Metropolitan Washington Council of Governments
National Capital Region Transportation Planning Board

State of the Art in Equilibrium Traffic Assignment

November 16, 2007

Vanasse Hangen Brustlin, Inc.
8300 Boone Blvd, Suite 700
Vienna, VA 22182
(703) 847-3071

Executive Summary

The Metropolitan Washington Council of Governments, National Capital Region Transportation Planning Board (TPB) engaged Vanasse Hangen Brustlin (VHB) to research the State of the Art in equilibrium traffic assignment as a follow-up to the FY 2006 research on the MPO State of the Practice on traffic assignment as a whole.

The FY 2006 research found that while a majority of MPOs nationally use equilibrium assignment, there are several outstanding problems with the equilibrium method, including a failure to reach closure within a reasonable number of iterations and instability in the assigned results. These problems with the widely-used Frank-Wolfe (F-W) equilibrium algorithm have been documented by other researchers, and further research has yielded some new algorithms that potentially overcome the issues with the F-W method. These new algorithms are just now starting to appear in major travel demand forecasting software packages and being applied by MPOs and other agencies. This memo summarizes our findings on the status of the State of the Art in equilibrium assignment and provides guidance for TPB going forward with model development on the specific issue of traffic assignment.

Given TPB's commitment to the Citilabs modeling platform, the next logical step is to pursue whatever run-time and convergence gains can be achieved under the TP+/Cube environment. In order to accomplish this, TPB must convert the v2.2 model to a form that can be used under Cube Cluster. At the time of this memo, TPB has purchased Cube Cluster and has been working with it. Staff has gotten a demo model to run, but have not yet gotten their regional travel model running under Cube Cluster. They must also identify either existing machines for cluster creation, or purchase new hardware. If new hardware is to be acquired, specifications must be created, and an appropriate level of investment for software and hardware upgrades to support future model applications determined. At this time of this memo, TPB has purchased a modeling server (it has two dual-core Xeon processors) to use for both Cube Cluster and general model runs. Any distributed processing work will also be conducted on this machine.

TPB modeling staff should maintain contact with representatives from St. Louis, Minneapolis / St. Paul, and Columbus as a resource throughout this process, and should follow-up immediately with Citilabs to discuss any issues encountered while moving the production model into Cube Cluster. The level of benefits reported by Cube Cluster users should improve the TPB model performance sufficiently while Citilabs implements alternatives to F-W in future versions of Cube. TPB should also consider using the model to test the efficacy of Citilabs' future implementation of any advanced assignment algorithms.

Background: Overview and Current TPB Traffic Assignment Methods

The final element of a traditional four-step travel demand forecasting model is traffic assignment. This step allocates travel demand (vehicle-trips, developed in the previous three steps of the model chain) to a transportation (usually highway and transit) network between origin-destination (OD) pairs according to a specified method. The qualities of a good assignment method include reasonable accuracy, fast and precise convergence, short computing time, and stable results. Among MPOs, the most-widely utilized assignment method is equilibrium assignment, which simultaneously solves for link flow and cost.[1] Equilibrium assignment is predicated on two underlying assumptions: 1) travelers have perfect information on conditions on all possible routes, and 2) travelers always make a rational route choice to minimize their travel time / cost. When the network reaches equilibrium, all trips are assigned to those paths with the minimum impedance (e.g., travel time or travel cost) between each OD pair, and no traveler can improve his or her travel time by switching to an alternate path. Like most MPO models, the current TPB production travel demand model utilizes equilibrium assignment, as will the next production model. Specifically, the TPB model uses the Frank-Wolfe (F-W) algorithm for equilibrium assignment.

Link-based Frank-Wolfe Algorithm

The F-W algorithm, also known as the convex-combination algorithm, is a classic algorithm in operations research and the most widely-applied equilibrium assignment algorithm in travel demand forecasting.[2] The F-W method views the traffic assignment problem as a minimization problem using linear programming. At each step the objective function is linearized and a solution is calculated to reduce the objective. In general, the F-W algorithm performs well during the first several iterations, but it slows down significantly when close to the minimum point (that is, approaching equilibrium) and never reaches its objective function's minimum.[3] Therefore, the algorithm may be best used to find an approximate solution rather than a true equilibrium.[4] Besides the commonly used performance measures – gap, relative gap and average excess cost, the stabilization of link flows from iteration to iteration gives the forecaster some assurance that an adequate approximation has been achieved.

In recent years the F-W method has been widely used for determining the equilibrium flows in transportation networks. Theoretically, true user equilibrium can only be achieved in an artificially small or virtually uncongested network; for a highly congested transportation network, equilibrium can only be closely estimated. Most travel demand forecasting software packages use the F-W method. Compared with other equilibrium assignment methods, it is easy for software developers to code and requires the least computer memory since at each iteration it deals with only a single path between each origin-destination pair.[5]

---

[1] See MWCOG (2006) and Spielberg and Shapiro (2006).
[2] See Boyce, *et al* (2004) and Slavin, *et al* (2006).
[3] See Dial (2006).
[4] The level of approximation is inversely related to the number of assignment iterations; that is, more iterations bring the solution closer to a true equilibrium, and therefore directly related to the level of computational power used to run the forecasting model.
[5] Dial, *ibid.*, and Jayakrishnan, *et al* (1994).

Several issues with the F-W algorithm have been reported in previous research, such as slow convergence, long computational time, and unstable assignments, in which a relatively small change to the travel network or other conditions produces unexplainable results across the whole network.[6] According to Wolfe, the unsatisfactory performance occurs because the search direction tends to become orthogonal to the steepest descent direction as the optimum solution is approached.[7] In addition, the F-W algorithm has no mechanism to avoid the introduction of cyclic flows. A cyclic network normally contains a cycle, a path from a node to itself, which may be one reason for the slow convergence of the F-W method.[8] The highlighting of all these issues has increased interest in alternatives to the F-W algorithm.

Literature Review / Overview of Emerging Equilibrium Assignment Algorithms

VHB conducted a literature review to obtain information on emerging equilibrium assignment algorithms. There are two major algorithms emerging (or in some cases, reemerging) as potential improvements over the link-based F-W algorithm: path-based and origin-based.

Path-based Algorithm

Different from a link-based solution, a path-based algorithm (also called a route-based algorithm) for equilibrium assignment provides a complete picture of the travel pattern and offers modelers the capability to keep track of the distribution of the O-D flows among the different routes as well as the corresponding turning details.

Path-enumeration algorithms were first proposed in the late 1960s.[9] At that time they were infeasible because of the computing power required to store all utilized paths from all origins to all destinations. As computing power became greater, cheaper, and more available over the last decade, path-based algorithms were re-examined as a solution for equilibrium assignment. The path-based algorithm currently uses a gradient projection method or other algorithms to reach convergence faster and more efficiently.[10] After an initialization with an all-or-nothing assignment, the path-based algorithm searches for other paths with shorter travel time between each origin-destination pair and shifts some traffic from previously identified paths to new shorter paths. In the course of the path search, the paths with zero flow are dropped. Path-based methods are still generally considered more computationally-intensive when compared to link-based and origin-based methods. PTV's VISUM modeling platform includes the option of a path-based algorithm as part of its equilibrium assignment module, as does Caliper's TransCAD software. It should be noted that particular formulations vary between path-based algorithms; for example, the path-based method in VISUM is different than the method used in TransCAD.

[6] See MWCOG (2006).
[7] See Wolfe (1970).
[8] See Janson and Zozava-Gorostiza (1987)
[9] See Dafermos and Sparrow (1969).
[10] See Bertsekas and Gafni (1983).

Origin-Based Algorithm

Origin-based algorithms (OBAs) attempt to retain the advantages of path-based algorithms; e.g., providing immediate route flow interpretation, while further reducing computational requirements. While the solution variables of both the F-W and path-based algorithm are link flow and path flow, the OBA defines the solution variables in an intermediate way between links and paths. The main variables for the OBA are origin-based approach proportions, which allow efficient storage of a complete description of the route flows. The OBA has three distinct advantages:

- The capability to deliver detailed solutions.
- Substantially lower computation time.
- Lower memory requirements compared with path-based algorithms.

Bar-Gera[11] presented and implemented the first OBA, for which the underlying concept is acyclic flows. An acyclic network does not include any cycles. A cycle could be a travel path around a city block or traversing opposite directions on the same roadway segment. The algorithm consists of two main steps: update the restricting acyclic subnetwork and shift flow within the subnetwork. The main solution variables are approach proportions and are updated when the flow shifts from high-cost alternatives to low-cost alternatives. Then the route proportions are determined as the product of approach proportions of all links along the route. Finally, the route flows are calculated using origin-destination flow and route proportion. An acyclic restricting subnetwork is maintained for every origin so that only the links that are included in this subnetwork are assigned approach proportions and unused links are removed. Therefore, only routes that are limited to the links in the subnetwork can be used. The computation efficiency of the OBA results from the following features:

- Boundary search procedure with well-estimated search direction
- Restricted acyclic origin-based subnetwork.
- Origin-specific topological ordering of the nodes.

Boyce[12] applied the Bar-Gera OBA to solve the user-equilibrium traffic-assignment problem in a practical large-scale roadway network in New Jersey. Different alternatives were tested to evaluate the addition of two proposed ramps. Compared to the F-W method, the OBA achieved highly converged solutions with significantly fewer iterations. In an OBA, the routes serving the OD pairs are efficiently identified and optimized because all the destinations for each origin are grouped together. In general, origin-based assignment methods require more computational resources than the F-W method but less than path-based methods. Other experimental results for medium and large model networks have demonstrated that the OBA can efficiently find a highly accurate solution for equilibrium assignment, but that the process still takes an excessive amount of computing time. There is a free, open-source OBA available for download from the Internet.[13]

---

[11] 1999.
[12] 2004.
[13] See http://www.openchannelfoundation.org/projects/Origin-Based_Assignment. The site includes executable code provided by Bar-Gera for the OBA as well as a Chicago transportation network and trip tables.

Dial introduced an updated path-based user-equilibrium traffic assignment algorithm which eliminated the need for path storage.[14] However, Caliper's later research work classified Dial's algorithm as an origin-based equilibrium method because it decomposes the UE problem into a sequence of single-origin problems on acyclic sub-networks or "bushes."[15] Using these simpler sub-networks, it efficiently locates and shifts flow from costly paths to the cheaper paths until the costs of all used paths are within a user-specified range of the cheapest path. Dial's algorithm has several primary benefits:

- It avoids oscillation when approaching equilibrium and achieves a precision unreachable by the F-W algorithm regardless of the network's size and congestion level.
- It avoids explicit path storage and enumeration by restricting attention to a relatively few path segments in sequential acyclic sub-networks; this in turn improves computational efficiency.
- It provides a "warm start" feature to compute a new equilibrium much faster using the solution obtained for a similar, previous problem.

Dial tested and reported the new algorithm's performances on two networks, and in both cases it significantly outperformed both the F-W algorithm and the Bar-Gera OBA. The new algorithm routinely achieves the precision F-W was never able to approach, and it reached a relative gap below 10^{-3} in less time than the OBA.

Building on their previous research, Caliper Corporation has implemented an origin-user equilibrium (OUE) method modified from Dial's "Algorithm B" which demonstrated superior performance in reaching tight equilibrium within much lower computational times.[16] The OUE establishes an order-dependent acyclic sub-network from each origin to all destinations and calculates shortest path more efficiently. During each iteration, the algorithm examines and updates the origin-based link flow to improve travel time. Caliper has done an empirical comparison of alternative traffic assignment methods which demonstrates the OUE method is a promising and feasible solution based on its high level of convergence, fast computing time, and modest memory requirements.[17] Caliper's OUE has been incorporated into TransCAD 5.0, which is currently in beta release and will have a final release this summer.

Status of New Equilibrium Assignment Techniques and Computational Advances among Major Forecasting Software Vendors

The impact of new assignment algorithms on computer hardware requirements and subsequently model run-time remains a chief concern among travel forecasters. Currently it takes approximately 18 hours to run the TPB travel demand model (v2.1D#50) on a computer with a 2.99 GHz dual-core processor with roughly 1GB of memory running Windows XP Professional. An increase in run-time is anticipated with the incorporation of the model improvements planned

[14] 1999 and 2006.
[15] For purposes of organization we have followed Caliper's characterization and included Dial's algorithm with the origin-based methods.
[16] See Slavin (2006).
[17] Ibid.

for the Version 2.3 model.[18] Machines with multi-core processors and/or multiple processors have become more widely available and more affordable in recent years, and travel demand software vendors have been working to take advantage of the increased computing power. While all traffic assignment methods can benefit from multithreading and/or distributed processing when more than one processor is available, there is particular benefit for advanced equilibrium algorithms due to the complexity and volume of calculations required for large networks. VHB contacted the major travel demand forecasting software vendors: Citilabs, Caliper, PTV, and INRO, to discuss their current implementation or plans for implementing both advanced computing processes and advanced traffic assignment methods. In general, the findings are as follows:

- Citilabs will include origin-based assignment in Cube Voyager 5.0, to be released in May 2008; their current release focuses on improving model run-time under F-W assignment by using distributed processing.
- INRO is working on incorporating advanced traffic assignment algorithms and advanced computational methods to its software, but did not say when these features would be available.
- PTV currently includes advanced traffic assignment algorithms in VISUM; advanced computational methods are under development.
- Caliper includes both advanced traffic assignment algorithms and advanced computational methods in TransCAD.

Citilabs

Citilabs' new Cube Cluster reduces run time by distributing modeling tasks across multiple processors. There are two methods to distribute model processes: intrastep distributed processing and multistep distributed processing. The former splits zone-based tasks from a single step into groups based on the availability of processors while the latter distributes the independent steps to available processors. Users may define the cluster range using model scripts. Cube Cluster will run on multiple computers which share Windows files or a computer with multiple processors. However, the hardware setup for Cube Cluster is not automatic. Cube Cluster does not impose scaling restrictions on the number of processors or machines in the cluster, although there are obviously practical limits due to physical space, cost, and other issues. A Cube Cluster license costs $1,500, plus $1,500 per node machine for licenses of Cube Voyager.

PTV

PTV's VISUM modeling software provides a path-based multiclass assignment implemented in the equilibrium procedure to distribute demand over the network. It keeps track of all utilized routes and equilibrates flow over different paths. Equilibrium is reached by multiple iterations based on an all-or-nothing assignment or an incremental assignment as a starting solution. The

[18] Nearly all of the increase in run-time is due to the implementation of the nested logit mode choice model with speed feedback within the v2.3 model; on a machine with a 3.73 GHz Xeon processor and 2GB of memory, run time was reduced to 12 hours compared with the statistics on v2.1D #50 above, but increased to 21 hours when using the nested logit mode choice model.

outer iteration step searches for the new routes in the system (those routes with lower impedances) while the inner iteration step balances the network by shifting vehicles among the competing routes. VISUM's path-based algorithm provides two advantages: 1) it stores the paths for later analysis, and this feature in turn allows for 2) path reloading, where a previous assignment is used as the starting point for a new assignment (so-called "warm starts"). Furthermore, VISUM takes advantage of path compression techniques to improve computing efficiency in both assignment processing and path storage.

The latest version of VISUM also includes a new continuous implicit path formulation for the user-equilibrium assignment problem developed at the University of Rome.[19] This method works with time-varying demand and time-varying supply and is reported to be an excellent choice for mesoscopic modeling with reasonable computational requirements and model run-times. So far PTV has released multithreading functionality for VISSIM microsimulation software, but not for the planning software VISUM. The VISUM development team has started to work on this capability, and it is expected that VISUM will multithread all highway assignment and demand modeling procedures over the next two or three years. VISUM's path-based equilibrium offers level of convergences (10^{-7}) that exceed those used in practice and has path reloading.

Caliper

Caliper has successfully reduced computational time in TransCAD by multi-threading the F-W algorithm with multiple processor or multiple-core machines. In TransCAD, some of the key procedures in travel forecasting are <u>automatically</u> multi-threaded if used in a multiple-processor environment; for example, traffic assignment, which runs almost twice as fast on a dual-processor computer than on a single-processor machine. TransCAD also supports distributed processing or clustering, similar to Citilabs' Cube Cluster. Distributed processing is not automatic due to the complexity and setup and implementation varies with different models. In terms of advanced assignment algorithms, the OBA in TransCAD 5.0 offers exceptionally tight convergence at levels much better than current standard practice.[20] Caliper's OBA also includes the "warm starts" feature, similar to that found in VISUM.

INRO

[19] See Gentile, *et al* (2005).

[20] Boyce (2004) recommends using convergence of at least 10^{-4} to reach stability of link-flow difference for a large-scale network. Anecdotal evidence suggests that most MPOs converge between 10^{-2} and 10^{-4} and / or fix their number of iterations (like TPB) after repeated tests end up in this range. In the recent TRB survey on the State of the Practice, only 32% of respondents indicating that they used equilibrium assignment also indicated their model's closure tolerance; indicating that many MPOs may use the default settings of their modeling software or don't know certain characteristics of their traffic assignment. Of those responding with closure tolerance, 96% indicate a convergence at a gap between 10^{-1} and 10^{-3} Yet even with the capabilities offered by TransCAD, Caliper recommends using convergence at a *minimum* relative gap of 10^{-3} for many applications and a relative gap of 10^{-5} or 10^{-6} for careful project evaluation.

INRO's current major release of EMME/3 does not include support for advanced traffic assignment algorithms nor advanced computing processes.[21] Both of these features are under development, but it is not known when they will be fully implemented in the software. Given EMME's history and foundation built on early adoption of the F-W method, it is reasonable to assume that INRO's next major release will include some implementation of both features in order to keep pace with other software vendors and meet the demands of their user base.

Table 1 shows a comparison of the major software packages and their features.
=

Table 1: Comparison of Major Forecasting Software Advance Assignment and Computing Features

	Cube/TP+	**VISUM**	**TransCAD**	**EMME/3**
Equilibrium Assignment	Link-based Frank-Wolfe Method	Frank-Wolfe Method, Path-based Multiclass Method	Frank-Wolfe Method, Origin-based Method	Frank-Wolfe Method
Stopping Criteria	GAP, RELATIVE GAP, AAD, RAAD, PDIFF, PDIFFVALUE, RMSE*	Relative Gap	Relative Gap	Relative Gap
Computational Capabilities	Cube Clusters	Warm Start	Clusters Multithreading Warm Start	N/A

* *GAP* – Relative difference in system cost between two iterations; RELATIVEGAP – An alternative GAP measure; *AAD* – Average absolute difference in volumes between two iterations; *RAAD* – Relative average absolute difference in volumes between two iterations; *PDIFF* – Fractional portion of links whose change in volume between two iterations is less than the value of PDIFFVALUE; *PDIFFVALUE* – The value to be used with PDIFF; *RMSE* – Root mean squared error of the difference in volumes between two iterations.

To date, both PTV and Caliper (in VISUM and TransCAD, respectively) have implemented assignment methods touted in the literature as converging more rapidly than the conventional link-based Frank-Wolfe method. TransCAD reports excellent runtimes with the origin-based assignment.[22] The performance of this new method in terms of convergence, runtime and network stability will be best reported by the planning practitioners who are actually using it.

The "warm starts" feature is particularly attractive to planning practitioners since it allows them to perform a series of model runs with feedback and analyze multiple scenarios when making slight changes to the land use or network facilities.

MPO Contacts / Use of New Features

VHB contacted several MPOs to discuss either their use of advanced algorithms for equilibrium assignment, or advanced computing options to improve model performance (run-time), or both (if applicable). Because these features are so new and not yet in wide use, additional contacts were made with state departments of transporation (SDOTs) and one county planning agency to

[21] Most of INRO's work developing EMME/3 focused on improving the user interface rather than adding model chain features.
[22] Bar-Gera reported much slower run-times using his own code.

capture the full scope of agencies known to be employing the new assignment techniques. The list of agencies using the new techniques was obtained from the software vendors, and VDOT was contacted due to their interaction with TPB's forecasting work.

Agency	Modeling Platform(s)
Virginia Department of Transportation (VDOT)	Cube
Ohio Department of Transporation (ODOT)	Cube Voyager (Cluster), VISUM
Metropolitan Council of the Twin Cities (Minneapolis / St. Paul MPO)	Cube Voyager (Cluster)
METRO (Portland, Oregon MPO)	VISUM
East-West Gateway Coordinating Council (St. Louis, MO MPO)	Cube Cluster
The Maryland-National Capital Park and Planning Commission, Prince George's Planning Department	TransCAD
Capital District Transportation Committee (Albany, NY MPO)	VISUM

Virginia Department of Transportation (VDOT)

VDOT's modeling group in Richmond currently has no plans to move to Cube Cluster, as 1) they are satisfied with the performance of their individual workstation licenses, and 2) VDOT's "one computer per user" computing policy effectively prohibits use of any cluster or distributed processing feature by anyone except IT staff. VDOT report that their on-call consultant (Corradino) has experienced significant performance enhancement running models using Cube Cluster. Finally, VDOT noted that the latest version of Cube Voyager has the ability to create path databases that store assignment information in a way similar to VISUM for easier select link and other analysis.

Ohio Department of Transporation (ODOT)

ODOT uses Cube Cluster to run a variety of models ranging from small (200 zones) and medium-sized (800 zones) MPO models (written in Voyager and Application Manager) to the Columbus model (1900 zones) and the Ohio statewide model (5000 zones). These last two are both activity-based models written primarily in Java, so Cube (and the distributed processing is used only for the network skims and assignment). Equilibrium closure gap used by ODOT range from 10^{-3} to 10^{-5}. ODOT's run-time gains are so significant that they will not run their large models <u>except</u> under Cube Cluster. ODOT has a significant investment in hardware – approximately $80,000 for a cluster of 9 machines, each with 2 dual-core processors.[23] ODOT's modeling staff reports minimal effort required to begin using Cube Cluster if the users already

[23] ODOT's system is a more robust version of the clusters at the Columbus transit agency (COTA) and MPO (MORPC), where the tour-based models were implemented prior to being used at ODOT and their statewide model developed. More details on those smaller systems can be found in the TPB FY 2006 memo on activity-based models.

are familiar with Cube, and offered to share sample scripts as well as discuss specific issues with TPB staff. VISUM assignment was used on a few projects in the area; however, the ODOT staff VHB contacted could not provide more detailed information on the performance of the software.

Metropolitan Council of the Twin Cities (Minneapolis / St. Paul MPO)

The Metro Council tested Cube Cluster by creating a cluster of their three existing modeling workstations – those machines had single-core processors with clock speeds ranging from 3.0 to 3.2 GHz and memory ranging from 0.5 GB to just under 1 GB. A typical model run on the agency's most powerful modeling workstation had a run-time of 35 to 40 hours. Testing with the initial cluster yielded run-times that ranged from approximately 35% to 61% of pre-cluster run-time, depending on the processing methodology and number of machines used in the cluster (the best performance used the multi-step methodology with the maximum available number of machines). The Metro Council was pleased with the test results and is planning to invest approximately $30,000 for a full implementation, with most of that expenditure going to purchase two quad-core workstations for an improved cluster. Cube Cluster will be used for all MPO modeling activities. The Metro Council also reported a relatively easy learning curve for Cube Cluster, noting that most of the startup time was for converting their model execution scripts into a format that best utilized the Cluster features (including some legacy FORTRAN programs). They plan to address these issues more directly in the next version of the model by having it fully implemented in Cube Voyager.

METRO (Portland, Oregon MPO)

METRO uses VISUM assignments for major corridor studies, some of them multimodal and at least one currently with tolls under consideration. In addition, many cities and counties within the Portland area also use VISUM for impact studies and system management plans, so there is compatibility between agencies. Prior to implementation, METRO compared the results of the VISUM assignment to those assigned using INRO's EMME/2 forecasting software. The results were sufficiently comparable that METRO moved ahead with their use of VISUM. VISUM's network structure allowed METRO to more precisely define intersection capacities, which has improved their analytical capabilities.

METRO noted that VISUM's path storage is a significant benefit when performing select link analysis; since all the paths are stored as part of the assignment, there is no need to run a new assignment for purposes of the analysis, and there are resulting time-savings for staff. Assignment run-time is about four hours for a regional network with 2013 zones and 25,000 one-way links, with 3-4 vehicle classes and a high degree of convergence. METRO's workstations use a 2.8GHz processor and 4GB of memory. Assignment results have been stable. METRO will continue using VISUM for regional forecasting work, including future analysis of the regional long range plan. They are also planning to move toward regional dynamic traffic assignment (DTA) in the longer term, which can be performed using VISUM. Finally, METRO noted that the next version of VISUM will use real numbers instead of integers for assignment, and this switch will lead to a faster and tighter convergence.

East-West Gateway Coordinating Council (St. Louis, MO MPO)

East-West Gateway uses Cube Cluster for all model runs on a three-machine cluster where each machine has the following specifications: single-core 3.6GHz processor and 2GB memory. They report significant time savings due to Cube Cluster and a minimal learning curve.[24] East-West Gateway staff indicated a willingness to answer further questions.

The Maryland-National Capital Park and Planning Commission, Prince George's Planning Department

TPB staff are already familiar with the Prince George's TranForM model, which is essentially the v2.1D model with a disaggregate zone structure for Prince George's County, a conflated, true-shape regional network, and a few model structure changes, all currently implemented in TransCAD 4.8 and soon to be in production using TransCAD 5.0.[25] However, by running in TransCAD, the Prince George's model takes advantage of the advanced assignment algorithms and advanced computational methods that are native to the software platform. The Prince George's model runs in about 2-3 hours, running 100 iterations with a relative gap of 10^{-2} and two feedback loops. The modeling hardware was recently upgraded to a quad-core workstation. After moving their production model to TransCAD 5.0, Prince George's will be able to use the Caliper multi-threaded UE, path-based, or OUE algorithm for its assignments.

Southern California Association of Governments (SCAG)

SCAG is currently using TransCAD 4.8 for its regional transportation model and is upgrading the model to TransCAD 5.0. There are 4149 internal zones (4191 total zones) and 65,000 links in the SCAG network. Congestion varies widely among the Los Angeles subregions. Peak period average freeway speed is about 30 mph. Due to the size of the model, the OUE feature is not used but will be tested for version 5.0.

Prior to moving to the TransCAD platform, SCAG used TRANPLAN for their year 2000 model validation, which was the basis for the 2004 Regional Transportation Plan (RTP). They utilized five feedback loops with flow smoothing between loops. For each loop, the assignments were done with a maximum of 30 iterations. For their year 2003 model validation, which is the basis for the 2008 RTP, SCAG is using standard user equilibrium assignment in TransCAD. The model is setup for up to 10 feedback loops and a maximum of 40 iterations with a relative gap of 10^{-2}. With a five loop application, it takes about 24 hours to complete their model on a quad-core PC.

Capital District Transportation Committee ([CDTC], Albany, NY MPO)

CTDC uses VISUM for typical MPO modeling applications, including corridor studies, scenario testing for the New York State Department of Transportation (NYSDOT), and testing of projects for the regional long-range plan. Processing time for a typical assignment is about one hour for 12 user equilibrium iterations with use of a feedback loop back to trip generation and 10-20 minutes without feedback. CTDC's model network contains 1,000 zones and 10,000 links.

---

[24] East-West Gateway did not specify current model run-time; however, during testing in a ten-processor cluster, model run-time decreased to eight hours from 48 hours.
[25] See Slavin, *et al* (2006) for more details.

Traffic assignments use VISUM's path-based algorithm. CDTC staff report satisfaction with the assignment results, stability, running time, and convergence.

The CTDC model was recently used for a series of different tests to improve computational time and network convergence when applying feedback.[26] When applied, the most successful methodology converged to 10^{-7} after between 15 and 20 feedback loops – a computational time of between 1.5 and 2 hours. A relative gap of 10^{-6} was reached after only six feedback loops with six user equilibrium iterations per loop. Tests using the most successful methodology with VISUM's "warm start" feature did not show any significant improvement in performance due to applying a previous solution rather than computing an initial solution for travel cost. This result may be in part due to the relatively small and less congested network in Albany (compared to TPB).

Comparison of Alternative Traffic Assignment Methods

Caliper's recent research summarized an empirical comparison of alternative user equilibrium traffic assignment methods on large-scale regional transportation networks.[27] The methods under the comparison were as follows:

- Caliper TransCAD UE using F-W[28]
- Caliper Path-Based
- Caliper Bar-Gera OBA
- Caliper OUE

The origin-based and path-based algorithms were coded and tested based on the existing literature. Modifications were made in the initial stage of implementation to improve the convergence performance. The tests revealed that the Bar-Gera origin-based method converged tightly but only after very long computational times while the path-based method did not converge well on medium to large size networks until modifications were made to the gradient search. The memory requirement and computing times were still issues for the path-based method on larger networks. The origin user equilibrium (OUE) method reached a tight equilibrium in significantly less computing time than F-W. Furthermore, the warm start feature of the OUE method requires much less time to reach a new equilibrium solution for a similar problem where the user previously obtained a good solution and saved those results. This confers significant benefits to practitioners when performing scenario analyses. The research concluded that OUE makes it feasible to calculate traffic assignments with relative gaps of 0.0001 or lower with reasonable computation times for virtually all large models in the U.S.

Caliper's most recent research, presented at the recent TRB Planning Applications conference in Daytona Beach, builds on the above work by comparing only the multithreaded UE F-W and the OUE side-by-side using the Prince George's TransForM model and performing multi-class

[26] See Boyce, *et al* (2007).
[27] Slavin, *et al* (2006).
[28] Caliper uses a proprietary implementation of F-W that reportedly runs faster than comparable algorithms in other modeling programs; the key procedures of this algorithm are now multithreaded to create TransCADs "standard" assignment algorithm.

assignments for different time periods and using feedback.[29] This work again shows significant benefits for both advanced assignment algorithms and advanced computing techniques.

Conclusion / Recommendations

It is important to understand that even with the same term "relative gap" or "origin-based method", the calculation equations and the implementation procedures could be totally different in different software packages, and the proprietary nature of software development makes it difficult to make true "apples to apples" comparisons between platforms. It is the practitioner's responsibility to ensure an adequate approximation of the equilibrium solution is achieved within a reasonable computation time in their model networks. Experience and professional judgment are needed to evaluate whether advanced application procedures for travel forecasting projects actually produce meaningful results. Both the Caliper and PTV research stress the importance of repeating their tests with other models and/or other platforms. Finally, other issues besides the assignment algorithm and computational efficiency may affect convergence and assignment run-time – these include model design, zone structure and size, delay functions, network capacities, and others.[30]

Given TPB's commitment to the Citilabs modeling platform, the next logical step is to pursue whatever run-time and convergence gains can be achieved under the TP+/Cube environment. In order to accomplish this, TPB must convert the v2.2 model to a form that can be used under Cube Cluster. At the time of this memo, TPB has purchased Cube Cluster and has been working with it. Staff has gotten a demo model to run, but have not yet gotten their regional travel model running under Cube Cluster. They must also identify either existing machines for cluster creation, or purchase new hardware. If new hardware is to be acquired, specifications must be created, and an appropriate level of investment for software and hardware upgrades to support future model applications determined. At this time of this memo, TPB has purchased a modeling server (it has two dual-core Xeon processors) to use for both Cube Cluster and general model runs. Any distributed processing work will also be conducted on this machine.

TPB modeling staff should maintain contact with representatives from St. Louis, Minneapolis / St. Paul, and Columbus as a resource throughout this process, and should follow-up immediately with Citilabs to discuss any issues encountered while moving the production model into Cube Cluster. The level of benefits reported by Cube Cluster users should improve the TPB model performance sufficiently while Citilabs implements alternatives to F-W in future versions of Cube. TPB should also consider using the model to test the efficacy of Citilabs' future implementation of any advanced assignment algorithms.

[29] See Slavin, *et al* (2007).
[30] These thoughts were echoed by Dick Walker of Portland METRO in his response to questions.

References

Bertsekas, D. P. and E. M. Gafni. Projected Newton Methods and Optimization of Multicommondity Flows. IEEE Transactions on Automatic Control, Vol. AC-28, pp. 1090-1096. 1983.

Boyce, David, O'Neill, Christopher, and Scherr, Wolfgang. New Computational Results on Solving the Sequential Procedure with Feedback. Paper presented at TRB 2007 Planning Applications Conference, Daytona Beach, FL, May 2007.

Boyce, David; Ralevic-Dekic, Biljana; and Bar-Gera, Hillel, "Convergence of Traffic Assignments: How Much is Enough?" *Journal of Transportation Engineer*, ASCE, January/February 2002.

Cube Cluster case studies from Citilabs website – http://www.citilabs.com, downloaded 2007.

Dafermos, S. C. and F. T. Sparrow. The Traffic Assignment Problem for a General Network. Journal of Research of the National Bureau of Standards, Vol. 73B, pp. 91-118. 1969.

Dial, Robert B. Algorithm B: Accurate Traffic Equilibrium (and How to Bobtail Frank-Wolfe, Volpe National Transportation Systems Center, Cambridge, MA, July 25, 1999.

Dial, Robert B., "A Path-Based User-Equilibrium Traffic Assignment Algorithm that Obviates Path Storage and Enumeration," *Transportation Research Part* B, March 13, 2006.

Email from Dick Walker, Portland METRO

Email from Guoxiong Huang, Southern California Association of Governments

Email from Lubna Shaoib, East-West Gateway Coordinating Council

Email from Mark Filipi, Metropolitan Council for the Twin Cities

Email from Jeremy Raw, VDOT

Email from Rebekah Anderson, ODOT

Email from Howard Slavin, Caliper Corporation

Email from Michael Clarke, Citilabs

Email from Michael Florian, INRO

Email from Fang Yuan, PTV America

Email from Chris O'Neill, CDTC

G. Bellei, G. Gentile and N. Papola. A within-day dynamic traffic assignment model for urban road networks, Transportation Research Part B: Methodological Volume 39, Issue 1, January 2005, Pages 1-29.

Janson, B., and Zozaya-Gorostiza, C. (1987). The Problem of Cyclic Flows in Traffic Assignment. Transportation Research, 21B, pp. 299-310.

Jayakrishnan, R., W. K. Tsai, J. N. Prashker, and S. Rajadhyaksha. A Faster Path-Based Algorithm for Traffic Assignment. In Transportation Research Record 1443, TRB, National Research Council, Washington, pp. 75-83. 1994.

Metropolitan Washington Council of Governments, National Capital Region Transportation Planning Board. Results of FY 2006 Travel Forecasting Research. Produced under contract by Vanasse Hangen Brustlin, Inc., 2006.

Slavin, Howard; Brandon, Jonathan and Andres Rabinowicz. "An Empirical Test of Alternative Traffic Assignment Methods," TRB 2006 Annual Meeting, Washington, D.C. January 2006.

Slavin, Howard; Brandon, Jonathan, Rabinowicz, Andres, and Srinivasan Sundaram. "Achieving Planning Model Convergence," TRB 2007 Planning Applications Conference, Daytona Beach, FL, May 2007.

Spielberg, F., and P. Shapiro._Metropolitan Planning Organization Travel Forecasting State of the Practice. Presentation to Transportation Research Board Annual Meeting, Washington, DC, 2007.

P. Wolfe. Convergence Theory in Nonlinear Programming. In Integer and Nonlinear Programming, edited by J. Abadie, pp. 1-36, North-Holland, Amsterdam. 1970.