been validated against the DTA outputs, an appropriate validation of the skims would be against origin-destination travel times. These data could be derived from a GPS data set for taxis or other fleet vehicles.

Possible Methodology

The dynamic skims should be specific to the vehicle classes used in the DTA, and specified at the same resolution used for the more temporally robust SF-CHAMP. The recommendation of the peer review panel was to use 30-minute increments, with the option to further aggregate the off-peak periods where there is little variation in travel time. The 30-minute increment provides an appropriate balance of measuring travel time at a level at which it is expected to change without creating an overly cumbersome process.

The process of creating dynamic skims is the same, regardless of whether it is done for only the PM peak period, or for a full 24-hour period. The choice obviously relates to whether or not a 24-hour DTA is implemented, with a 24-hour DTA providing for more consistency between the two models but potentially requiring impractical runtimes. If dynamic skims are used only for the PM peak period, static skims would continue to be used in SF-CHAMP for the remaining periods.

To create a usable set of skims for input to SF-CHAMP, the travel time information must be extracted from the DTA simulation for all zone pairs in San Francisco and for all external stations. This involves setting up the Dynameq procedures to create this output following the completion of the DTA run. Because the normal simulation does not calculate shortest paths for zone pairs that do not have trips, this process ensures that paths are built between all zone pairs for all time periods. In this process, it is also important to consider the time step used to create the skims relative to the time step used in the DTA. The DTA currently averages travel times over 5 minute steps. To generate skims for a 30-minute period, it is necessary to create an average experienced travel time across six 5-minute steps. All skims will be computed based on the trip departure time. The result of this process will be a set of impedance matrices with $N + E$ rows and columns, where N is the number of TAZs within San Francisco and E is the number of external stations at the county boundary.

The internal dynamic skims are not sufficient to feed back to SF-CHAMP, because they will only be available for San Francisco, whereas SF-CHAMP operates for the whole 9-county Bay Area. Therefore, it is necessary to merge these internal dynamic skims with the static skims from the rest of the Bay Area.

To provide additional input to this process, Cube will be used to generate a set of static skims between all TAZs outside San Francisco, as well as all external stations. The result of this skimming will be a set of impedance matrices with $O + E$ rows and columns, where O is the number of TAZs outside San Francisco and E is the number of external stations at the county boundary.

The end goal is to produce an impedance matrix with $N+M$ rows and columns that merges the dynamic and static skim information. Table 9 shows a representation of how this matrix would be configured. The $N \times N$ portion of the internal dynamic skim can be directly inserted into the top-left quadrant of the matrix. The $M \times M$ portion of the static skim can be directly inserted into the bottom-right quadrant of the matrix. Travel time for the remaining cells are part static and part dynamic and must be merged into a hybrid travel time.

Table 16. Hybrid Dynamic-Static Skim Matrix

Dynamic Skims (N×N)	Hybrid Skims (N×M)
Hybrid Skims (M×N)	Static Skims (M×M)

To calculate the appropriate hybrid travel time it is necessary to know which external station is used to go between a Bay Area zone and a San Francisco zone. Fortunately, this can be calculated using the process outlined in Section 2.4 of this document on adding external geographic representation. By running a series of select link analyses during the subarea extraction process, the (static) path from a Bay Area zone m , to a San Francisco zone n , is known to go through external station e . To get the total travel time from m to n , the static travel time from m to e is added to the dynamic travel time from e to n . This value is inserted into the appropriate cell in the bottom-left quadrant of the matrix. The reverse logic holds true for zone pairs in the top-right quadrant.

If more than one external station is used between m and n , then the path through each must be considered. In such cases, the impedance from the shortest total path should be selected. While this leads to a slight inconsistency with the original select link run, it is appropriate based on the assumption that the DTA is able to provide a better travel time estimate than the static model.

This process is repeated for each 30-minute time period. The static skims will be at a resolution longer than 30 minutes so the same static travel time will be drawn for multiple 30-minute periods.

Note that the DTA travel times are all based on the trip departure time. This presents no problem for the internal-to-internal portion of the skims, or for the outbound skims from San Francisco to the rest of the Bay Area. For the latter, the static travel time is just tacked onto the end. It does result in an offset for inbound trips, though. A trip from the Peninsula to San Francisco at 5:00 PM might spend 30 minutes to get to the county line, and another 30 minutes in the city. That trip should combine the PM peak static travel time to the external station with the dynamic travel time from the external station for the 5:30 PM period. Therefore, the inbound skims do not always draw from the same dynamic period, but are offset by the travel time to reach the external station. Note that an inverse offset should occur for inbound trips when they are converted to DTA demand matrices originating from the external stations.

What would it take? (Range of Level of Effort)

Table 17 shows the estimated level of effort for this task.

Table 17. Estimated Level of Effort for Create Skims that Use DTA and an Appropriate Level of Aggregation

	Low Estimate			High Estimate		
	Senior Modeler	Modeler	Total Cost	Senior Modeler	Modeler	Total Cost
Loaded Rate	\$160	\$100		\$160	\$100	
1. Extract Matrices from DTA	20	60	\$9,200	40	80	\$14,400
2. Evaluate & Update Temporal Aggregation	40	40	\$10,400	40	80	\$14,400
3. Combine with Static Skims	60	100	\$19,600	80	160	\$28,800

Total Cost	\$19,200	\$20,000	\$39,200	\$25,600	\$32,000	\$57,600
------------	----------	----------	----------	----------	----------	----------

5. Conclusions and Prioritization

This document has outlined a series of future research topic related to the further development of SF-DTA, SF-CHAMP, and their eventual integration. These topics will be pursued as resources allow. It is important to note that the DTA model in its current form will continue to be used in application, and will not wait for the continuing enhancements. The research topics and priorities may be revised or re-prioritized based on what is learned during model applications.

One important consideration is how to prioritize the research topics discussed here. Figure 1 shows an initial guess as to the potential schedule, level of effort and dependency among the proposed topics. Those topics listed farther down on the chart are expected to be scheduled farther into the future, and those tasks listed farther to the right are expected to involve a higher level of effort.

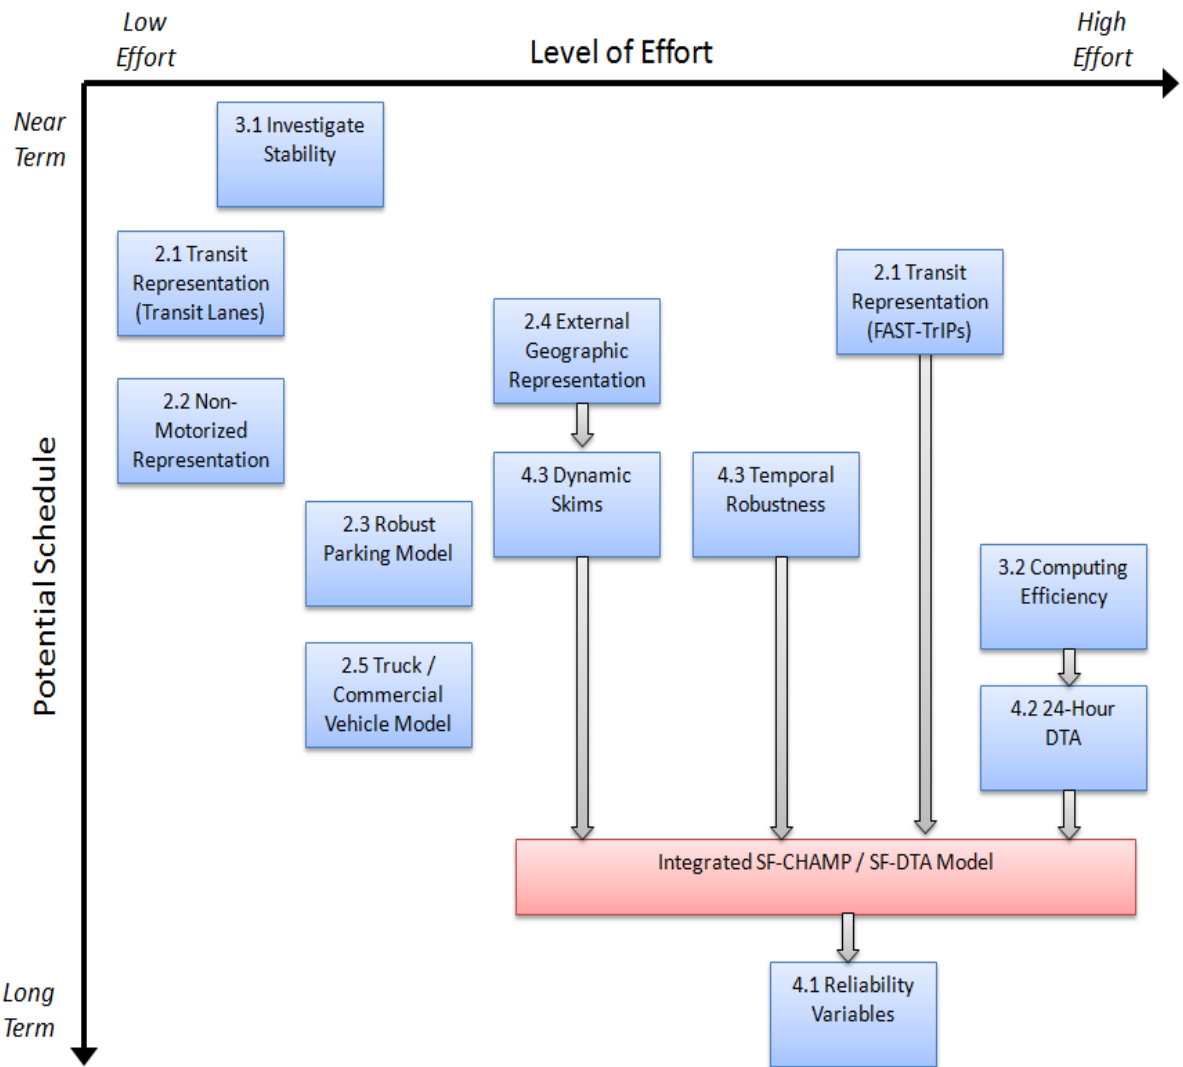


Figure 1. Potential Schedule, Level of Effort, and Dependency among Proposed Topics

The arrows show tasks that must be completed prior to other tasks. For example, the external geographic representation is needed before the hybrid dynamic-static skims can be created. Similarly, computing efficiency improvements are necessary before a 24-hour DTA is practical. The red box shows the “holy grail” of an integrated SF-CHAMP and SF-DTA model, as well as the tasks that flow into it. The purpose of this diagram is not to provide a definitive ordering of topics, but rather to provide a general understanding of their relationships and how that might affect which should be completed first. The independent boxes (2.1, 2.2, 2.3 and 2.5) can be moved anywhere as each is stand-alone. One important note from this assessment is that both the external geographic representation and the temporal robustness are central to the future of an integrated model.

To support the short-term usability of the model, the research priority should be given to investigating the stability of the model across scenarios (3.1) and identifying any mitigation strategies that may be warranted to ensure that the reported differences are due predominantly to the project being tested and not to simulation noise.