University of Kentucky

From the SelectedWorks of Gregory D. Erhardt

September 30, 2013

# Development and Application of a Dynamic Traffic Assignment Model for San Francisco

Gregory D. Erhardt, *Parsons Brinckerhoff*
Elizabeth A. Sall, *San Francisco County Transportation Authority*
Lisa Zorn, *San Francisco County Transportation Authority*
Daniel Tischler, *San Francisco County Transportation Authority*
Renee Alsup, *Parsons Brinckerhoff*, et al.

Available at: https://works.bepress.com/gregory-erhardt/11/

# DEVELOPMENT AND APPLICATION OF A DYNAMIC TRAFFIC ASSIGNMENT MODEL FOR SAN FRANCISCO

Gregory D. Erhardt
Parsons Brinckerhoff
Elizabeth A. Sall
San Francisco County Transportation Authority
Lisa Zorn
San Francisco County Transportation Authority
Daniel Tischler
San Francisco County Transportation Authority
Renee Alsup
Parsons Brinckerhoff
Neema Nassir
University of Arizona

## 1. INTRODUCTION

San Francisco recently completed the development of a citywide dynamic traffic assignment (DTA) model. The approach combines a mesoscopic DTA model with an explicit representation of every hill, street, traffic signal, stop sign, and transit route in the city.

This paper focuses on San Francisco's experience developing and applying the DTA model, on lessons learned from the process, and on recommendations for future improvements. It is structured as follows:

First, the introduction covers the policy context and motivation to pursue DTA, as well as a brief review other known DTA projects and how the approach used here relates to practice elsewhere. Second, the model development section describes an open-source toolkit developed to automate the conversion of key data DTA inputs, challenges faced in calibrating a model in a congested urban environment, and the resulting model validation. Third, the results of several policy scenarios are presented, including how the DTA model results differ from the static assignment model. Finally, the authors give their lessons learned that may be more broadly useful to others implementing DTA, and pose several open questions about how the process could potentially be improved.

Complete reports and all source code written for this project are available on the project website at: https://code.google.com/p/dta/.

© AET 2013 and contributors

## 1.1. Policy Context and Motivation

San Francisco County has a population of approximately 800,000 within a region populated by approximately 7.2 million people. Surrounded on three sides by water and on the final side by a small mountain range, San Francisco is connected to the rest of the region via two bridges: the Golden Gate to the North and the Bay Bridge to the East, two freeways to the south, a Transbay Tube carrying heavy transit rail, and a network of ferries.

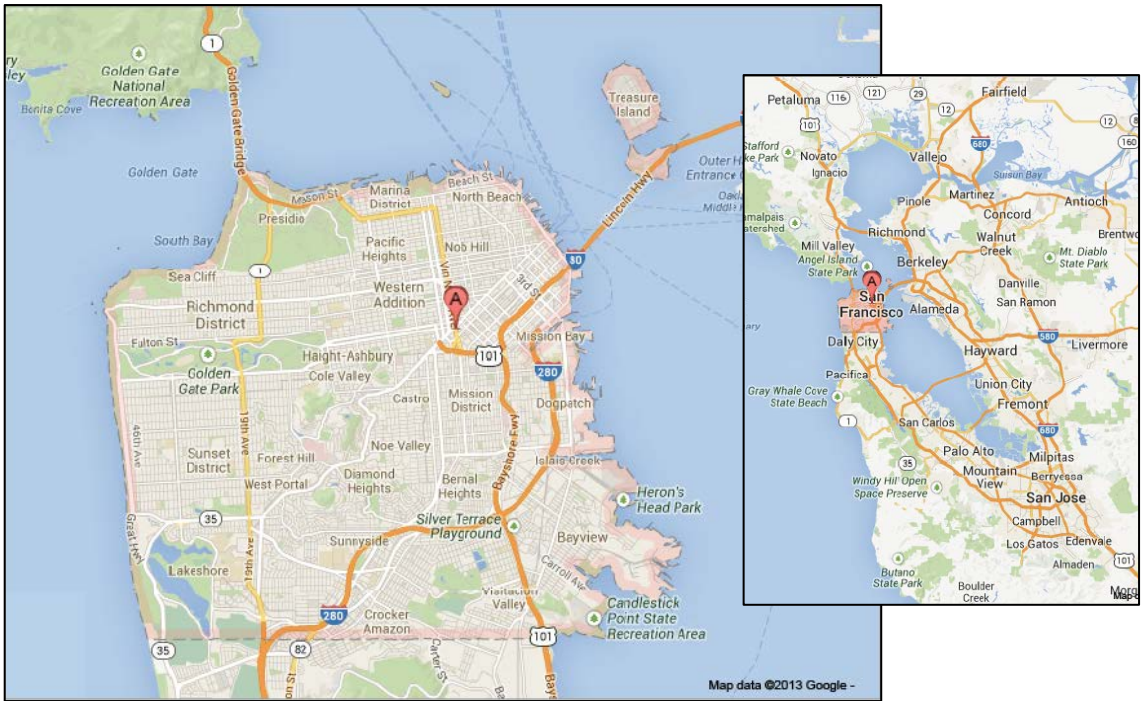


**FIGURE 1** City of San Francisco (left) and location in broader metropolitan area (right)

The San Francisco County Transportation Authority (the Authority) is the sponsor of this project. The Authority administers proceeds of a local sales tax dedicated to funding transportation improvements, and staff provide analytical support to planning projects throughout the city. In order to meet local goals for green house gas reduction, economic vitality, and livability, San Francisco is considering effective demand reduction strategies such as cordon-based congestion pricing as well as investment in cost-effective transit efficiency strategies such as bus rapid transit (BRT).

For over a decade, the Authority has used an activity-based travel model, known as SF-CHAMP, for project evaluation (Jonnalagadda, et al 2001; Outwater and Charlton 2006; Erhardt, et al 2008). While SF-CHAMP has generally served the agency's needs well, the lack of realism in the static assignment model it uses has been limiting for certain applications.

© AET 2013 and contributors

Therefore, the Authority sought to develop a citywide dynamic traffic assignment (DTA) model.

At a project-level, planners have found DTA to be a useful post-process to static traffic assignment to (1) understand the effect of projects that can be measured at the mesoscopic level, and (2) serve as a tidier transition from static traffic assignment into a traffic microsimulation model. Because many of the projects being evaluated are specifically designed to reduce the auto demand, it is important that the demand model be able to provide new trip tables to the DTA model for each scenario.

At a higher level, planners are interested in incorporating DTA's ability to more accurately evaluate reliability, transit/auto interactions, and operational treatments into an integrated modeling framework that is capable of evaluating changes in travel behavior (an integrated activity-based demand model and dynamic network model).

## 1.2 Overview of Current Practice

Over the past 10 years, DTA has transitioned from a research endeavour to a viable tool that is commercially available in a number of packages. Each package uses a somewhat different approach, but once a package is selected, the assignment methods are largely established (Chiu, et al 2011). Therefore, the majority of the effort for a practical implementation lies in developing the inputs (networks and trip tables) and parameters (traffic flow parameters and various other settings) to be used.

Commonly, networks will be imported from a static model, a base layer, or some other primary source. Then intersection detail and traffic control information is added to the network. This typically leads to a moderate to large amount of manual network editing to code missing intersections or resolve errors and special cases.

Trip tables are typically taken from the demand model, and then run through an origin-destination matrix estimation (ODME) process to adjust the demand such that it better fits the traffic counts. This gives calibrated trip tables correspond to conditions in the base-year (Balakrishna, et al 2007; Zheng, et al 2013).

The result is a stand-alone DTA model that can be operated independently of the demand model to test network alternatives. Our own initial DTA

© AET 2013 and contributors

implementation for the northwest quadrant of San Francisco largely followed this formula (Xyntarakis, et al 2010, Sall, et al 2010).

Recently, projects have sought to integrate DTA with activity-based travel models. This includes large projects in Jacksonville and Sacramento which explored approaches to feedback and the stability of convergence in the integrated models (Castiglione, et al 2012; Cambridge Systematics, et al 2010). They do not, however, deal with the challenges of working in a congested urban environment.

### 1.3. Model Development Approach

Based largely on our needs for project applications, we were left unsatisfied with the approach of building a stand-alone DTA model and using ODME to calibrate the demand. There are several problems that arise from that approach.

On the network side, maintaining two copies of the same network in different formats leads to divergence of the networks and challenges in managing network coding. The reliance on manual editing inevitably leads to mistakes on a large network.

On the demand side, ODME is known to result in non-unique solutions. It is also not clear whether the same matrix adjustments can be safely applied to future year trip tables or to trip tables from a scenario with significantly different demand. Finally, ODME serves to mask any underlying errors in either the demand model or the DTA model. For example, in our initial northwest quadrant DTA implementation, the matrix adjustment process dramatically increased the trips to the Presidio, a former military base that has been transitioned to civilian use. Investigation revealed a mistake in the land-use file in which the employment in the Presidio was left out. Once the mistake was corrected, the model results aligned much better and the ODME was no longer necessary. While we are open to its use as a diagnostic tool, our preference is to identify these types of issues and correct them at their source rather than mask them with ODME.

For these reasons, the approach taken for this project is to establish an automated pipeline of information from SF-CHAMP to the San Francisco DTA model (SF-DTA). The networks are converted for each scenario, such that a project coded in the static network would be represented in an equivalent way in SF-DTA. Manual coding is avoided and any special cases are handled in an override file such that they are documented and reproducible. The

© AET 2013 and contributors

demand is extracted from SF-CHAMP for each scenario, allowing the demand effects of the project to be reflected in SF-DTA. ODME is not used, which means that the demand is consistent with SF-CHAMP and future scenarios or scenarios which significantly change the demand can be run without worry about whether the same matrix adjustments would apply.

The result is that project evaluations will use a paired application of these two models. It is not an integrated model because there is no feedback from DTA to the demand model, but it does lay the groundwork for future integration by ensuring that the information flow in the forward direction is uncorrupted.

## 2. MODEL DEVELOPMENT

This section describes the development process for building the citywide DTA model. A key challenge in developing a DTA model is that a mesoscopic simulation of a congested grid network is highly sensitive to network coding and decisions as small as how centroid connectors are placed. As such, the team collaboratively developed a toolset in order to accurately capture, manage, and translate information about every traffic signal, stop sign, roadway, and transit vehicle operating in the City of San Francisco. This section describes that toolset, as well as the derivation of traffic flow parameters, which were estimated from real data as much as possible. After an initial model was up an running, a range of adjustments were made during model calibration, related primarily to the model's representation of certain complexities of the urban environment. Finally, the model validation results are summarized.

### 2.1. Model Overview

SF-DTA utilizes the Dynameq package for assignment, which our past experience has shown to produce reasonable results, have a mature user interface, and a responsive developer.

The model covers all of San Francisco, as shown in FIGURE 1. The road network includes every street in the city, with 976 traffic analysis zones (TAZs) and 22 external stations. The network includes actual phasing and timing plans for nearly every traffic signal in the city, about 1,115 in total (certain mid-block or pedestrian-only signals are ignored). It also includes stop control at about 3,726 intersections based on a GIS layer of stop signs in the city.

The demand for the San Francisco DTA is taken from a subarea extraction of the trip tables produced by SF-CHAMP. SF-CHAMP covers the entire nine-county Bay Area, which is why the subarea extraction is necessary. The

© AET 2013 and contributors

demand is currently segmented into four user classes: autos, trucks, autos with toll, and trucks with toll. There are currently no tolls within the city, so the toll matrices serve as placeholders for congestion pricing scenarios.

The simulation is currently run for a PM peak period. Demand is loaded from 2:30-7:30 pm, with the simulation continuing until all traffic clears. The five-hour demand period includes a one-hour warm-up period, a three-hour P.M. peak period, and a one-hour cool-down period. With warm-up time, the reliable period of results is approximately 3:30-6:30 pm, corresponding to the three-hour P.M. peak assignment in SF-CHAMP.

## 2.2. Software, Network and Trip Table Development

The inputs to the DTA model are prepared using the "DTA Anyway" codebase - a set of open-source tools we developed for the purposes of this project. Key functionality includes the ability to: read and write Dynameq ASCII files, read static network files (in the form of a Cube network), write shapefiles of links and movements (essential for debugging), and other network manipulation tasks (i.e., splitting links, finding links and movements according to various attributes, setting movement priority hierarchies for an intersection, etc.). It is designed to be extensible to other packages and other projects.

FIGURE 2 shows the core scripts used in DTA Anyway, which include:

- *Defining the DTA scenario and converting the static network.* This includes a number of sub-steps, starting with defining the vehicle classes, traffic flow parameters, and generalized cost functions. The static network is read and attributes such as capacity are mapped to relevant DTA network attributes, and shape is added to the links. HOV and transit-only lane restrictions are applied. All intersection movements are added unless known turn prohibitions prevent them. Virtual links are inserted between centroids and road nodes, and the centroid connectors are moved to mid-block locations. Overrides for specific lane configurations or operating characteristics are read from a file. Some basic clean up is performed, including fixing overlapping links, links shorter than the longest vehicle length, and deleting unused zones. This step gives a DTA network, but without intersection control specified.
- *Importing transit routes.* The Cube TP+ transit network files are read and those busses and light rail vehicles operating on-street are converted to Dynameq format. This allows their effect on congestion to be counted, and also allows the model to output the travel times for the bus routes in different scenarios—an important performance measure in San Francisco.

© AET 2013 and contributors

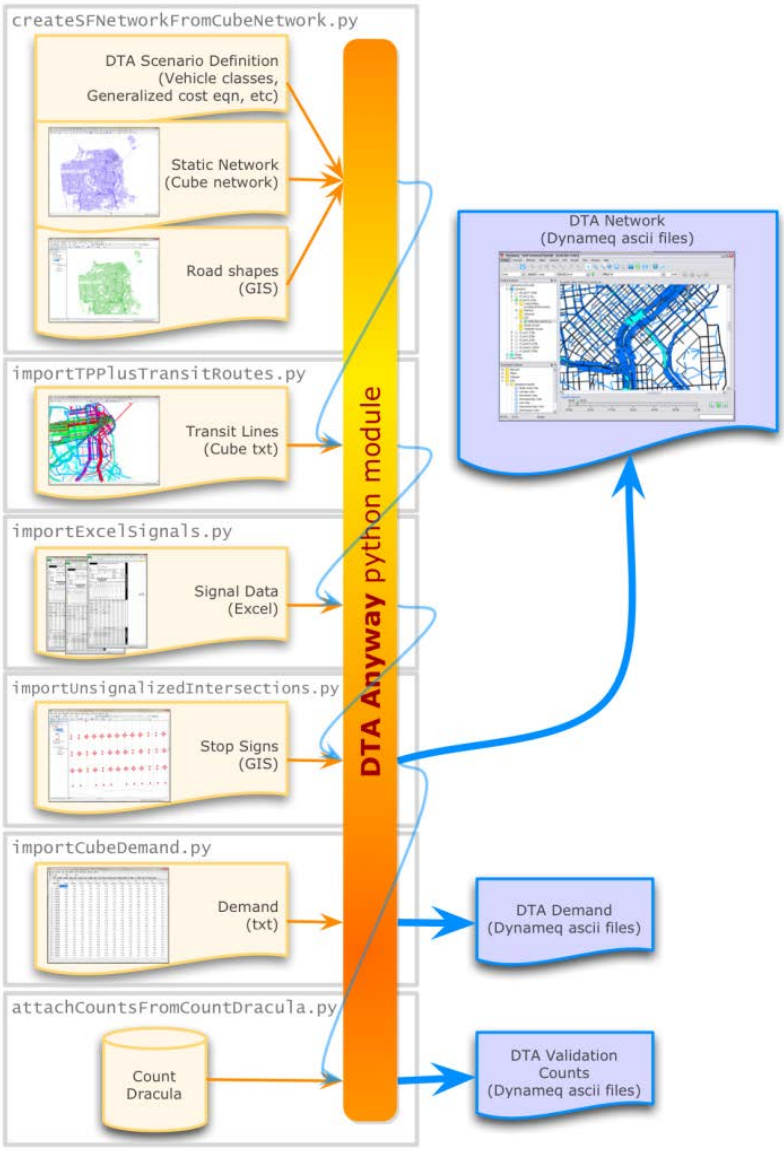


**FIGURE 2** Overview of the script sequence used to create the San Francisco DTA network

- *Importing signals.* The San Francisco Municipal Transportation Authority (SFMTA) maintains a spreadsheet showing the signal phasing and timing for each of the 1,100 traffic signals in the city. Known as the signal cards, these are updated when the signal timings are set in the field. DTA Anyway, reads these signal cards converts them to Dynameq format, and matches them to the appropriate intersections and movements.
- *Importing stop signs.* Stop sign control is added after the signalized intersections. The location and type (all-way or two-way) of stop signs is read from a shape file, and matched to the network. For two-way stops, the priority is defined based on facility type where they differ, and based on a custom priority where they are the same.

© AET 2013 and contributors

- *Extracting subarea demand.* To generate dynamic demand matrices, a subarea extraction is run using the static assignment model. This is necessary because the static model covers a 9-county region, whereas the DTA model is only for San Francisco. Midday (9:00 AM to 3:30 PM), PM peak (3:30 PM to 6:30 PM), and evening (6:30 PM to 3:00 AM) static assignments are run for this purpose, and a peaking profile is applied based on counts taken in 15-minute increments.
- *Importing counts.* Traffic counts are read from the Authority's count database, known as CountDracula, and associated with the links or movements in the DTA network.

### 2.3. Traffic Flow Parameters

Among the primary aspects of the model that can be calibrated are the traffic flow parameters. In Dynameq, the fundamental flow-density diagram is approximated in a simplified triangular form, as shown in FIGURE 3. The parameters shown can be defined from this curve, including the free-flow speed, the saturation flow rate, the response time, the backwards wave speed and the jam density (inverse of the effective vehicle length).

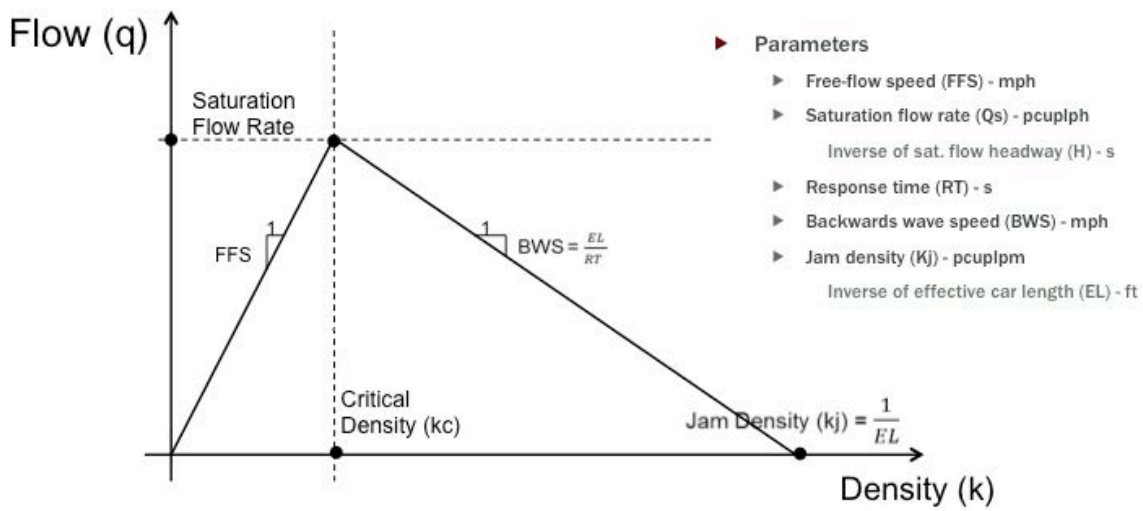


**FIGURE 3** Flow-density relationship used in Dynameq

A valuable data resource for parameter estimation in this study is the inductance loop detector information from the Caltrans Performance Measurement System (PeMS), which collects lane-by-lane counts, speeds and densities in 30-second increments on Bay Area freeways. As illustrated in FIGURE 4, linear regression was used to fit lines through the uncongested and congested phases of these observed flow-density data to define the triangular flow-density relationship used in the model.

© AET 2013 and contributors

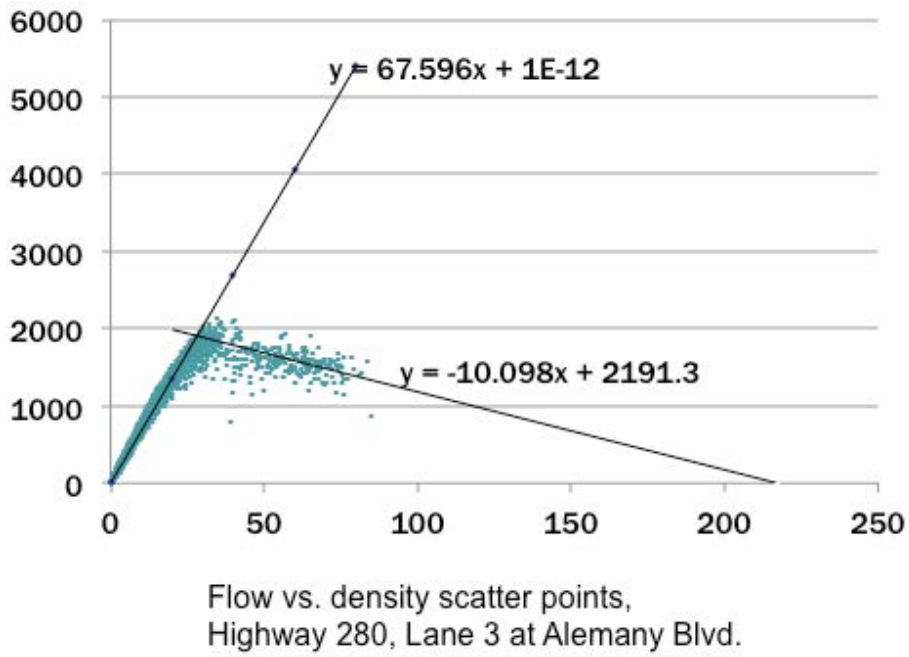


**FIGURE 4** Example of estimating traffic flow parameters from loop detector data

Off the freeways, loop detector data is not available, so the traffic flow parameters were derived using field observations of the effective vehicle length of queued vehicles, the saturation flow rate of vehicles departing at traffic signals, and the free-flow speed on uncongested segments.

## 2.4. Calibration

The model calibration process involved conducting a series of base-year model runs, with an evaluation against observed data for each. Based on that evaluation, adjustments were made to the model in an effort to improve its expected performance. Examples of the types of changes made include: fixing network coding errors, handling special cases for traffic signal inputs, adjusting traffic flow parameters, and changing the generalized cost function.

One pattern observed in calibrating a congested DTA model is that a backup in one location can cause the entire network to break down. FIGURE 5 shows an example of such an instance, where a problem at a single intersection (identified with the blue arrow) caused gridlock (shown by the thick red lines) to extend across the entire downtown area. In this example, the problem was that a traffic signal with an optional pedestrian phase was imported assuming that the pedestrian phase was always at its maximum length. Much of the effort of calibrating the model involved identifying and resolving these issues.

© AET 2013 and contributors


**FIGURE 5** Example of gridlocked area

Several important issues were found and addressed through the calibration process. These include:

- *Speed-flow parameters.* While the data sources provided some guidance on setting the speed-flow parameters, there was some additional testing that was required to determine the correct vehicle length and response time values. The model was found to be particularly sensitive to the truck vehicle length, and it was determined that driving behavior on urban freeways such as US-101 can be much more aggressive than the default settings initially used.
- *Traffic dispersion.* Several early calibration runs revealed traffic volumes much lower than the corresponding counts. This result was surprising, because the static assignment model produced reasonable results with the exact same demand. Further investigation revealed that traffic in the DTA model was being dispersed much more broadly throughout the grid network than traffic in the static model was, as shown in FIGURE 6. To resolve this issue, changes were made to better account for the acceleration and deceleration time at stop-controlled intersections, and to reflect drivers' preference for staying on main routes.

© AET 2013 and contributors

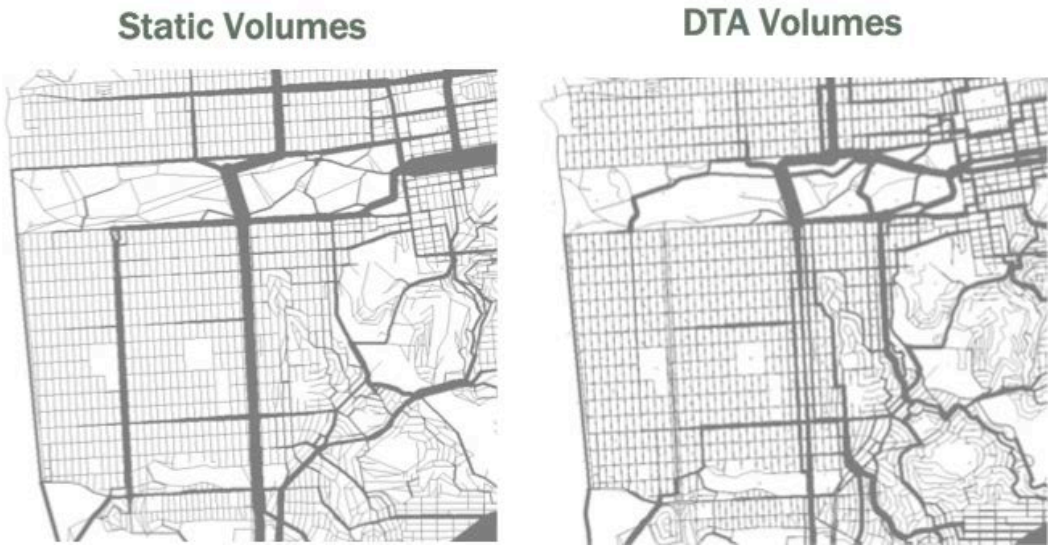


**FIGURE 6** Early example of traffic dispersion through grid network

- *Pedestrian friction.* In much of the downtown area, turning vehicles must wait for a herd of pedestrians to cross before making their turn. Not only is there a delay to that vehicle, but to the vehicles behind it. To serve as a proxy for these conditions, the traffic flow parameters were adjusted as a function of area type, with the denser area types given a larger penalty to account for heavier pedestrian traffic.
- *Transit-only lane permissions.* The transit-only lanes in San Francisco are restricted to allow transit vehicles as well as vehicles making right turns. Dynameq does not currently allow movement-specific vehicle class restrictions, so this was initially modeled by excluding all autos. To better simulate the true condition, a link-splitting method was introduced. In this method, all vehicles are allowed to use the "front" half of a link approaching an intersection, but only transit vehicles are allowed to use the bus lane on the "back" half of the link. Therefore, autos cannot make a through movement from the bus lane. FIGURE 7 shows an example of this coding.

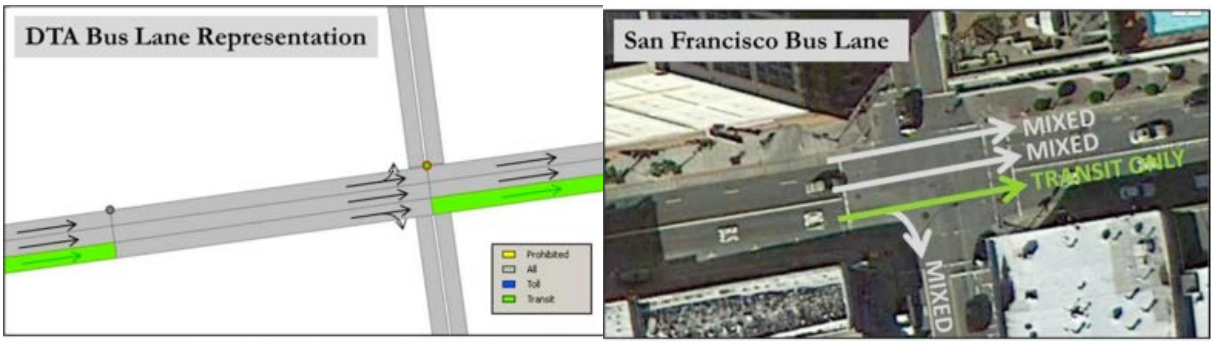


**FIGURE 7** Representation of bus lanes and accommodation for right-turning autos

© AET 2013 and contributors

- *Temporal profiles.* The disaggregate output from the activity-based demand model contain a rich set of information for each trip, including purpose, vehicle occupancy, departure time and value-of-time. The subarea extraction process, however, involves aggregating this information into a limited number of matrices for the static assignment. In this process, there is a significant loss of information, particularly in regards to the ability to apply different temporal profiles to trips with different purposes or specific departure time constraints. For this project, a uniform departure time profile was applied to all trips by vehicle class (auto versus truck) based on counts taken in 15-minute increments. The effect of this assumption on the model results is unclear, as it may be the case that certain areas, such as the financial district, have trips that all leave within a short time of each other creating shockwaves in the system.

## 2.5. Validation

The calibrated model was validated against observed traffic counts and speeds, as well as by visual inspection, as described here.

First, the model convergence was evaluated. While static assignment models can achieve very tight relative gap measures, the stochasticity inherent in the traffic simulation means that a relative gap less than about 5% can be considered sufficient. FIGURE 8 shows the critical gap convergence for the final calibration run. Across time periods, the maximum gap is 6.2%, and the mean gap is 2.7%, with stable gaps for almost 20 iterations.

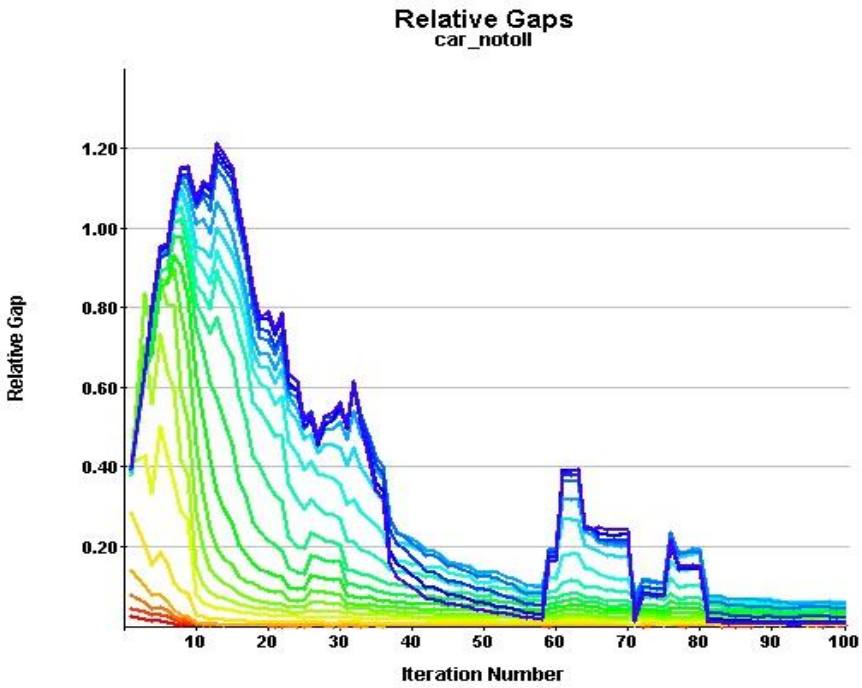


**FIGURE 8** Relative gap convergence plot for final base-year assignment

© AET 2013 and contributors

FIGURE 9 shows a scatter plot comparing the modeled and observed volumes for all links with counts. The yellow lines show the California Department of Transportation (Caltrans) targets for maximum desirable deviation, and the black lines show a linear curve fit through the data. Caltrans recommends that 75% of count locations on freeways and principal arterials fall within the target. In the calibrated DTA model, 76% of major arterials and 77% of minor arterials fall within the maximum desirable deviation.

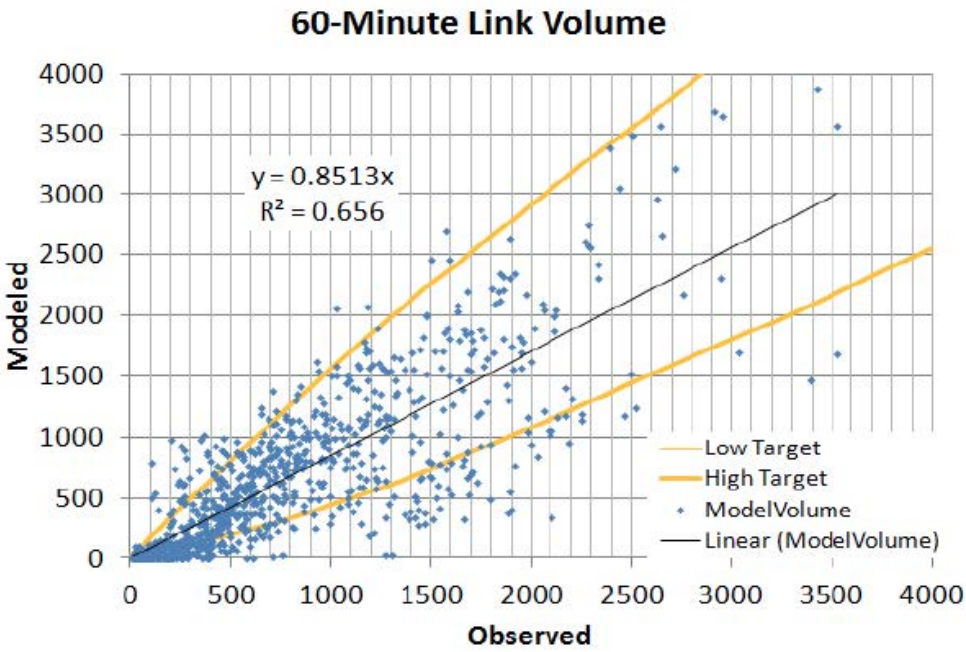


**FIGURE 9** Modeled versus observed link volumes

FIGURE 10 shows a scatterplot of observed vs. modeled route travel times for one of the calibration tests. There is a general trend toward modeled speeds being slightly lower than observed. This result is driven by a handful of outliers, whose location suggests it may be related to queueing on freeway ramp approaches.

© AET 2013 and contributors

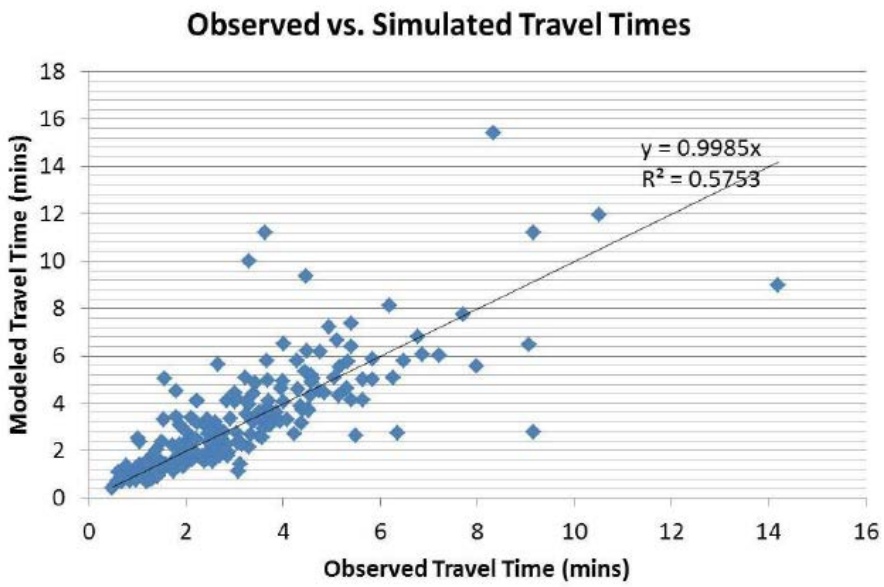


**FIGURE 10** Modeled versus observed travel times

As a basic check of the flows for the network as a whole, we compared bandwidth maps of the DTA results to equivalent output from SFCTA's static model (FIGURE 11). The DTA model has the same overall pattern of flows as the static model, providing some level of comfort.

**FIGURE 11** Bandwidth plots of assigned link volumes (static assignment on left and dynamic assignment on right)

## 2.6. Stability Testing

As part of the model validation process, the model was tested to ensure stable results with small changes to the system. This included testing a change to

© AET 2013 and contributors

the random number seed used by the traffic simulation and testing a trivial network change. Both tests resulted in changes in vehicle miles travelled (VMT) on the order of 0.01% and changes in vehicle hours travelled (VHT) on the order of 1%. Most link volumes changed by less than 100 vehicles over the course of the simulation, although the maximum change was 600 vehicles for the random seed test and 900 vehicles for the network change test, both occurring on high-volume freeways.

While fairly small, such differences should be considered when evaluating projects that are expected to produce a small change in the network as a whole. Additional strategies are being investigated to maximize the stability of the models, such as running the model multiple times and average the results.

## 3. SCENARIO TESTING

Several policy scenarios were run through both SF-CHAMP and SF-DTA to understand the types of results that can be expected and how the DTA results differ from those of the static assignment. None are official policies of the Authority, but were selected to represent the types of applications that could be run.

### 3.1. Bus Rapid Transit

The BRT Scenario test performed was the addition of a center-running BRT on Mission St in both directions between 14th St and Cesar Chavez Street in the Mission District, reducing the auto traffic to one northbound lane, with no southbound lanes. In the base scenario, there are no transit-only lanes on Mission. Previously, Mission St included two lanes in each direction. In order to account for this change, South Van Ness Ave was also modified from providing two lanes in each direction for auto traffic to providing one northbound lane and three southbound lanes. FIGURE 12 shows the lane configuration changes, where red indicates fewer lanes and blue indicates additional lanes.

© AET 2013 and contributors


**FIGURE 12** Lane changes near Mission Street for BRT scenario (red links lost lanes, blue links added lanes)

The demand model predicts an increase of about 1,400 boardings on the relevant bus lines, drawing both from competing rail service and from auto travel. FIGURE 13 shows a map of the change in highway volumes from both the static assignment model and the DTA model. Both show auto diversion off Mission Street where there are lane reductions and onto parallel routes, but the DTA model predicts significantly more diversion for this scenario. This is likely because the DTA model does not allow the volume to exceed the capacity for the chokepoint locations in the corridor, so must respond by re-routing traffic. In addition, the static model does not reflect the constraints of the signals, especially those at the north end of the project, highlighted by the black circles.

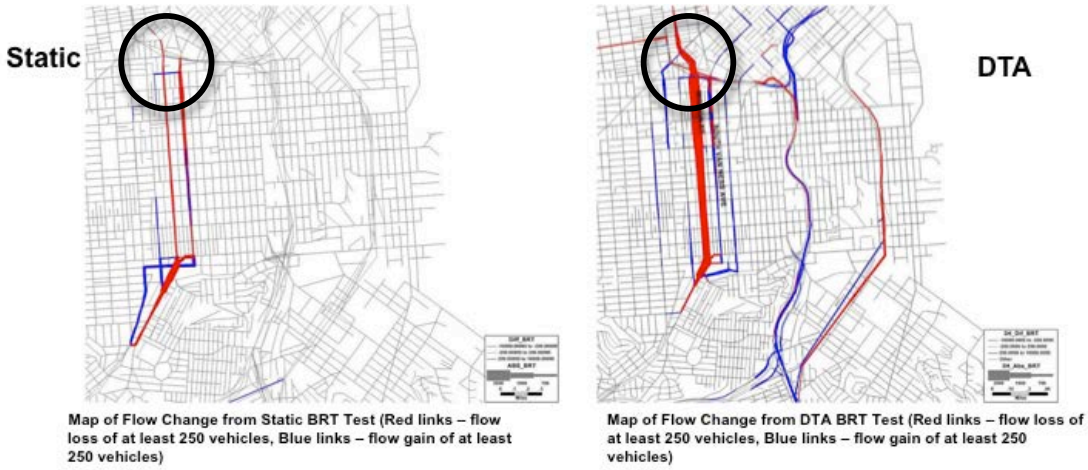


**FIGURE 13** Flow difference for BRT scenario (static assignment on left and dynamic assignment on right)

© AET 2013 and contributors

### 3.2. Congestion Pricing

The congestion pricing scenario institutes a cordon circling the northeast quadrant of San Francisco, as shown in FIGURE 14. Any vehicle crossing the cordon during the PM peak period is subject to a $3 congestion fee (in 2012 US dollars). SF-CHAMP is used to predict the demand response, and the cordon charge is reflected in SF-DTA using the value-of-time from the generalized cost function.

**FIGURE 14** Congestion pricing cordon subject to $3 fee (in 2012 US dollars) during the PM peak period

The demand model predicts an approximately 8% drop in total auto trips across the DTA model subarea for the simulation time period. This includes changes in trip making and tour formation, changes in destination, mode shifts and time-of-day shifts. The static assignment model predicts a 19% decrease in VMT and a 22% decrease in VHT, whereas the DTA model predicts a 15% decrease in VMT and a 19% decrease in VHT. Both use the same demand matrices, so the VMT difference between the two models is due to more circuitous routes in SF-DTA.

Figure 15 shows the auto the flow changes predicted by static and DTA models. While both models show a shift in flow away from the pricing cordon, the static model has less diversion to local roads in the central region and also displays some counter-intuitive flow shifts on eastbound links in the northern region, and a handful of large flow decreases without a clear explanation. The DTA model predicts “smoother” changes in these areas, some of which may

© AET 2013 and contributors

be due to its inclusion penalties on left and right turns, which incentivize linear rather than zig-zag paths through the network.

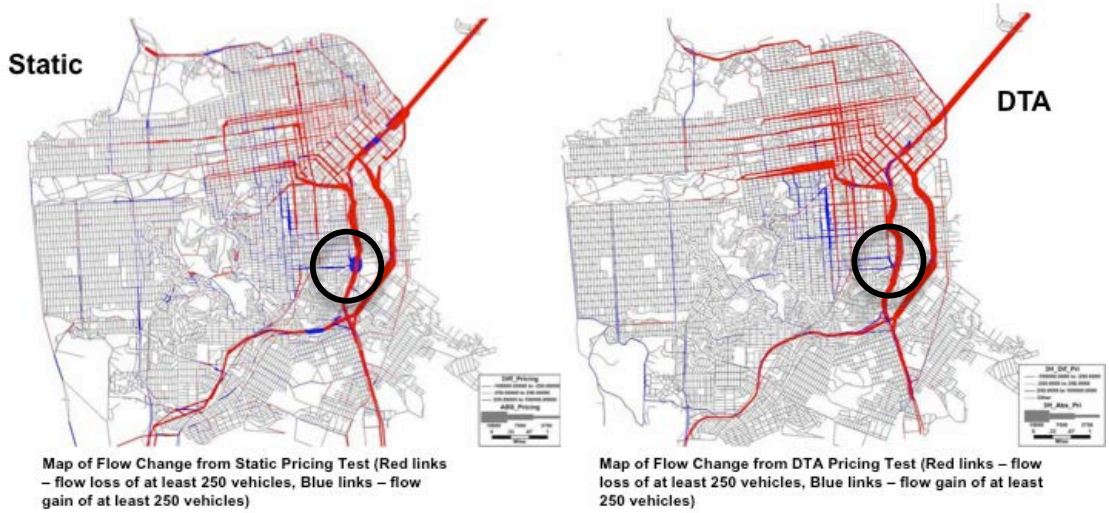


**FIGURE 15** Flow difference for congestion pricing scenario (static assignment on left and dynamic assignment on right)

Also noteworthy is the last freeway ramp before crossing the pricing cordon, highlighted by the black circles. The static model predicts large volume increase on this ramp, shown in blue, causing the volume-capacity ratio to exceed 6. The DTA model does not predict this over-assignment, and instead routes traffic elsewhere.

### 3.3. General Purpose to High-Occupancy Vehicle Lane Conversion

A third scenario involved converting general purpose lanes to high occupancy vehicle (HOV) lanes along I-80 in the Northeast region of the network. This section of freeway is the approach to the Bay Bridge, and among the most congested segments in the region.

Both the baseline and build scenarios were run for 2020 conditions. While the 2020 baseline test converged after approximately 100 iterations and was able to clear all vehicles from the network, the 2020 HOV scenario was not converged after 150 iterations and did not appear likely to converge. Several steps were taken to investigate the failed convergence, with the conclusion that the reduction in capacity was simply too great relative to the reduction in demand predicted by SF-CHAMP. This result highlights the need for feedback from SF-DTA to SF-CHAMP, in order to get a bigger response out of the demand model.

© AET 2013 and contributors

## 4. CONCLUSIONS

This project demonstrates the feasibility of a citywide DTA model for a congested urban network. The DTA Anyway toolkit involved a significant upfront investment in software, but proved to be valuable by making the input development replicable and providing a framework where numerous details can be traced back to their source when issues were encountered in model calibration. The direct extraction of demand matrices from SF-CHAMP allowed the model to be readily applied to scenarios—such as congestion pricing or bus rapid transit—that are expected to change the demand. While OD matrix estimation would produce a tighter fit against count data, the model produces reasonable results without matrix estimation.

The immediate plans for this model involve using it for project applications, some of which may involve more detailed corridor-specific calibration and validation. In the longer term, a plan has been identified to move towards feedback from SF-DTA to SF-CHAMP. Having an established data pipeline from SF-CHAMP to SF-DTA lays the groundwork for that future feedback, although several intermediate steps are needed, such as implementing DTA for a full 24-hour period.

The remainder of this section describes some lessons we learned during the project and our associated recommendations for future developers, as well as several open questions that warrant further investigation.

### 4.1. Lessons Learned and Recommendations for Future Developers

Perhaps the biggest challenge faced developing this model was the high level of sensitivity in the model due to small changes in the model inputs or parameters. Changes in capacity (such as the treatment of transit only lanes) or traffic flow parameters (such as the vehicle length) caused the model to shift between predicting gridlock and predicting relatively smooth flowing conditions. It is thought that this is due, at least in part, to the high level of congestion in the real world network and the inherent instability of heavy traffic flow, which the model is capturing. Nonetheless, it is a less forgiving model to work with than a static assignment model. Therefore, we recommend that others working in congested urban environments allow sufficient time to resolve gridlock issues, and also test a range of model parameters to better understand which are most critical to DTA on their specific network.

Along those same lines, we learned greatly from the stability testing and scenario testing conducted for this project. While showing the fit against traffic

© AET 2013 and contributors

counts is useful, it is not sufficient for understanding how the model will perform for project evaluation. Therefore, we recommend that stability testing and scenario testing be treated as an integral part of model validation.

The DTA model did produce notably different results in scenario testing than the static assignment model, especially where a specific intersection or ramp served as a chokepoint. In this regard, the DTA model results appear quite reasonable, although we do not have data to know the "right" level of diversion that would occur for these scenarios. The trade-off is that it is possible to develop scenarios, such as the HOV lane conversion, that simply do not converge, leaving somewhat of a dilemma as to what to report to the planning team about the scenario.

While we found modelling the City of San Francisco, rather than the entire 9-county metropolitan area, to be plenty of geographic scope for our purposes, the subarea extraction process caused an important loss of information from the activity-based model outputs. This is particularly limiting with respect to understanding the departure times of individual trips. We recommend (and intend to pursue) developing a refined technique to allow the subarea extract to be related back to the disaggregate demand model outputs, thus allowing the DTA trip tables to be created across more flexible range of dimensions.

### 4.2. Open Questions and Topics for Future Investigation

Several issues arose for which we do not have a clear answer or recommendation, but instead serve as topics for further investigation.

Further investigation is needed to better understand how the stochasticity observed in the stability tests relates to the magnitude of changes for smaller projects being evaluated. If the stochasticity turns out to be a significant portion of the project difference, strategies are needed to manage it. One option may be to run the model multiple times and average the results, and another may be to cut a smaller subarea for specific projects.

The model calibration process used worked for the project, but it is labor intensive and requires quite a bit of analyst judgment about what to investigate or try next. Further thought is required to identify strategies to make the model calibration process more systematic. Specifically, we wonder if it may be more valuable to calibrate the model to match observed link density, rather than traffic flow. The fact that the same flow occurs at two different density points on the traffic flow curve (FIGURE 3) means that when calibrating to counts it is not immediately clear whether the flow is on the

© AET 2013 and contributors

congested or uncongested phase of the curve. Such an approach would require a shift in how we think about what a model should produce.

Finally, it is not clear what to do when a model scenario results in gridlock. This may be a sign of the need for feeding the DTA travel times back to the demand model in order to cause a reduction in demand. In the current situation where feedback has not been implemented, what can be said about the scenario, other than that it does not work? In a future case in which feedback has been implemented, a gridlocked DTA model still will not produce reasonable travel times. In such a case, demand model results based on unreasonable travel times cannot themselves be trusted, so what should be fed back? Perhaps a method of successive averages will resolve the issue, but it would be valuable to identify a more efficient approach.

## ACKNOWLEDGEMENTS

This work was funded, in part, by a grant from the US Federal Highway Administration. The authors acknowledge the contributions of Michalis Xyntarakis to the formative stages of this project. Michael Mahut, Brian Gardner, Jim Hicks and Joel Freedman served as valuable resources to the development. A peer review panel, including Joe Castiglione, Bruce Griesenbeck, Vassillis Papayannoulis, David Stanek and Xueson Zhou provided important feedback improve the final model.

## BIBLIOGRAPHY

Balakrishna, R., Antoniou, C., Ben-Akiva, M.E., Koutsopoulos, H.N. and Wen, Y. (2007) Calibration of microscopic traffic simulation models: methods and application. Transportation Research Record 1999, 198-207.

Cambridge Systematics, Inc., Sacramento Area Council of Governments, University of Arizona, University of Illinois at Chicago, Sonoma Technology, Inc. and Fehr & Peers Associates (2010) SHRP C10B – Partnership to develop an integrated, advanced travel demand model and a fine-grained, time-sensitive network model design plan. Project report prepared for Strategic Highway Research Program.

Castiglione, J., Grady, B., Outwater, M., Lawe, S., (2012) Sensitivity testing and assessment of an integrated regional travel demand and traffic microsimulation model. Presented at the 91st Annual Meeting of the Transportation Research Board, Washington, DC.

© AET 2013 and contributors

Chiu, Y.-C., Bottom, J., Mahut, M., Paz, A., Balakrishna, R., Waller, T. and Hicks, J. (2011) Dynamic traffic assignment: a primer. Transportation Research Circular Number E-C 153.

Erhardt, G.D., Charlton, B., Freedman, J., Castiglione, J. and Bradley, M. (2008) Enhancement and application of an activity-based travel model for congestion pricing. Presented at Transportation Research Board Innovations in Travel Modeling Conference, Portland, Oregon.

Jonnalagadda, N., Freedman, J., Davidson, W.A. and Hunt, J.D. (2001) Development of microsimulation activity-based model for San Francisco: destination and mode choice models. Transportation Research Record 1777, 25-28.

Outwater, M.L. and Charlton, B., (2006) The San Francisco model in practice: validation, testing, and application. Proceedings from the Innovations in Travel Demand Modeling Conference, Austin, Texas.

Sall, E., Juran, I., Xyntarakis, M., Mahut, M., Charlton, B., Florian, M. and Velan, S. (2010) Rapid implementation and validation of dynamic traffic assignment in the Doyle Drive corridor, the southern approach to the Golden Gate Bridge. Presented at the 3rd Innovations in Travel Modeling Conference, Tempe, AZ.

Xyntarakis, M., Sall, E., Hicks, J. and Charlton, B., Regional dynamic traffic assignment for real world travel demand models – trying it out in San Francisco. Presented at the 89th Annual Meeting of the Transportation Research Board, Washington, DC.

Zheng, H., Nava, E., Chiu, Y-C, Xyntarakis, M. and Sabina E.E. (2013) Development and calibration of a daily regional dynamic traffic assignment model – a case study for Denver region. Presented at the 92nd Annual Meeting of the Transportation Research Board, Washington, DC.

© AET 2013 and contributors