Proceedings of the 2009 INFORMS Simulation Society Research Workshop
L.H. Lee, M.E. Kuhl, J.W. Fowler, and S. Robinson, eds.

EMBEDDED AD HOC DISTRIBUTED SIMULATION FOR TRANSPORTATION SYSTEM MONITORING AND CONTROL

Michael Hunter
School of Civil and Environmental Engineering
Georgia Institute of Technology
Atlanta, GA 30332, USA

Jason Sirichoke
Hewlett-Packard Company
Colorado Springs
Colorado 80919, USA

Richard Fujimoto
Ya-Lin Huang
Computational Science and Engineering
Georgia Institute of Technology
Atlanta, GA 30332, USA

1 INTRODUCTION

Sensors, mobile computing, and wireless communication networks have witnessed dramatic advances over the past few years. Simultaneously, on-line simulations, fed by real time data, have been incorporated into diverse areas such as manufacturing (Lendermann et al. 2005), supply chain optimization (Turner, Cai, and Gan 2000), and transportation system operations, to mention a few. The increasing presence of mobile and ubiquitous computing combined with nearly universal communications is resulting in exciting new opportunities in transportation system management (Werner 2004a, Werner 2004b). One can envision a distributed, adaptive, self-optimizing transportation infrastructure that can automatically reconfigure itself to maximize efficiency and minimize the effects of unexpected events ranging from everyday crashes to catastrophic natural or human generated disasters.

In this current effort we envision embedding on-line distributed simulations operating within transportation network elements (participating vehicles, roadside cabinets, etc.) in an attempt to create a transportation management and prediction system distributed across in-vehicle computer systems, roadside computers, and traffic management centers. In this approach participating vehicles play an active role in the monitoring and prediction of future states of the transportation network. Throughout a vehicle's trip it obtains information from sensor networks and other vehicles, performs real-time simulations and predictive functions, and shares its projections with other vehicle's simulators to create an aggregate view of the individual vehicle's area of interest and the transportation system as a whole. This distributed simulation approach represents a shift in transportation systems management, moving to the vehicle much of the intelligence, computational power, and responsibility for network management. Our current focus is set on investigating whether an embedded real-time distributed system is capable of efficient and accurate system monitoring and producing faster than real time future state forecasts.

2 AD HOC DISTRIBUTED SIMULATIONS

Our current research efforts are aimed at developing an embedded, distributed, simulation based transportation management system, combining in-vehicle simulators with information servers and simulations running within the roadside infrastructure. Future efforts hope to remove the need for the fixed infrastructure component, resulting in an entirely vehicle based management system. We envision a collection of in-vehicle simulations, each publishing projections of future system states, and utilizing projected state information from other simulators, real-time embedded traffic sensor data, and historical traffic behavior patterns. In our current implementation each participating vehicle contains a simulator that models the roadway network in the immediate vicinity of the vehicle, illustrated in Figure 1. Thus, as the vehicle traverses the network it will be simulating a dynamic roadway topology that must be continuously updated to reflect the current vehicle position. An advantage of this approach is that an embedded distributed simulation operates in close proximity to real-time data, allowing predictions to be based on detailed, up-to-date data collected from nearby sensors. In additional, it is anticipated that multiple vehicles will be simulating overlapping areas, resulting in significant redundancy, offering greater robustness and resilience to failures.

A key aspect of this approach is that the distributed simulation models operate in an asynchronous fashion, that is, a vehicle's on-board simulator is not required to operate in synchronous lock step with other participating vehicles, allowing for largely autonomous operation. Data and results sharing between the simulations are implemented by an approach inspired by the Time Warp algorithm (Jefferson 1985). We refer to this collection of embedded, overlapping, in-vehicle and roadside based autonomous, on-line simulations as an *ad hoc distributed simulation* (Fujimoto et al. 2007, Hunter et al. 2009). It is noted that while sharing similarities with a conventional distributed simulation, an ad hoc distributed simulation possess several distinguishing characteristics. While a conventional distributed simulation is partitioned into non-overlapping areas the vehicle simulators within the ad hoc network may cover overlapping areas, as seen in Figure 1. For instance, two or more simulators may both cover one portion of the network while other sections of the network may not be covered at all. Further, the network area modeled by a vehicle simulator can vary over time, as the vehicle traverses the network. Finally, the set of participating simulators is dynamic as new vehicles can join and existing vehicles leave during the analysis period.

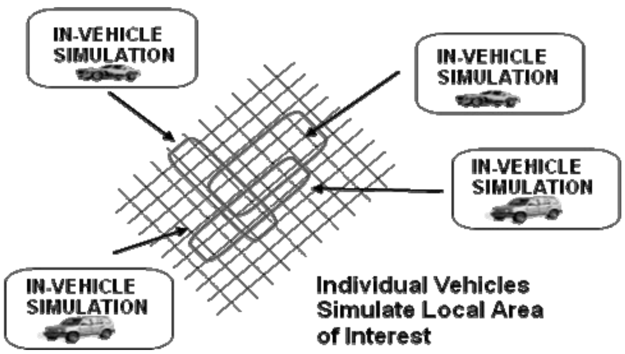

Figure 1: In-Vehicle Simulation

To ultimately implement an embedded ad hoc distributed system several key research problems must be addressed:

- *Individual Vehicle Simulator Self-Configuration and Area Identification.* It is necessary that as each vehicle enters the system that it senses its geographic area and automatically configures itself to simulate the appropriate region. One central question is the size and location of the traffic network that should be modeled by each participating vehicle.
- *Dynamic Simulation Topology*. As a vehicle moves, new portions of the network may have to be incorporated in the on-board model, while models of other portions may no longer be needed. The capability to create dynamic, adaptive simulation models that automatically change the modeled area based on vehicle movement and incorporating new data into the simulation is needed.
- *Simulation Self-calibration*. The simulations must have a means to automatically correct themselves based on live, measured data and results for other simulations or data sources within the ad hoc network.
- *System Reliability and Resilience*. The ad hoc simulation approach should offer significance resilience to failures due to the inherent redundancy of the approach. However, the effects of simulator failures and the potential to create cascading failures, either due to processor or communication failures require investigation.
- *Bandwidth Reduction Techniques*. Vehicle networks supporting ad hoc distributed simulations will be constrained by limited bandwidth in the foreseeable future. Techniques to reduce the bandwidth requirements of ad hoc distributed simulations are needed.
- *Aggregation of Ad Hoc Distributed Simulations*. Ad hoc simulation performance will depend on factors such as the fraction of vehicles that are instrumented (penetration ratio), vehicle density, simulated area, and accuracy of measurement data, among other factors. A thorough evaluation is required to evaluate the effectiveness of this approach in realistic traffic conditions, settings, and scenarios.

Our initial approach to implement an ad hoc distributed simulation system is a client/server architecture, where global state objects (i.e. simulation results from participating vehicles and sensor data) and associated composition functions (i.e. methods to aggregate data from the distributed simulations) are implemented at the server level (Fujimoto et al. 2007, Hunter et al. 2009). The server receives and stores state updates, compute composite values, and disseminates this information to other simulators. The key elements of our initial approach include Space Time Memory, State Aggregation, and Rollback Based Synchronization. Space-time Memory (STM) is used to coordinate interactions among simulators, storing current state information (e.g., traffic flow rates on various links) and projections of future states provided by the embedded vehicle simulators. The STM maintains multiple, time-stamped versions of state variables (Fujimoto 1989). State Aggregation is used to combine values provided by the multiple vehicle simulators to form a composite value of the system state. Finally, the pro-posed distributed simulation mechanism relies on an optimistic Rollback Based Synchronization for self-calibration when predictions are invalidated by sensor data or unexpected events, e.g., crashes resulting in new congestion points. When a simulator is rolled back, a mechanism similar to Time Warp (Jefferson 1985) anti-messages is used to retract incorrect messages sent by vehicle based simulators to the STM based on rolled back computations. This allows revised predictions to propagate rapidly across the network as new information becomes available. A key question to the use of a roll back mechanism is the development of thresholds to be utilized to identify when predictions from individual vehicle simulator are sufficiently in conflict with other data that they should be rolled back.

3 INITIAL EXPERIMENTS

Our preliminary work assumed that the vehicle simulators were already configured to model the designated scenario, and focused on predicting the effects of changes in traffic patterns. Initial experiments were conducted on a notional west-to-east ten intersection arterial model of a single roadway using a cellular automata simulator (Fujimoto et al. 2007, Hunter et al. 2009). Ten vehicles were assigned to simulate the west half of the corridor (i.e. the 5 west most intersections) and ten vehicles were assigned simulate the east half of the corridor (i.e. the 5 east most intersections). The links at the midpoint of the corridor were modeled by all twenty vehicles. Four experiments were undertaken, three with constant arrival rates on the corridor (testing the impact of differing demand levels) and one with an initial arrival rate that is then significantly increased during runtime. This last experiment was intended to model a sudden influx of traffic at one end of the network to evaluate the responsiveness of the system to sudden changes. Initial results from these experiments demonstrated that the ad hoc distributed simulation shows good potential to predict the behavior of transportation systems both under steady state conditions and in the face of a rapidly changing traffic load. Results for the increase in traffic demand showed that the ad-hoc approach is able to capture the increase in traffic and pass this information to the downstream vehicle simulators. That is, while the vehicles simulating the downstream intersections did not receive any field data to indicate the likely future increase in traffic flow they are able to capture the impact of the flow increase on future performance with updated predicted demands from the upstream vehicle simulators.

These initial experiments were followed by a second round utilizing a larger network. Specifically, a Manhattan-style 10 x 10 grid with two-lane, two-way roads. The grid contains 100 equally spaced intersections (1320 ft. apart) with each intersection containing a shared through/right turn lane and a 600 ft left turn bay. All intersections operate under a four-phase fixed-time 120 sec cycle (i.e. 20 sec east/west left turn, 50 sec east/west through, 15 sec north/south left turn, and a 35 sec north/south through). Turn rates are constant at each intersection, with 85% of the vehicles on each approach traveling through the intersection, 5% turning left and the remaining 10% turning right. All vehicles are assumed to have a maximum velocity of 45 ft/sec. Four main arterials are defined that have a significantly higher volume of traffic than other roads in the network. These arterials occur at the horizontal and vertical midpoints of the network, cross the entire length of the network.

Forty participating vehicles (i.e. vehicles with on-board simulators) are distributed over the network at initialization, each simulating a group of intersections and connected roads, with the arrival rate on all boundary input roads initialized to a specified value. Each client simulation area depends upon the vehicle location and the direction in which the vehicle is traveling. The simulation region for each client is defined as the 5 x 3 intersection area in front of the vehicle assigned to it. If the network area in front of the client is smaller than the region, then the client only simulates the smaller extent.

Experiments were conducted to model the network under steady state conditions (see Figure 2). Three different volume scenarios were tested: low, medium, and high. Each experiment was conducted for two sets of thresholds, the first equivalent to the earlier initial experiments and a second set of tighter thresholds, allowing for an exploration of the balance between potential improved predication capabilities and higher computational demands due to additional roll backs. Ten replicate simulations of the entire network were also generated to compare to the ad hoc distributed simulation performance. As expected, under low volume conditions, the projections from the ad hoc version correlated well with the replicated average over the course of the simulation. These simulations required no synchronization rollbacks, indicating the individual vehicles provided reasonable predictions without receiving additional prediction data from other participating vehicles. At higher traffic volume input rates, rollbacks do occur across the different sets of thresholds. In the case of the medium input rate simulations, even with an increase in rollbacks when using a tighter threshold, the results using the two sets of thresholds closely resemble each other over the course of the simulation. In some cases the calculated arrival rates exactly matched over significant periods of time in the simulation, despite using different threshold sets.

The next set of experiments tested a case where the input arrival rate changes suddenly for the same traffic network used for the steady state experiments. The purpose of this experiment was to understand how well the distributed simulation approach captures the effect of a change in arrival rates. In this experiment, the network was initially set to match the original low input rate state. At simulation time 1200, the arrival rate of all border input roads on the west side of the traffic network were increased to 300 cars/hr on the minor arterials and 500 cars/hr for the main arterial. Unlike the steady state simulations described in the previous section which experienced no rollbacks at the low input rate settings, a distributed simulation will require rollbacks in order to successfully simulate the increase in traffic.

The graph in Figure 3 focuses on a typical road segment along the east bound main artery. The set of experiments at each in-put rate is represented in the figure by three lines showing the average input rate over time of the replicated trials, the ad hoc simulation, and a single ad hoc distributed client (which illustrates how individual clients are reacting to the change in the arrival rate). The ad hoc simulation approach is able to track the unexpected change in input rate. Measurements on a variety of points in the network demonstrated that the ad hoc approach was able to track this change in traffic flow at different locations throughout the network.

4 SUMMARY

Transportation is the largest industry in the world. Our transportation system significantly impacts every individual and the welfare of our entire nation in terms of economics, health, and quality of life to name a few. However, for many decades improvements to our ability to actively manage our surface transportation system in real-time have been stagnant. Today, wide-spread deployment of sensors, computers, and communications in vehicles and roadways is creating new challenges and opportunities to effectively exploit the wealth of real-time data and information that

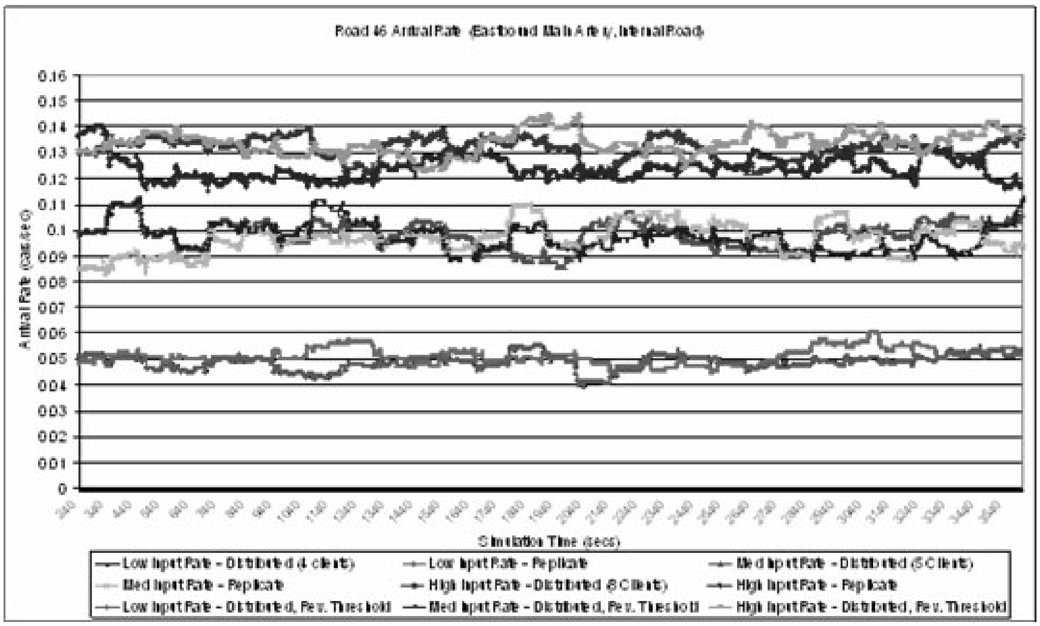

Figure 2: Steady state experiments, arrival rates on internal road in eastbound main artery

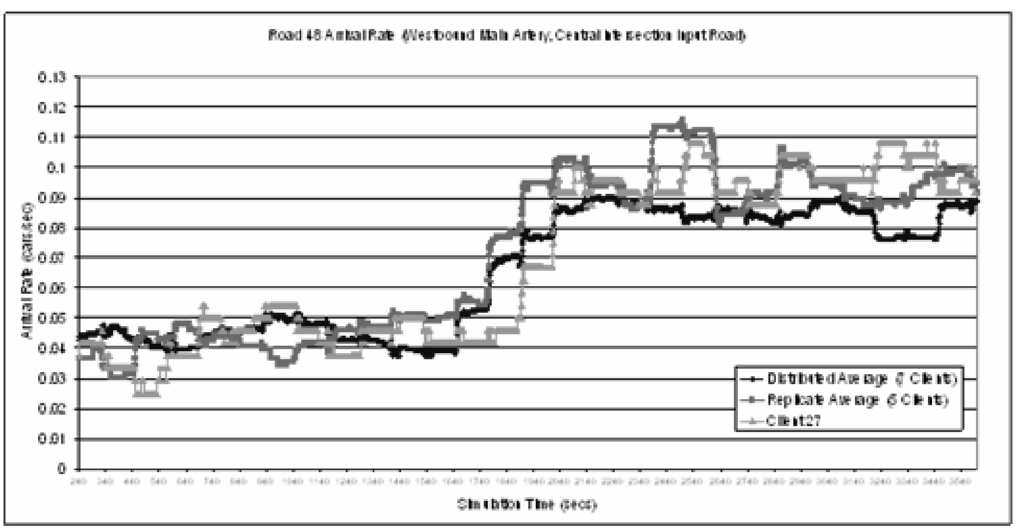

Figure 3: Input rate modification experiments, arrival rates on central intersection input road in eastbound main artery

are becoming available. We attempt to capitalize on these rapid technology and communications advancements using ad hoc distributed simulations, that features dynamic collections of autonomous simulations interacting with each other and with real-time data in a continuously running, real-time, distributed simulation environment. Initial results support the potential of this approach.

ACKNOWLEDGMENTS

This research was supported under NSF grants CNS-0540160 and EFRI-0735991.

REFERENCES

Lendermann, P., M. Y. H. Low, B. P. Gan, N. Julka, L. P. Chan, L. H. Lee, S. J. E. Taylor, S. J. Turner, W. Cai, X. Wang, T. Hung, L. F. McGinnis, and S. Buckley. 2005. An Integrated and Adaptive Decision-Support Framework for High-Tech Manufacturing and Service Networks. In *Proceedings of the 2005 Winter Simulation Conference*, ed. M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines, 2052–2062. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Turner, S. J., W. Cai, and B. P. Gan. 2000. Adapting a Supply-Chain Simulation for HLA. In *Proceedings of the 4th IEEE Workshop on Distributed Simulation and Real-Time Applications*, 71–78. Washington DC: Institute of Electrical and Electronics Engineers Computer Society.

Werner, J. 2004a. More Details Emerge about the VII Effort. *Newsletter of the ITS Cooperative Deployment Network*. Available via <http://www.ntoctalks.com/icdn/vii_details_itsa04.html> [accessed May 17, 2009].

Werner, J. 2004b. USDOT Outlines the New VII Initiative at the 2004 TRB Annual Meeting.

Jefferson, D. 1985. Virtual Time. *ACM Transactions on Programming Languages and Systems* 7:404–425.

Fujimoto, R., M. Hunter, J. Sirichoke, M. Palekar, H. Kim, and W. Suh. 2007. Ad Hoc Distributed Simulations. In *Proceedings of the 21th International Workshop on Principles of Advanced and Distributed Simulation*, 15–24. Washington DC: Institute of Electrical and Electronics Engineers Computer Society.

Hunter, M., H. Y. Kim, W. Suh, R. Fujimoto, J. Sirichoke, and M. Palekar. 2009. Ad Hoc Distributed Dynamic Data-Driven Simulations of Surface Transportation Systems. *Simulation* 85: 243–255.

Fujimoto, R. M. 1989. The Virtual Time Machine. In *Proceedings of the First Annual ACM Symposium on Parallel Algorithms and Architectures*, 199–208. New York, NY: Association for Computing Machinery.

AUTHOR BIOGRAPHIES

MICHAEL HUNTER is an Assistant Professor at the School of Civil and Environmental Engineering at Georgia Institute of Technology. His primary teaching and research interests are in transportation operations and design, specializing in adaptive signal control, traffic simulation, and real-time arterial corridor operations. Dr. Hunter obtained his Ph.D. and M.S. in Civil Engineering from the University of Texas at Austin in 2003 and 1994. He obtained his B.S. in Civil Engineering from Rensselaer Polytechnic University in 1992. He has also worked in private industry conducting traffic impact studies, signal timing projects, freeway operation evaluations, toll plaza analyzes, etc. His email is <michael.hunter@ce.gatech.edu>.

JASON SIRICHOKE is currently working as a systems/software engineer at Hewlett Packard in the StorageWorks Division. He received both his B.S. and M.S degrees in Computer Science from the Georgia Institute of Technology in 2006 and 2007. His current work at Hewlett Packard is focused on the performance of enterprise-level storage virtualization platforms.

RICHARD FUJIMOTO is Regents' Professor and Chair of the Computational Science and Engineering Division of the College of Computing at the Georgia Institute of Technology. He received the Ph.D. and M.S. degrees from the University of California (Berkeley) in 1980 and 1983 (Computer Science and Electrical Engineering) and B.S. degrees from the University of Illinois (Urbana) in 1977 and 1978 (Computer Science and Computer Engineering). His research is concerned with the execution of discrete-event simulation programs on parallel and distributed computing platforms. This research has included work on platforms ranging from mobile distributed computing systems to cluster computers to supercomputers. This work has included several application areas including transportation systems, telecommunication networks, multiprocessor systems, and defense systems. He lead the working group that was responsible for defining the time management services for the Department of Defense High Level Architecture (HLA) effort (IEEE Standard 1516). His email is <fujimoto@cc.gatech.edu>.

YA-LIN HUANG is a Ph.D. student in the College of Computing at Georgia Institute of Technology. She obtained her M.S. degree in Network Engineering and B.S. degree in Computer Science and Information Engineering from National Chiao Tung University in 2007 and 2005. Her current research focuses on distributed simulations. Her email is <huangyl@gatech.edu>.