Title
HPC4Mobilty w/ UCB

Authors
Nugent, P
MACFARLANE, J

Publication Date
2020-05-06

Peer reviewed

I High Performance Computing and Big Data

I.1 High Performance Computing for Mobility (HPC4Mobility)

Jane Macfarlane, Lead Principal Investigator
Lawrence Berkeley National Laboratory
1 Cyclotron Road
Berkeley, CA 94780
E-mail: jfmacfarlane@lbl.gov

Bin Wang, LBL, Principal Investigator
E-mail: wangbin@lbl.gov

Prasad Gupte, DOE Technology Development Manager
U.S. Department of Energy
E-mail: prasad.gupte@ee.doe.gov

Start Date: October 1, 2017 End Date: September 30, 2019
Project Funding (FY18): $250K DOE share: $250K Non-DOE share: $0

Project Introduction

The purpose of this project is to examine the energy impact of urban-scale traffic for the Los Angeles Basin by developing and implementing a scalable traffic assignment model. An energy optimization function will be posed and when integrated into the optimization code for travel assignment it can be mathematically proven to converge. The energy optimization function can then be compared to the typical travel time optimization that is traditionally used in traffic assignment models.

The analysis will begin with static traffic assignment models with the routing for all origin and destinations computed in parallel on high performance computing facilities. Convergence of the numerical methods rely on the solution of convex programs (or extensions of these). This step will mostly consist of demonstrating the ability to parallelize the Frank Wolfe algorithm on various platforms.

This work will contribute to LBNL’s efforts to develop new processes, analytical tools, program designs, and business models to advance the state of the art in next-generation sustainable transportation solutions.

Objectives

Starting with a small and well researched part of the LA Basin that is known as the Connected Corridor, demand data (provided by SCAG and in collaboration with LA Metro) will be used as an input to the Traffic Assignment code. Robust code that has been developed by ITS will be used to implement the Frank Wolfe algorithm. This code will be further developed to include the energy optimization case. The proposed energy model is a combination of the CMEM emission model with a traditional BPR function. To ensure convergence, the speed – fuel consumption curve will be slightly adjusted.

A geospatial area that represents 28,000 road links will be investigated. A demand model of 100,000 Origin/Destination pairs will be applied to this road network.

Four key cases will be investigated:

- Energy optimized at the system level
- Energy optimized selfishly at the vehicle level
- Travel time optimized at the system level
- Travel time optimized selfishly at the vehicle level.

Once complete, we will to move to larger scale road networks, dynamic assignment models, and information aware routing.

Approach

Mobiliti, a proof-of-concept simulator enables distributed memory parallelism with an asynchronous execution strategy. Distributed memory parallelism allows for the use of multiple nodes in a high-performance computer or cluster, which equates to the ability to apply more processors and more memory to the problem. Asynchronous execution strategy removes the serial bottleneck of thread synchronization and increases the progress of the simulation. However, it requires an additional causality control in order to rollback events that are out of sync. This approach is the subject of current parallel distributed simulation research e.g. Time Warp protocol. Mobiliti can model networks of the scale of 9.5M trip legs over a road network with 1.1M nodes and 2.2M links (eg. the size of the SF Bay Area), processing 2.4B events in less than 30 seconds using 1,024 cores on a supercomputer.

Results

In order to handle even larger simulations, we have updated our vehicle routing algorithm to use contraction hierarchies instead of Dijkstra’s algorithm or the A* algorithm. As a result of this update, our routing is much faster, and it is capable of identifying optimal routes through the network instead of approximate ones. This work will enable us to study the impact of dynamic re-routing since we can now support more efficient routing during the simulation, rather than simply as a preprocessing step.

After a coordination discussion with local transportation officials from San Francisco and San Jose, we have coordinated with the SFCTA to ingest their validated demand model consisting of 27.6M trip legs, representing the traffic demand from 7.3M residents. Of these original trip legs, 21.7M result in vehicle trips through the road network, while the remaining are satisfied through walking or cycling. Since the SFCTA model works with demand at a granularity of TAZ to TAZ, we adapted the model to generate node-level inputs for Mobiliti via random node selection within the origin and destination TAZs. Since demand may not be uniformly distributed among each TAZ, we are identifying ways to modify the resulting demand to further increase the fidelity of our simulation. We have additionally optimized the parallel decomposition network partitioning strategy for the SFCTA demand model. Excluding initialization time, we can simulate 21M trip legs during one simulated day in less than 30 seconds when utilizing 32 nodes (1,024 compute cores) of the Cori supercomputer.

Extended TAP:

Instead of only optimizing the system based on the travel time, we extended the existing algorithms for the traffic assignment program (TAP) with new objective functions to incorporate the vehicle fuel consumption. Specifically, we extract vehicle fuel consumption curve, i.e. fuel consumption rate vs. speed curve from the developed afore-mentioned data-driven energy models. We conduct multiple experiments to investigate the patterns of different traffic assignment methods, i.e. time-based user-equilibrium (UET), time-based system-optimal (SOT), fuel-based user-equilibrium (UEF) and fuel-based system-optimal (SOF). Preliminary visualization programs are developed to perform exploratory analysis on these four different cases. In addition to the reduced network used by the original TAP code, we also processed the bay area network of a much larger size into the TAP, using the Mobiliti framework.

Conclusions

The new network has 2 million road links and the new traffic demand includes 22 million origin-destination pairs. The preliminary performance results, indicating the total solving time of 45 mins, were collected on the Cori supercomputer with single computing node and 64 threads.

Key Publications

References

None.

Acknowledgements [Use EERE Heading 3]

Insert Acknowledgements Text
[Use Body Text]

{Opportunity to provide recognition for individuals (NETL manager, co-authors, etc.) that aren’t a PI or the DOE Technology Manager. NOT REQUIRED. DELETE header if not used.}