

Ad Hoc Distributed Simulation of Surface Transportation Systems

R.M. Fujimoto, R. Guensler, M. Hunter, K. Schwan, H.-K. Kim,
B. Seshasayee, J. Sirichoke, and W. Suh

Georgia Institute of Technology, Atlanta, GA 30332 USA
{fujimoto@cc, randall.guensler@ce, michael.hunter@ce,
schwan@cc}.gatech.edu

Abstract. Current research in applying the Dynamic Data Driven Application Systems (DDDAS) concept to monitor and manage surface transportation systems in day-to-day and emergency scenarios is described. This work is focused in four, tightly coupled areas. First, a novel approach to predicting future system states termed ad hoc distributed simulations has been developed and is under investigation. Second, on-line simulation models that can incorporate real-time data and perform rollback operations for optimistic ad hoc distributed simulations are being developed and configured with data corresponding to the Atlanta metropolitan area. Third, research in the analysis of real-time data is being used to define approaches for transportation system data collection that can drive distributed on-line simulations. Finally, research in data dissemination approaches is examining effective means to distribute information in mobile distributed systems to support the ad hoc distributed simulation concept.

Keywords: surface transportation systems, ad hoc distributed simulations, rollback operations.

1 Introduction

The Vehicle-Infrastructure Integration (VII) initiative by government agencies and private companies is deploying a variety of roadside and mobile sensing platforms capable of collecting and transmitting transportation data [1-3]. With the ongoing deployment of vehicle and roadside sensor networks, transportation planners and engineers have the opportunity to explore new approaches to managing surface transportation systems, offering the potential to allow the creation of more robust, efficient transportation infrastructures than was possible previously. Effective and efficient system management will require real-time determinations as to which data should be monitored, and at what resolutions. Distributed simulations offer the ability to predict future system states for use in optimizing system behaviors both in day-to-day traffic conditions as well as in times of emergency, e.g., under evacuation scenarios. Data collection, data processing, data analysis, and simulations performed by system agents (sub-network monitoring systems, base stations, vehicles, etc.) will lessen communication bandwidth requirements and harness surplus computing

capacity. Middleware to manage the distributed network, synchronize data and results among autonomous agents, and resolve simulation output conflicts between agents using disparate data sets become critical activities in such a system. Dynamic, data-driven application systems (DDDAS) offer the potential to yield improved efficiencies in the system that can reduce traffic delays and congestion, pollution, and ultimately, save lives during times of crisis.

We are addressing this challenge through a distributed computing and simulation approach that exploits in-vehicle computing and communication capabilities, coupled with infrastructure-based deployments of sensors and computing equipment. Specifically, we envision a system architecture that includes in-vehicle computing systems, roadside compute servers (e.g., embedded in traffic signal controllers) as well as servers residing in traffic management centers (TMCs). The remaining sections provide an overview of specific research examining key elements of this system. The next section describes a concept called ad hoc distributed simulations that are used to project future system states.

2 Ad Hoc Distributed Simulations

Consider a collection of in-vehicle simulations that are interconnected via wireless links and (possibly) wired network infrastructure. Individually, each simulation only models a portion of the traffic network – that which is of immediate interest to the “owner” of the simulator, Figure 1.

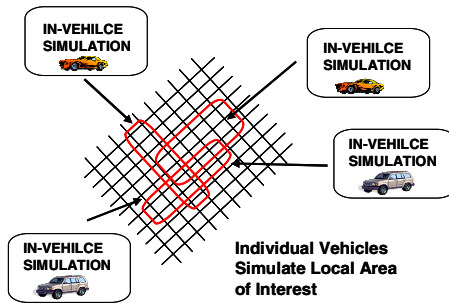


Fig. 1. In-Vehicle Simulation

Collectively, these simulations could be used to create a kind of distributed simulation system with the ability to make forecasts concerning the entire transportation infrastructure as a whole. One can envision combining in-vehicle simulators with simulations running within the roadside infrastructure, e.g., within traffic signal controller cabinets, simulations combining sub-regions of the transportation network, and simulations running in traffic management centers to create a large-scale model of a city’s transportation system, as shown in Figure 2.




Fig. 2. Ad Hoc Distributed Simulation Structure

We term a collection of autonomous, interacting simulations tied together in this fashion an *ad hoc distributed simulation*. Like a conventional distributed simulation, each simulator within an ad hoc distributed simulation models a portion of the overall system, and simulators exchange time stamped state information to collectively create a model of the overall system. However, in a conventional distributed simulation the system being modeled is designed in a *top-down* fashion. Specifically, the system is neatly partitioned into non-overlapping elements, e.g., geographic regions, and a process is assigned to model each element. By contrast, an ad hoc distributed simulation is created in a *bottom-up* fashion with no overarching intelligence governing the partitioning and mapping of the system to processes. Rather, the distributed simulation is constructed in an “ad hoc” fashion, in much the same way an arbitrary collection of mobile radios join together to form an ad hoc wireless network. The elements of the physical system modeled by different simulators in an ad hoc distributed simulation may overlap, leading to possibly many duplicate models of portions of the system, as seen in Figure 1. Other parts of the system may not be modeled at all. For example, an in-vehicle transportation simulator may be only modeling the portion of the road network along the vehicle’s intended path to reach its destination. Thus, ad hoc distributed simulations differ in important fundamental ways from conventional distributed simulations.

Ad hoc distributed simulations are on-line simulation programs, meaning they are able to capture the current state of the system through measurement, and then execute forward as rapidly as possible to project a future state of the system. By assumption, each of the simulators making up the distributed simulation can simulate some portion of the system faster than real time.

Ad hoc distributed simulations require a synchronization protocol to coordinate interactions among other simulations. For this purpose we have developed an optimistic (rollback-based) synchronization protocol designed for use in these systems. Each in-vehicle simulator utilizes information concerning traffic conditions and predictions of future system states (e.g., road flow rates) to complete its

simulation. If this information changes beyond certain parameterized limits, the simulator rolls back, and corrects its previously computed results. Based on this protocol, a prototype ad hoc distributed simulation system has been developed using both a custom-developed cellular automata traffic simulator as well as a commercial simulation tool called VISSIM, described next. Further details of this work are presented in [4].

3 Transportation Simulation Models

An ad hoc transportation simulation based on a cellular automata model was developed for our initial investigations. The simulation consists of agents modeling vehicles, traffic signal controllers and traffic lights. The simulation operates in a timestep fashion. At every timestep, each agent decides what operation to perform next based on its own characteristics and the state of the system at the end of previous interval. Each vehicle agent includes characteristics such as origin, destination, maximum speed of the vehicle, and driver characteristics such as aggressiveness. Each vehicle has complete knowledge of the road topography around it. At the end of each timestep each vehicle agent makes decisions whether to move, stop, accelerate, or decelerate based on the following four rules:

- *Acceleration*: if the velocity v of a vehicle is lower than max velocity and if the distance to the next car ahead is larger than $v + \text{gap}$, the speed is increased: $v = v + \text{acceleration}$.
- *Slowing down* (due to other cars): If a vehicle at site X sees the next vehicle at site $X + j$ (with $j < v$), it reduces its speed to j : $v = j$
- *Randomization*: The velocity a vehicle, if greater than zero, is decreased with probability p_{dec} : $v = v - \text{deceleration}$.
- *Car motion*: each vehicle is advanced v cells.

Similarly, the traffic controller may also change its state, specifically the stop-go flag, at the end of each time step.

We conducted experiments in the management of 20 client simulations using the cellular automata simulator covering a 10 intersection corridor. Under steady conditions, the distributed simulation client provided a system representation similar to that of a replicated trial experiment of the entire network, demonstrating that the ad hoc approach offers potential for accurately predicting future systems states. Further, when a spike in the traffic flow was introduced into the network, the distributed clients again successfully modeled the traffic flows; however there was a short delay (up to approximately four minutes) in the upstream client transmitting the increased flows to the downstream clients.

While the cellular automata simulation model allowed for the development of an understanding of the behavior of the ad hoc distributed approach it is desirable that future experimentation be conducted with a significantly more detailed, robust transportation simulation model. Further, the ability to adapt existing commercial simulation software for use in ad hoc distributed simulations would significantly improve the likelihood that the technology would be used. An initial investigation into implementing the ad hoc strategy using the off-the-shelf transportation simulation

model VISSIM was conducted. VISSIM is widely used by private firms and public agencies for transportation system analysis. VISSIM is a discrete, stochastic, time step based microscopic simulation model developed to model urban traffic (freeway and arterial) and public transit operations. The model is capable of simulating a diverse set transportation network features, such as facility types, traffic control mechanism, vehicle and driver types, etc. Individual vehicles are modeled in VISSIM using a psycho-physical driver behavior model developed by Wiedemann. The underlying concept of the model is the assumption that a driver can be in one of four driving modes: free driving, approaching, following, or braking [5]. VISSIM version 4.10 was utilized for this investigation.

Investigation into the use of VISSIM has proven very hopeful. Utilizing the VISSIM COM interface [6] it is possible to access many of the VISSIM object, methods, and properties necessary to implement simulation roll backs and automate VISSIM clients disbursed among workstations connected by a local area network. The initial investigation discussed in the section focuses on the ability to implement a roll back mechanism in VISSIM. Through the VISSIM COM interface it is possible at anytime t_i during a simulation run to save the current state of the simulation model. At any later time t_j , where $j > i$, it is possible to stop the simulation run, load the simulation state from time t_i , update the simulation attributes indicated by the roll back (i.e. arrival rate on some link), and restart the simulation from time t_i . An initial experiment was conducted using VISSIM demonstrating that the roll back algorithm was successfully implemented. This simulator was driven by traces of simulation-generated traffic flow data, and demonstrated to accurately predict future states (as represented from the trace data). Further details of these experiments are presented in [4].

4 Real-Time Dynamic Data Analysis

The precision of the real-time data varies depending on the level of data aggregation. For example, minute-by-minute data are more precise than hourly average data. We examined the creation of an accurate estimate of the evolving state of a transportation system using real-time roadway data aggregated at various update intervals.

Using the VISSIM model described in the previous section, a simulation of the transportation network in the vicinity of the Georgia Institute of Technology in Atlanta, Georgia was utilized to represent the real world, with flow data from the Georgia Tech model provided to a smaller simulation of two intersections within the network. The “real-world” flow data (i.e. flow data measured from the large scale Georgia Tech model) was aggregated in different intervals and used to dynamically drive the two-intersection model. The desire of this study was to explore how well the small-scale simulation model is able to reflect the real world scenario when fed data at different aggregation levels. This work explored congested and non-congested traffic demand at data five different aggregation time intervals: 1 sec., 10 sec., 30 sec., 60 sec., and 300 sec.

For the non-congested conditions, there existed minor differences in the average value of the considered performance metrics (arrival time and delay) and the performance metric difference values for the tested scenarios. However, there was a

clear trend of increasing the root mean square error (RMSE) as the aggregation interval increased. Varying the upstream origin of the arrival streams also tended to influence the RMSE values more than the average values. From these results it can be seen that under non-congested conditions, the average of performance metrics alone are likely not good indicators of the ability of a data driven simulation to reflect real world operations. Measures of variation such as RMSE should also be considered.

Unlike the non-congested conditions, the average values of the performance metrics in congested conditions were considerably different for the large real world simulation than the local simulation. The RMSE values also were significantly greater than those in the non-congested scenarios. There is also not a clear trend of the local simulation providing an improved reflection of the large simulation when given the smaller aggregation intervals. For the tested scenarios, the impact of congestion dominated the impact of a selected aggregation interval and upstream arrival pattern. The use of outflow constraints significantly improved the local model performance. These constraints helped capture the impact on the local simulation of congestion that occurs outside the local model boundaries. Where the boundaries of the congested region fall outside of the local simulation it becomes readily apparent that both the inflow and outflow parameters of the simulation must be dynamically driven to achieve a reasonable reflection of the real world conditions.

In the deployment of in-vehicle simulations these experiments highlight the need for realistic field measured inflow and outflow data streams, which are not currently widely available. However, as sensor technologies have advanced, the amount of available real-time field data is increasing dramatically. The quantity of available real-time data is expected to continue to climb at an ever-increasing rate. This tidal wave of real-time data creates the possibility of a wide variety of data driven transportation applications. This effort has begun to examine some of the innumerable potential uses of this data. Further details of these results are described in [7].

5 Data Dissemination

The proposed DDDAS system relies on mobile, wireless ad hoc networks to interconnect sensors, simulations, and servers. The experiments reported earlier were conducted over a local area network, and thus do not take into account the performance of the wireless network. However, the performance of the network infrastructure can play an important role in determining the overall effectiveness of the ad hoc distributed simulation approach. Thus, a separate effort has been devoted to examining the question of the impact of network performance on the results produced by the ad hoc distributed simulation.

Clearly, a single hop wireless link between the vehicles and servers severely limits the wireless coverage area, motivating the use of a multihop network. However, standard routing protocols designed for end-to-end communication in mobile ad hoc networks cause each node to maintain the state of their neighboring nodes, for efficient routing decisions. Such a solution does not scale to vehicular networks,

where the nodes are highly mobile, and route maintenance becomes expensive. Other routing protocols such as those discussed in [8,9] attempt to address this problem using flooding or optimistic forwarding techniques.

Vehicle ad hoc network (VANET) data dissemination protocols are typically designed to address data sharing among the vehicles, and related applications such as multimedia, control data, etc. Distributed simulations, however, define a different data transfer model, and consequently, the solutions designed for data sharing applications may not perform well when transposed to the demands of simulations. Additionally, data transfers in VANETs is inherently unreliable, and drops and delays in message delivery can be highly dependent on factors such as traffic density and wireless activity, thus having a strong impact on the simulation itself.

To address this challenge, a data dissemination framework for addressing the routing demands of a distributed simulation in the VANET environment has been developed. Our framework uses a combination of geographic routing and controlled flooding to deliver messages, with no organization enforced among the vehicles. The design parameters of the framework are currently under study in order to assess how this impacts the overall accuracy and reliability of simulation results.

6 Future Research

Research in ad hoc distributed simulations and their application to the management of surface transportation systems is in its infancy, and many questions remain to be addressed. Can such a distributed simulation make sufficiently accurate predictions of future system states to be useful? Can they incorporate new information and revise projections more rapidly and/or effectively than conventional approaches, e.g., global, centralized simulations? Does networking the individual simulators result in better predictions than simply leaving the individual simulators as independent simulations that do not exchange future state information? How would an ad hoc distributed simulation be organized and operate, and what type of synchronization mechanism is required among them? Will the underlying networking infrastructure required to support this concept provide adequate performance? These are a few of the questions that remain to be answered.

A broader question concerns the applicability of the ad hoc simulation approach to other domains. While our research is focused on transportation systems, one can imagine simulations like those discussed here might be used in other on-line simulation applications such as management of communication networks or other critical system infrastructures.

Acknowledgement

The research described in this paper was supported under NSF Grant CNS-0540160, and is gratefully acknowledged.

References

1. Werner, J.: Details of the VII Initiative 'Work in Progress' Provided at Public Meeting(2005)
2. Bechler, M., Franz, W.J., Wolf, L.: Mobile Internet Access in FleetNet. In KiVS (2003)
3. Werner, J.: USDOT Outlines the New VII Initiative at the 2004 TRB Annual Meeting(2004)
4. Fujimoto, R. M., Hunter, M., Sirichoke, J. Palekar, M. Kim, H.-K., Suh, W. :Ad Hoc Distributed Simulations. Principles of Advanced and Distributed Simulation(2007)
5. PTV, VISSIM User Manual 4.10. 2005, PTV Planung Transport Verkehr AG: Karlsruhe, Germany(2005).
6. PTV, VISSIM COM, User Manual for the VISSIM COM Interface, VISSIM 4.10-12. 2006, PTV Planung Transport Verkehr AG: Karlsruhe, Germany(2006).
7. Hunter, M. P., Fujimoto, R. M. , Suh, W., Kim, H. K. :An Investigation of Real-Time Dynamic Data Driven Transportation Simulation. In: Winter Simulation Conference(2006)
8. Rahman, W. Olesinski, Gburzynski, P.: Controlled Flooding in wireless ad-hoc networks. In: Proc. of IWWAN(2004)
9. Wu, H., Fujimoto, R., Guensler, R., Hunter, M.: MDDV: A Mobility-Centric Data Dissemination Algorithm for Vehicular Networks. In: Proceedings of the VANET Conference(2004)