

HOMES: Highway Operation Monitoring and Evaluation System (Demo Paper)

Chang-Tien Lu, Arnold P. Boedihardjo, Jing Dai, Feng Chen
 Department of Computer Science, Virginia Tech, USA
 {ctlu,aboediha,daij,chenf}@vt.edu

1. INTRODUCTION

Increasing population, rising fuel costs, and stricter environmental policies have all contributed to the greater demand for improved roadway capacity and efficiency. A 2007 national study has found that inefficiencies within the roadway system commonly associated to traffic congestions have costed the U.S. \$78 billion in productivity loss and 2.9 billion gallons of wasted fuel [2]. A key component in improving the efficiency of the roadway network is the availability of an advanced traffic information system to provide real-time monitoring and long-term evaluation of the roadway conditions. To that endeavor, we have developed the *Highway Operation Monitoring and Evaluation System (HOMES)*, a traffic visualization system that allows for browsing the spatial-temporal dimension hierarchy via integrated data cube operations. The Virginia Department of Transportation (VDOT) currently utilizes HOMES to monitor traffic incidents, analyze roadway behaviors, generate operation strategies, and develop/verify highway designs. In general, HOMES can identify traffic patterns, rules, and anomalies to achieve the following benefits: efficient roadway designs; objective evaluation metrics of traffic policies; improved management of operations and emergent events; and better utilization of the roadway network.

HOMES currently monitors highways I-66 and I-395 within the Washington Metropolitan Area. It collects streaming information (arrival rate = 1 observation/sensor/min) from approximately 850 radar sensors and loop detectors. An important challenge in the development of HOMES is the provision of real-time query processing for both current and historical data under various cube operations. To address this issue, HOMES employs a spatial data warehouse approach as the underlying traffic data management strategy [3]. Fast cube-based query response times (less than 5 sec. average) are achieved by maintaining concurrent sets of aggregated and non-aggregated sensor information. Quick traffic data updates are accomplished by an incremental approach for computing the updated aggregate and non-aggregate representations of the traffic data streams.

To further enhance the query response times, we employed a caching strategy to reduce the operation costs associated with database communication and query processing. Caching allows for the reduction of resultant data in network transit and the number of direct and redundant accesses to the database. Because accessing data in main memory (i.e., data with higher locality) is more efficient (i.e., lower response times) than accessing a remotely located database, some of the traffic data will be temporarily stored in the system's application memory pool for future recurrent accesses. A critical consideration for the caching strategy is the use of an appropriate replacement policy that adapts well to the system's data request behavior. The Least Recently Used (LRU) has been adapted as the replacement policy in the HOMES caching scheme. As a result, HOMES allows traffic personnel to monitor and assess highway conditions in real-time while supporting efficient data operations for long-term traffic analysis.

2. VISUALIZATION OVERVIEW

A traffic measure (e.g., speed) is observed at a particular time and space. The following describes the supported types of traffic measure and its associated time/space aggregations:

Traffic measure – The core traffic measures provided by the detectors are speed (avg. vehicle velocity), volume (# of vehicles), and occupancy (ratio of time vehicle is detected). The derived measures are travel time, travel time index (ratio of current to historical travel time), flow (# of vehicles detected per hour), vehicle miles traveled (total miles traveled by all vehicles), vehicle hours traveled (total hours accumulated by all vehicles), and delay (difference of current travel time and ideal travel time).

Time – Time is aggregated into time of day (time point or time range), T_{TD} , and day of week (Monday to Sunday), T_{DW} .

Space – Spatial entities are hierarchically categorized into detectors, S_D , station (a set of spatially continuous detectors), S_{ST} , link (set of spatially continuous stations), S_L , and section (set of spatially continuous links), S_{SE} .

For each traffic measure, associated visualizations are given from several elements of the pair-wise combination set of T and S aggregation hierarchy. For example, a contour plot ($T_{TD}S_L$) can be shown for all sets of available links in a highway for a given time range. Hence, roll-up and drill-down operations are supported by navigating the various hierarchies of the time and space dimensions. Figure 1 shows the supported visualization dimension hierarchy of HOMES.




Figure 1. HOMES visualization dimension hierarchy. Each node represents a supported visualization.

Each class of traffic application users will use a different set of visualizations for their particular tasks. For instance, a traffic operations manager may invoke a S_L geographic map visualization to constantly monitor the current condition of the roadway system. A traffic planner, on the other hand, may not be interested in the current conditions, but rather the weekly patterns (T_{DW}) in order to design new optimization strategies. HOMES provides a comprehensive suite of visualization to meet the specific requirements of various classes of traffic users.

3. CASE STUDY – INCIDENT MONITORING AND ANALYSIS

In this section, we demonstrate the use of HOMES visualization system for detecting, monitoring, and evaluating a traffic incident that occurred at I-66 westbound on February 5, 2008 10:50 AM. The incident was an overturned vehicle that resulted in a blockage of the left lane.



Figure 2. S_L map of I-66 link-level travel time index (TTAi) for Feb 5, 2008 at 10:50 AM. The black box shows the incident region.



Figure 3. S_D map of detector speeds (MPH) at 10:50 AM on Feb 5, 2008. Black box shows traffic incident.

Incident detection: Figure 2 provides the S_L map snapshot of I-66 on Feb 5, 2008 at 10:50 AM. The area inside the black rectangle is the region of traffic incident. Under normal circumstance, the link within the rectangle would exhibit a travel time index of 1. However, the exhibited link travel time is about

50% more than usual ($TTA_i > 1.5$) which would indicate a potential incident to the traffic operators. Figure 3 gives the drill-down detector view (S_D) of the incident link. Certain incidents can exhibit localized congestion within a particular lane, however, the dramatic speed drops of all the surrounding lanes imply that the observed situation is more severe. In short, these multiple views allow traffic operators to rapidly detect and provide quick evaluations of a traffic incident.

Monitoring incident recovery: As the incident progresses, traffic operators can monitor the real-time status of the affected roadway segment. This ability allows the operators to dynamically assess the effectiveness of its recovery effort and invoke necessary protocols to minimize the total recovery time. Figure 4 provides the T_{TD} speed and travel time plots of the incident link compared to its average value of the past two weeks. At around 11:50 AM, the speed begins to rise indicating the first measurable effects of the incident recovery process. Traffic models employing flow theory can be applied to this data sample to extrapolate the amount of time for which traffic will resume to normal speed [1]. If the observed recovery starting point occurred at a much later time, then the likelihood of the congestion coinciding with the evening rush hour is increased. At that point the traffic operator may invoke a protocol calling for more on-site emergency personnel. An additional benefit provided by the T_{TD} visualization is that it allows the users to precisely determine the incident time. Although the incident was reported to have occurred at 10:50 AM, the graph indicates that it actually occurred 20 minutes earlier at 10:30 AM.




Figure 4. T_{TD} speed plot (left) and T_{TD} travel time (in minutes) plot (right) of incident link from 10:20 AM to 12:00 PM. Red line is the current traffic and green line is the avg. traffic for the past two weeks.

Another important task is to monitor the effects of the incident to its neighboring links and lanes. Figure 5 gives the $T_{TD}S_L$ and $T_{TD}S_D$ plots of the adjacent links and detectors. This type of information gives a broader picture of the incident's impact and provides traffic operators with concrete actionable information to minimize the incident recovery time (e.g., activating ramp meters to reduce inbound traffic flow).




Figure 5. $T_{TD}S_L$ speed plot (left) of incident link (162) and its adjacent neighbors. $T_{TD}S_D$ speed plot (right) incident lane and its neighbors.

Post-incident analysis: One critical question that is posed by traffic engineers and planners is, "What was the total impact and cost of the incident?" The cost is commonly derived from the total delay observed which in turn is translated to a dollar measure of productivity loss. Figure 4 provides a part of that answer by displaying the total delay incurred at a small and localized link region. To analyze the impact of the incident for the total highway section, HOMES provides $T_{TD}S_{ST}$ contour plots (Figure 6) which gives the overall impact to vehicular speed for the entire span of I-

66 westbound. The x-axis gives the time, y-axis represents the milepost location, and color indicates the speed values. Compared to a non-incident day (previous day), the impact of the traffic can be seen to affect approximately 10 miles of I-66 (from mileposts 60 to 70) for the period of 2.5 hours. The exact speed values can be obtained by invoking the HOMES data export function. With this data, the cost of the productivity loss can be more accurately computed.




Figure 6. $T_{TD}S_{ST}$ speed contour plots for the entire stretch of I-66 between 12:00 AM 2:00 PM. Black box indicates incident region.

Performing a drill down operation on the contour plot by taking the vertical cross-sections of all highways at times 11:20 AM and 12:00 PM produces the plots in Figure 7. In this figure, one can observe with greater resolution the effects of the incident on the spatial dimension at specific snapshots of time.




Figure 7. S_L speed plots of I-66 at 11:20 AM (left) and 12:00 PM (right). Black box indicates affected region of incident.

4. SUMMARY

Visualization has shown to be an important tool for identifying patterns, trends, and anomalies in massive spatial data sets. HOMES is an attempt to develop visualization techniques for monitoring and analyzing traffic data. HOMES enables a comprehensive analysis of highway conditions and patterns by implementing high-performance cube operations coupled with an extensive set of visualizations. As exemplified by the case study, HOMES can facilitate VDOT by reducing their response time to an emergent event, providing real-time feedback of their recovery efforts, and quantifying their overall performance. Furthermore, HOMES' visualization components support short/long-term analysis to improve traffic operation strategies, verify/develop highway designs, and enhance utilization of the roadway network. The provided visualization components alleviate much of the time-consuming and manual tasks of knowledge discovery, and thus promoting the effective and efficient use of the transportation data.

5. REFERENCES

- [1] N. Bellomo, V. Coscia, and M. Delitala, "On the Mathematical Theory of Vehicular Traffic Flow I: Fluid Dynamic and Kinetic Modelling," *Mathematical Models and Methods in Applied Sciences*, vol. 12, pp. 1801-1843, 2002.
- [2] D. Schrank and T. Lomax, "The 2007 Urban Mobility Report [<http://mobility.tamu.edu/ums/report/>]," Texas Transportation Institute, 2007.
- [3] S. Shekhar, C. T. Lu, X. Tan, and S. Chawla, "Map Cube: A visualization Tool for Spatial Data Warehouses," in *Geographic Data Mining and Knowledge Discovery*, H. J. Miller and J. Han, Eds.: Taylor and Francis, 2001.