

## STATISTICAL ISSUES IN AD HOC DISTRIBUTED SIMULATIONS

Ya-Lin Huang  
Richard Fujimoto

Georgia Institute of Technology  
Computational Science and Engineering  
Atlanta, GA 30332, USA

Wonho Suh  
Michael Hunter

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

Christos Alexopoulos

Georgia Institute of Technology  
Industrial and Systems Engineering  
Atlanta, GA 30332, USA

### ABSTRACT

An ad hoc distributed simulation is a collection of online simulators embedded in a sensor network that communicate and synchronize among themselves. Each simulator is driven by sensor data and state predictions from other simulators. Previous work has examined this approach in transportation systems and queueing networks. Ad hoc distributed simulations have the potential to offer greater resilience to failures, but also raise a variety of statistical issues including: (a) rapid and effective estimation of the input processes at modeling boundaries; (b) estimation of system-wide performance measures from individual simulator outputs; and (c) correction mechanisms responding to unexpected events or inaccuracies of the model itself. This paper formalizes these problems and discusses relevant statistical methodologies that allow ad hoc distributed simulations to realize their full potential. To illustrate one aspect of these methodologies, an example concerning rollback threshold parameter selection is presented in the context of managing surface transportation systems.

### 1 INTRODUCTION

Uses of online simulations are emerging as the needs for solving problems in operational systems in real-time become more and more common. Many real-world problems are sufficiently complex that analytical solutions may not exist or may be too complicated to be solved efficiently. Simulations offer an alternate approach where a model that mimics the underlying physical system is created and driven by sensor data and problem-specific configurations. A successful online simulation solution must satisfy real-time-related constraints and be equipped with the ability to capture the dynamics of the physical system.

Online simulations are also referred to as dynamic data-driven application systems (DDDAS), symbiotic simulations (Fujimoto et al. 2002), and cyber-physical systems in the literature. They have been widely studied and applied to various science and engineering disciplines for a diverse collection of purposes (Davis 1998). A typical application is to optimize the operation of a physical system. For example, in an emergency situation alternate evacuation scenarios may be modeled and evaluated in order to minimize evacuation time. The evacuation plan may need to adapt as the evacuation evolves with new unforeseen events arising. Other online simulation applications include path planning for unmanned aerial vehicles (Kamrani and Ayani 2007), parameter tuning for computer networks (Ye et al. 2008), management of

semiconductor manufacturing systems (Low et al. 2005), and optimization of surface transportation systems (Hunter et al. 2009a; Hunter et al. 2009b). Online simulations are also used to better understand and gain insights into the physical systems that are difficult or impossible to observe. Example applications include identifying an accident using cell phone data (Madey, Szabo, and Barabási 2006; Madey et al. 2007) and determining the boundary conditions of fluid-thermal systems (Knight, Rossman, and Jaluria 2006; Knight et al. 2007).

Many existing approaches to online simulations are centralized: sensor data are transmitted to a specific location and the simulation results (instructions or stimuli) are used to optimize the physical system. Communication can be problematic in so far as they consume energy and may incur large delays. For some sensor nodes the ratio of power consumption for communication compared to computation can be up to 10,000 to 1 (Zhao and Guibas 2004). Also communication failures may impact the effectiveness of online simulations; they must be expected in sensor networks especially under critical environmental conditions (e.g., inclement weather resulting in battery fail and radio interference). Scalability is another concern. In the centralized paradigm, modeling an enlarged physical system requires more computing resources and communication bandwidth in order to produce results in a timely fashion. These observations motivate replacing the centralized method by embedding simulations within the sensor network itself.

Ad hoc distributed simulations are an approach to real-time system monitoring, analysis, and operation optimization that involves embedding online simulations into the sensor network covering the physical system of interest. An ad hoc distributed simulation contains a collection of autonomous simulators (or logical processes in distributed simulation terminology), each modeling a portion of the physical system, referred to here as *coverage* or *modeling area*. The coverage areas can overlap. The simulation is constructed in a bottom-up fashion, as opposed to the top-down approach used by traditional distributed simulations where the physical system is partitioned into non-overlapping segments. Instead, in an ad hoc distributed simulation each logical process (LP) selects the modeling area based on its local objectives while collectively the LPs model the entire physical system. An optimistic, i.e., rollback-based synchronization approach is used where the simulation allows the LPs to advance ahead of other LPs. Figure 1 depicts an ad hoc distributed simulation. The grid represents the physical system; the LPs are co-located with the sensors, and the modeling areas are the regions shown on the grid. This figure also serves as an illustrative example for the following where the features of ad hoc distributed simulations are discussed.

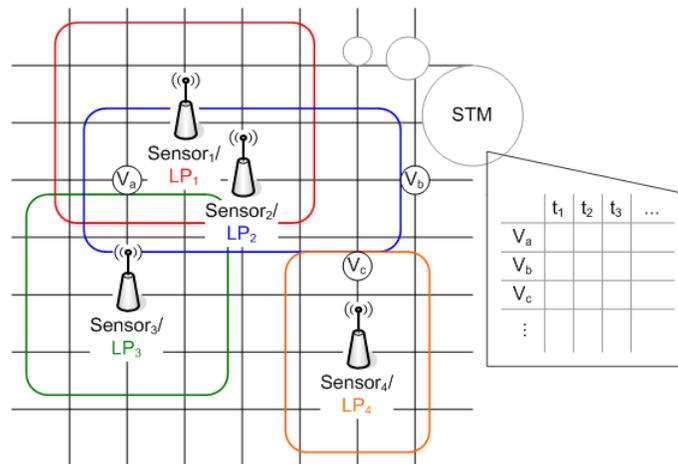


Figure 1: Illustration of an ad hoc distributed simulation

The designation of the area(s) modeled by each LP can be arbitrary. In some cases, an LP may select its modeling area in order to perform some local monitoring tasks such as predicting the travel time for a particular vehicle in a transportation network. Our previous work serves as an example where the LPs reside in vehicles while the sensors may be co-located with road-side cabinets or traffic signals (Hunter et

al. 2009a; Hunter et al. 2009b). In other cases, the model may simply predict future states of the area covered by a sensor. In Figure 1 each LP models the designated rectangular region covered by the corresponding sensor. This approach distributes the computational load and potentially reduces data transmission costs. It may lead to situations where some parts of the physical system are modeled by multiple LPs (e.g.,  $V_a$  in Figure 1) while others are left uncovered (e.g.,  $V_b$  in Figure 1). The former case introduces redundancy, offering the potential for greater robustness (resilience to failures), which differentiates the ad hoc approach from traditional distributed simulations where the system is perfectly partitioned so that each segment is modeled by exactly one LP. This may also increase prediction accuracy because multiple LPs will be used to provide state predictions. In addition, widely varying predictions by different LPs may be indicative of a malfunctioning LP (e.g., caused by incorrect model assumption or inaccurate sensor data) or an LP that has detected changes in system behavior in advance of other LPs.

LPs in an ad hoc distributed simulation may change their modeling areas as the simulation progresses, as might be the case for applications involving mobile components. For example, LPs monitoring a surface transportation system may be deployed in vehicles or handheld computing devices (e.g., GPS devices or smart phones). An LP might adjust its modeling area based on vehicle movements and driver's desires to receive predictions concerning the remainder of the planned route.

Current and future state predictions are shared among LPs via a construct called space time memory (STM; see Figure 1). The STM holds time-stamped system state updates from different LPs; the time interval associated with an update specifies when the update is valid. Other LPs may then read a system state by specifying the name of the desired variable and a timestamp. For example both  $LP_1$  and  $LP_2$  in Figure 1 update  $V_a$  while  $LP_3$  reads  $V_a$ . This approach is different from that used in traditional distributed simulations where the LPs are notified by events exchanged via messages. The STM aggregates state updates for the same variable produced by different LPs. Realization of the STM depends on the physical system in which the ad hoc distributed simulation executes. Ideally, the STM should be implemented in a distributed fashion over the simulation infrastructure, but could, in principle, be implemented centrally.

The optimistic synchronization method used in ad hoc distributed simulations allows LPs to predict future system states without necessarily waiting for other LPs. If certain desired state information is not available upon request, LPs may approximate the missing state information based on past state information or similar variables/objects. The approximation could be inaccurate and hence LPs rely on the rollback mechanism to detect disagreement between the used data and the predictions later produced by other LPs or the field data from sensors. Similarly, if the data previously received by an LP is determined to be in error, the rollback mechanism may be utilized to correct the erroneous data. The rollback mechanism, which is not unlike that used in Time Warp, rewinds the LP back to a time prior to when it started using the invalid data. It restores the previous system state and restarts the simulation with the "correct" data according to the currently available updates. Furthermore, the predictions produced by the LP prior to this rollback might be contaminated by the invalid input data; they should be revoked, resulting in re-computing the projections for the affected system states. Cascaded rollbacks are possible.

In this paper, we explore several open research issues that arise in ad hoc distributed simulations. The discussion is organized around the life cycle of a modeling and simulation study. Then, the remainder of the paper examines one specific issue, namely the design of the rollback triggering mechanism in the context of online management of traffic networks. Section 2 presents the analysis issues, Section 3 focuses on a transportation example, and Section 4 ends with concluding remarks.

## 2 ANALYSIS ISSUES

Ad hoc distributed simulations introduce a number of analysis issues. A typical simulation study includes the following steps: (1) problem formulation, (2) conceptual model development/validation and data collection, (3) simulation program development, verification, and validation, (4) experimental design and execution, (5) output analysis, and (6) result documentation (Law 2007). Here, we are primarily concerned with steps 2, 4, and 5 where the choice of simulation analysis methodology influences the design.

The issues regarding input data analysis (data collection and data model construction) are discussed in Section 2.1. Section 2.2 explores experiment design, Section 2.3 considers model execution and adaptation, and Section 2.4 focuses on output analysis.

## **2.1 Input Data Analysis**

Each LP in an ad hoc distributed simulation is essentially a conventional sequential simulation, but its input processes must be estimated using real-time data and it is subject to rollbacks. In particular, the input data that are used to construct appropriate input process models are generated by other LPs and/or sensors. A central problem concerns the definitions of these input processes. The data may have not only complex autocorrelation structures but also cross-dependencies due to the overlapping nature of the various LPs.

One possible approach for input approximation in an ad hoc distributed simulation was applied in transportation systems (Fujimoto et al. 2007; Hunter et al. 2009a; Hunter et al. 2009b). These papers listed three possible data sources for input process estimation: (1) projected state information from other LPs, (2) real-time traffic sensor data, and (3) historical traffic behavior patterns. In those experiments, boundary input rates of LPs were estimated based on the aggregated traffic flow rate predictions within rolling time windows.

The operations research literature contains several state-of-the-art methods for estimating dependent input processes (Biller and Ghosh 2006, Chick 2006) as well as methods for generating sample paths from such processes (Devroye 2006, Leemis 2006). These methods are typically time-consuming and are not designed for data sets that are generated from statistically dependent simulations. Therefore, their applicability in the dynamic setting of ad hoc distributed simulations is challenging, if not prohibitive.

Huang et al. (2010, 2011) adopted the paradigm of Whitt (1982, 1983) to model open queueing networks with independent routings. The arrival processes of the links between LPs' modeling areas are approximated by renewal processes with gamma-distributed interarrival times; the distribution parameters are estimated based on the data observed within rolling time windows. This approach is limited to open queueing networks, and its applicability to closed queueing systems, queueing networks with state-dependent routings, and transportation systems is currently under investigation.

Another issue concerning input approximation relates to the sampling mechanisms and sizes of the pertinent real-time data sets. If the constituent simulations are in steady state, extending the sampling period allows for improved input estimates. However, long sampling periods introduce challenges, including (1) higher computational requirements to perform input data analysis and (2) slower response to system transitions from stable conditions. The latter is a consequence of LPs modeling only portions of the entire physical system: changes to a specific part of the system would not be revealed to other LPs until the corresponding LPs share the information. This slow transient tracking may be undesirable and error-prone especially for the applications intended to capture changes in system behavior or anomalous events.

A method for obtaining more observations within a fixed sampling period is by deploying multiple LPs to model the overlapping area. The data grow in proportion to the number of LPs, allowing state variables to be estimated from an aggregation of the data from multiple LPs. However, the potential overlaps between the LPs' coverage areas result in correlated data streams; such correlations must be handled with care. Simple aggregation methods based on the mean or median are not suitable in general. In addition, regardless of the underlying aggregation method, the predictions from various LPs may be weighted differently. Greater weight could be given to the LPs projecting with higher fidelity, which is affected by the sensor capability, the computational power, or the modeling level of detail, to mention a few.

Finally, recall that ad hoc distributed simulations are online simulations embedded in sensor networks. Prediction accuracy is an objective, but speedy (faster than real-time) execution is the key to their applicability. The execution efficiency is controlled by many factors and the best approximation to the underlying data model is determined by not only the quality of the data sets but also the effectiveness of the input analysis methods under the computational resource constraints.

## 2.2 Experiment Design

The first issue encountered in designing experiments using the ad hoc approach concerns the assignment of coverage areas to LPs. We refer to this as the *system partitioning problem*. One solution is to assign each LP the area within which sensor data are locally available. This is cost effective with respect to communication as sensor data are consumed locally; long-distance transmission is not required. However, local data alone may be insufficient for making useful predictions or recommendations. Some applications may benefit from separating the physical system based on component similarity rather than geographical location. Examples include the systems involving a high level of interactions between components of the same category but very few across categories, e.g., modeling of computer virus spread in the Internet and information diffusion in social networks. A more general consideration would be the computational resource restrictions imposed by the devices where simulations are executed. The balance between the resource restrictions, the LP deployment/configuration complexities, and the desired output measures creates complex dependency between LPs, which is not yet completely understood.

A related issue concerns the choice of shared information. Ideally, an LP would share as much information as possible. However, communication and computational limitations (the simulation must execute faster than real-time in most applications) impede the ideal configuration. The chosen information must be sufficient to drive the input processes of other LPs and allow for the identification of state changes so that rollbacks, when necessary, may be utilized efficiently to catch up to system dynamics.

The *system coverage problem*, which refers to how many LPs are required to model a segment, has a potentially significant influence over output measure accuracy. Intuitively, a large number of LPs are preferred as this should increase the rate at which a system transient is recognized. However, a key aspect of multiple LPs will be the ability to distinguish between the LPs that are indicating changes in system states versus “outliers,” not reflecting a true system trend.

## 2.3 Model Execution and Adaptation

Sensor networks, where ad hoc distributed simulations are most likely to be carried out, are highly constrained environments with significant limitations and deficiencies. Limited battery energy necessitates a balance between the sensing and the simulation tasks, within which both computation and communication compete with each other for resources such as power. The computation is demanded for real-time response while the CPU clock speed and the memory storage may not be as powerful as in modern day computers. Wireless communications suffer from limited bandwidth, potential large latencies, and are error-prone. Since LP failures could be frequent, the presence of replicates is important to improve the robustness of the ad hoc approach.

The rollback mechanism in ad hoc distributed simulations enables recovery from the usage of incorrect input models. Optimistic synchronization mechanisms, i.e., those based on rollback, allow LPs to advance as much as possible without being held back by slower LPs. If a requested system state (as input to the simulation) is unavailable, the LP may approximate the input rather than wait. Approximation can be based on inference from historical or real-time data that are readily available; each method has different influences on the accuracy of the interim system state predictions. Nevertheless, the final prediction relies heavily on how the invalid input models are identified and corrected by the rollback mechanism, or specifically the rollback triggering criterion.

Good rollback triggering criterion must successfully differentiate between uses of invalid input data and normal statistical fluctuations. LPs do not need to be rolled back if differences between the adopted input model and the corresponding system state are the results of pure randomness or the expected patterns of fluctuations. As a consequence, one major issue is to define the acceptable level of difference between two models. In some cases it may be useful to evaluate the sensitivity of output metrics to variations in input data. A straightforward method would be to specify a tolerance range with fixed width. However, the appropriate width varies across cases; there is no general guideline to follow. Setting the width in a relative sense is a more reliable alternative (e.g., allowing a 10% deviation). With regard to sta-

tistical hypothesis tests, increasing the tolerance can be done by choosing a lower significance level, that is, lowering the probability of rejecting the null hypothesis that two models are statistically identical. However, this would potentially increase the type II error (i.e., failing to reject the false hypothesis of two models being considered identical).

Resource constraints add complexities to designing a rollback mechanism. The rollback mechanism clearly impacts the amount of required computation and communication resources, or more generally the amount of power to be consumed. When power is a limiting factor, it must be considered how to achieve the best mapping of an ad hoc distributed simulation over the available computing resources. In some cases, restricted optimism might be a solution to sacrifice execution efficiency in order to save resources from being excessively consumed by rollbacks.

## **2.4 Output Analysis**

Statistical analysis of output data from ad hoc distributed simulations introduces several challenges not present in conventional simulations. First, the overall predictions from an ad hoc distributed simulation are tightly coupled with the complex mechanisms involved in input process approximations as well as the rollback mechanism. While the simulation literature contains methods for adjusting point estimates and confidence intervals for the performance measures from conventional simulations to account for uncertainties in input parameters (Chick 2006, Zouaoui and Wilson 2004), these procedures are not suitable for the dynamic nature of ad hoc distributed simulations.

In steady-state simulations, removal of the initial transient effects is complicated by the rollback mechanism. In conventional simulations, this problem can be addressed by batching (Alexopoulos and Goldsman 2004). The batch means method is not directly applicable in ad hoc distributed simulations because the output processes can exhibit temporary departures from stationary state (e.g., due to communication failures or incident-driven changes to the field data).

Estimation of metrics spanning multiple LPs is also challenging. For example, estimating the distributions of travel times in an ad hoc traffic or queueing network simulation involves various complications. Consider the estimation of the mean travel time of a unit across a route, a rather trivial problem for conventional simulations. In an ad hoc distributed simulation, a route may cross areas covered by different (potentially overlapping) LPs. If the route is acyclic one can form an optimal linear combination of the estimators from corresponding LPs. If units can follow cyclic routes, as in queueing networks with probabilistic routing mechanisms, this estimation problem becomes even more difficult.

## **3 AD HOC TRANSPORTATION SIMULATIONS**

As pointed out in Sections 2.2 and 2.3, rollback threshold selection is complicated by the tradeoffs among simulation objectives, rollback mechanism designs, and computing constraints. Selected threshold measures need to be sensitive to state changes so that rollbacks can be triggered to capture system dynamics, but not so sensitive as to result in a large number of unnecessary rollbacks. This section explores the performance of different rollback threshold selections under various traffic conditions.

### **3.1 Previous Work**

Our initial experiments were conducted on a notional intersection arterial model (Fujimoto et al. 2007; Hunter et al. 2009a; Hunter et al. 2009b). In these experiments, three stationary configurations with constant arrival rates were tested along with a configuration where the arrival rate was significantly raised over time. Initial results from these experiments demonstrated the positive potential of ad hoc distributed simulations with regard to predictions in transportation systems. While the previous experiments focused on overall performance of the ad hoc approach in different traffic conditions with the predetermined parameter setup, here we delve more deeply into the parameter selection, evaluating the performance of the ad hoc approach with different rollback threshold configurations under various traffic conditions. The goal of this sensitivity analysis is to give greater insight into the ad hoc approach.

### 3.2 Experimental Environment and Traffic Network

VISSIM, a widely used off-the-shelf traffic simulation program, is utilized for this study. VISSIM is a discrete, stochastic, time-stepped microscopic simulation software package that models a wide range of traffic systems including freeways, arterials, and public transit operations (PTV 2009). The experiments are performed over a set of heterogeneous personal computers connected through a local area wireless network. In total, 11 machines are utilized: one for the STM and each of the remaining ten assigned one LP. Each machine is a laptop computer with a 2.2 GHz dual-core CPU and 2 GB RAM.

Figure 2 illustrates the modeled traffic network. This Manhattan-style 3-by-6 grid network consists of a two-way, 8-lane road (Fifth Street) with all other roads being two-way, 4-lane facilities. Each LP models a 3-by-3 grid network. Five LPs (LPs 1–5) simulate the 3-by-3 grid network covering the western half of the network (white box in Figure 2), and the remaining five (LPs 6–10) model the eastern half of the network (grey box in Figure 2). It is assumed that LPs 1–5 are preconfigured to model the designated scenario area at the start of a run. LPs 6–10 may experience a rollback when the currently used boundary link flow rate and the average of corresponding predictions provided by LPs 1–5 differ by a preset threshold. The duration of each experiment is 90 minutes in simulation time, including a 30-minute warm-up to allow the system to reach steady state; the data from the warm-up period are ignored.

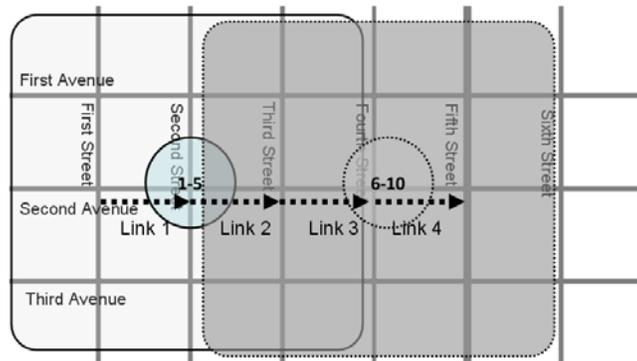


Figure 2: Modeled traffic network

### 3.3 Scenarios

The value of rollback threshold may influence the communication overhead and the output accuracy. Intuitively, a smaller threshold value would improve the accuracy while introduce more rollbacks and communication overhead. Thirty scenarios are constructed, which involve five different threshold values, four different initial traffic flow rates, and four different levels of flow rate increase; details are summarized in the first four columns of Table 2. For example, Scenarios 1–5 have different thresholds (60, 150, 240, 330, and 420 veh/hr, respectively) while their traffic conditions are identical: initial flow rate 100 veh/hr and a 400 veh/hr increase in flow rate at simulation time 20 minutes. The threshold of 60 veh/hr is equivalent to 2 vehicles per signal cycle length (2 minutes). Although these scenarios cannot represent all possible traffic conditions, they serve as a general view on the relationship between the rollback threshold value and the prediction accuracy. Ten replications are performed for each scenario. The results are compared with the “ground truth,” which is generated by modeling the entire network in a sequential fashion; similarly, 10 independent replications are executed for each corresponding traffic condition.

### 3.4 Results and Discussion

We study the influences of different rollback threshold values by evaluating the mean absolute error (MAE) of the flow rate estimates of Link 4. The definitions are in Equation (1) and Table 1:

$$\text{MAE} = \frac{1}{T} \sum_{t=1}^T \text{MAE}_t \quad \text{where} \quad \text{MAE}_t = \frac{1}{10} \sum_{i=1}^{10} |\text{AdHocRun}_{i,t} - \text{GroundTruthAverage}_t| \quad (1)$$

Table 1: Measurements summary

Notation	Description
$\text{MAE}_t$	Mean absolute error of the flow rate estimates of Link 4 for simulation time interval $t$
$\text{AdHocRun}_{i,t}$	Mean flow rate estimates of Link 4 for simulation time interval $t$ , produced by LPs in $i$ th ad hoc run
$\text{GroundTruthAverage}_t$	Mean flow rate estimates of Link 4 for simulation time interval $t$ , averaged over 10 sequential replicated runs
$T = 60$	Number of simulation intervals

Experiment results are shown in Table 2. In general, as the threshold value decreases (i.e., stricter rollback criteria), the overall output accuracy improves while more rollbacks occur. As expected, the smallest threshold value (60 veh/hr) provides the most accurate predictions over all tested traffic conditions. However, in some cases a larger threshold results in better performance in output accuracy than smaller thresholds. For example, the threshold of 330 veh/hr in Scenario 4 produces higher accuracy than the smaller thresholds (150 veh/hr and 240 veh/hr in Scenarios 2 and 3, respectively) between 31–50 minutes. We believe that this inconsistency is due to the fact that the difference between the initial flow rate and the increased flow rate is just a little larger than the selected threshold in Scenario 4. Therefore the LPs in Scenario 4 were able to publish their flow rate predictions with better accuracy than in Scenarios 2 and 3; see Figure 3 (a). A similar phenomenon is revealed in Scenarios 8 and 17 (compared to Scenarios 7 and 16, respectively), where the threshold of 240 veh/hr offers higher prediction accuracy than the threshold of 150 veh/hr. However, it is also noted that the threshold of 330 veh/hr is unable to capture the traffic flow changes from 100 veh/hr to 400 veh/hr (Figure 3 (b)) and 200 veh/hr to 500 veh/hr (Figure 3 (c)) in Scenarios 9 and 18, respectively. Since the flow rate increment (300 veh/hr) is smaller than the selected rollback threshold (330 veh/hr), rollbacks cannot be triggered.

Table 2 also shows the number of rollbacks triggered in each scenario. As the threshold value increases, the number of rollbacks decreases as expected. Since the communication overhead is positively correlated to the number of rollbacks, the tradeoff between the output accuracy and the communication overhead is obvious. The communication overhead could be a major issue in real-world applications of ad hoc distributed simulations. Although 60 veh/hr (the smallest threshold value in the experiments) provides the best accuracy, small thresholds may not be appropriate in practice. They might lead to excessive rollbacks, especially when the threshold values are lower than the expected variations in traffic flow conditions, such as the fluctuations in flow rate due to traffic control. Recall that the objective of the rollback threshold is to identify changes in traffic conditions, not any slight variation.

## 4 CONCLUSIONS

The great complexity of improving operational systems in real-time motivates new techniques for online simulations. While traditional centralized approaches limit the interaction between the simulation and the physical world, ad hoc distributed simulations provide an opportunity to improve efficiency, effectiveness, flexibility, scalability, and robustness by embedding the online simulation into the sensor network that monitors the physical system. In this paper, we explored analysis issues that arise in the ad hoc approach. Issues include input process estimation, rollbacks for model adaptation, and prediction measurement and evaluation. Furthermore, we studied how rollback triggering criterion affect execution efficiency and prediction accuracy in the online management of traffic networks. As anticipated we found that

more stringent criterion often resulted in better output accuracy but worse efficiency in performance. This study also illustrated some of the tradeoffs that arise in an ad hoc distributed simulation as functions of the system parameters.

Table 2: Mean Absolute Error in flow rate of link 4

Scenario ID	Initial Traffic Flow Rate	Increased Traffic Flow Rate (at 20 minute)	Rollback Threshold	Mean Absolute Error (veh/hr)			Number of Rollbacks
				Before Increase 1–20 minutes	During Increase 21–30 minutes	After Increase 31–50 minutes	
1	100	500	60 vph (2 veh/cycle)	7.1	22.6	26.0	203
2			150 vph (5 veh/cycle)	7.5	24.2	41.4	100
3			240 vph (8 veh/cycle)	9.2	35.2	61.0	50
4			330 vph (11 veh/cycle)	7.2	38.2	29.4	50
5			420 vph (14 veh/cycle)	11.0	67.6	323.3	0
6	100	400	60 vph (2 veh/cycle)	9.2	20.0	19.9	151
7			150 vph (5 veh/cycle)	9.0	25.1	84.3	65
8			240 vph (8 veh/cycle)	8.6	24.6	20.2	50
9			330 vph (11 veh/cycle)	11.6	52.4	247.6	0
10	100	300	60 vph (2 veh/cycle)	6.6	22.8	20.0	109
11			150 vph (5 veh/cycle)	6.8	26.1	19.9	50
12			240 vph (8 veh/cycle)	5.2	48.7	163.2	0
13	100	200	60 vph (2 veh/cycle)	7.3	18.9	22.0	54
14			150 vph (5 veh/cycle)	7.0	28.6	83.7	0
15	200	500	60 vph (2 veh/cycle)	10.1	20.1	31.6	135
16			150 vph (5 veh/cycle)	10.3	26.7	60.3	65
17			240 vph (8 veh/cycle)	9.8	35.4	33.4	50
18			330 vph (11 veh/cycle)	11.2	50.7	239.7	5
19	200	400	60 vph (2 veh/cycle)	8.8	17.5	24.4	120
20			150 vph (5 veh/cycle)	7.1	24.0	25.1	50
21			240 vph (8 veh/cycle)	9.5	30.7	164.8	5
22	200	300	60 vph (2 veh/cycle)	11.3	26.5	25.4	90
23			150 vph (5 veh/cycle)	14.9	34.3	83.3	0
24	300	500	60 vph (2 veh/cycle)	12.6	19.6	25.7	215
25			150 vph (5 veh/cycle)	13.4	17.8	28.8	50
26			240 vph (8 veh/cycle)	9.7	21.6	147.6	10
27	300	400	60 vph (2 veh/cycle)	9.1	19.1	27.7	75
28			150 vph (5 veh/cycle)	19.5	27.6	82.1	5
29	400	500	60 vph (2 veh/cycle)	20.5	21.6	23.5	161
30			150 vph (5 veh/cycle)	17.3	12.6	86.7	5

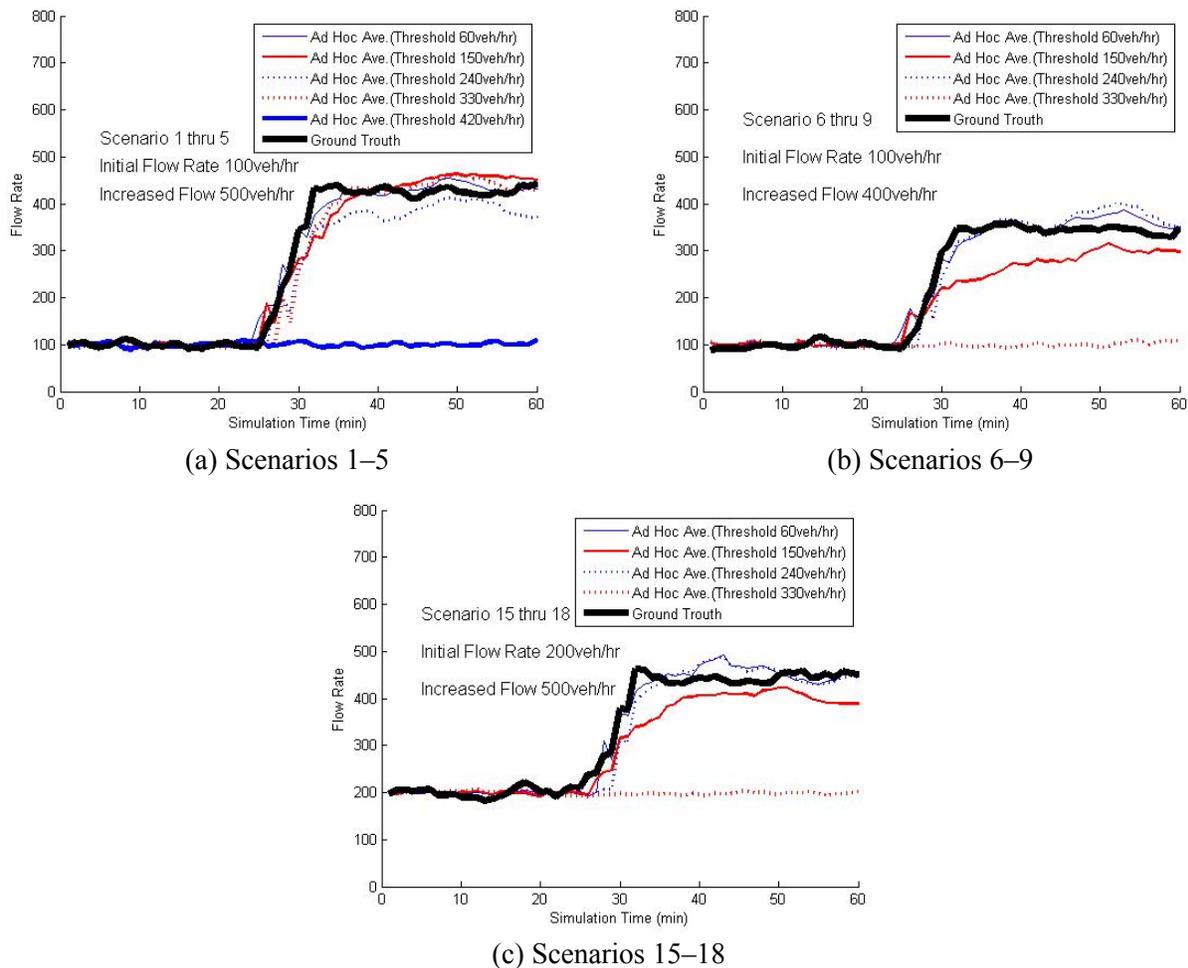


Figure 3: Flow rate estimates of link 4

## REFERENCES

- Alexopoulos, C., and D. Goldman. 2004. "To Batch or Not to Batch?" *ACM Transactions on Modeling and Computer Simulation* 14: 76–114.
- Billar, B., and S. Ghosh. 2006. "Multivariate Input Processes." In *Handbooks in Operations Research and Management Science, Volume 13: Simulation*, edited by S. G. Henderson and B. L. Nelson, 123–153. Amsterdam: Elsevier.
- Chick, S. E. 2006. "Subjective Probability and Bayesian Methodology." In *Handbooks in Operations Research and Management Science, Volume 13: Simulation*, edited by S. G. Henderson and B. L. Nelson, 225–257. Amsterdam: Elsevier.
- Davis, W. J. 1998. "On-line Simulation: Need and Evolving Research Requirements." In *Handbook of Simulation: Principles, Methodology, Advances, Applications, and Practice*, edited by J. Banks, 465–518. New York: Wiley.
- Devroye, L. 2006. "Nonuniform Random Variate Generation." In *Handbooks in Operations Research and Management Science, Volume 13: Simulation*, edited by S. G. Henderson and B. L. Nelson, 83–121. Amsterdam: Elsevier.

- Fujimoto, R., D. Lunceford, E. Page, and A. M. Uhrmacher. 2002. "Grand Challenges for Modeling and Simulation: Dagstuhl Report." Technical Report No. 350, Dagstuhl Seminar No. 02351, Schloss Dagstuhl, Germany.
- Fujimoto, R., M. Hunter, J. Sirichoke, M. Palekar, H. Kim, and W. Suh. 2007. "Ad Hoc Distributed Simulations." In *Proceedings of the 21<sup>st</sup> International Workshop on Principles of Advanced and Distributed Simulation*, 15–24. Washington, District of Columbia: Institute of Electrical and Electronics Engineers Computer Society, Inc.
- Huang, Y.-L., C. Alexopoulos, M. Hunter, and R. M. Fujimoto. 2010. "Ad Hoc Distributed Simulation of Queueing Networks." In *Proceedings of the 24<sup>th</sup> International Workshop on Principles of Advanced and Distributed Simulation*, 57–64. Washington, District of Columbia: Institute of Electrical and Electronics Engineers Computer Society, Inc.
- Huang, Y.-L., C. Alexopoulos, R. Fujimoto, and M. Hunter. 2011. "On the Accuracy of Ad Hoc Distributed Simulations for Open Queueing Network." In *Proceedings of the 25<sup>th</sup> International Workshop on Principles of Advanced and Distributed Simulation*, 163–168. Washington, District of Columbia: Institute of Electrical and Electronics Engineers Computer Society, Inc.
- Hunter, M., H. K. Kim, W. Suh, R. Fujimoto, J. Sirichoke, and M. Palekar. 2009a. "Ad Hoc Distributed Dynamic Data-driven Simulations of Surface Transportation Systems." *Simulation* 85:243–255.
- Hunter, M., J. Sirichoke, R. Fujimoto, and Y.-L. Huang. 2009b. "Embedded Ad Hoc Distributed Simulation for Transportation System Monitoring and Control." In *Proceedings of the 2009 INFORMS Simulation Society Research Workshop*, edited by L. H. Lee, M. E. Kuhl, J. F. Fowler, and S. Robinson, 13–17. Hanover, MD: INFORMS.
- Kamrani, F., and R. Ayani. 2007. "Using On-line Simulation for Adaptive Path Planning of UAVs." In *Proceedings of the 11<sup>th</sup> IEEE International Symposium on Distributed Simulation and Real-time Applications*, edited by D. J. Roberts, G. K. Theodoropoulos, and A. E. Saddik, 167–174. Washington, District of Columbia: Institute of Electrical and Electronics Engineers Computer Society, Inc.
- Knight, D., T. Rossman, and Y. Jaluria. 2006. "Evaluation of Fluid-thermal Systems by Dynamic Data Driven Application Systems." In *Proceedings of the 2006 International Conference on Computational Science*, edited by V. N. Alexandrov, G. D. van Albada, P. M. A. Sloot, and J. Dongarra, 473–480. Berlin: Springer Verlag.
- Knight, D., Q. Ma, T. Rossman, and Y. Jaluria. 2007. "Evaluation of Fluid-thermal Systems by Dynamic Data Driven Application Systems – Part II." In *Proceedings of the 2007 International Conference on Computational Science*, edited by Y. Shi, G. D. van Albada, J. Dongarra, and P. M. A. Sloot, 1189–1196. Berlin: Springer-Verlag.
- Law, A. M. 2007. *Simulation Modeling & Analysis*. 4th ed. New York: McGraw-Hill, Inc.
- Leemis, L. M. 2006. "Arrival Processes, Random Lifetimes and Random Objects." In *Handbooks in Operations Research and Management Science, Volume 13: Simulation*, edited by S. G. Henderson and B. L. Nelson, 155–180. Amsterdam: Elsevier.
- Low, M. Y. H., K. W. Lye, P. Lendermann, S. J. Turner, R. T. W. Chim, and S. H. Leo. 2005. "An Agent-based Approach for Managing Symbiotic Simulation of Semiconductor Assembly and Test Operation." In *Proceedings of the 4<sup>th</sup> International Joint Conference on Autonomous Agents and Multiagent Systems*, edited by F. Dignum, V. Dignum, S. Koenig, S. Kraus, M. P. Singh, and M. Wooldridge, 85–92. New York: Association for Computing Machinery, Inc.
- Madey, G. R., G. Szabo, and A.-L. Barabási. 2006. "WIPER: The Integrated Wireless Phone Based Emergency Response System." In *Proceedings of the 2006 International Conference on Computational Science*, edited by V. N. Alexandrov, G. D. van Albada, P. M. A. Sloot, and J. Dongarra, 417–424. Berlin: Springer-Verlag.
- Madey, G. R., A.-L. Barabási, N. V. Chawla, M. Gonzalez, D. Hachen, B. Lantz, A. Pawling, T. Schoenharl, G. Szabó, P. Wang, and P. Yan. 2007. "Enhanced Situational Awareness: Application of DDDAS Concepts to Emergency and Disaster Management." In *Proceedings of the 2007 Interna-*

- tional Conference on Computational Science*, edited by Y. Shi, G. D. van Albada, J. Dongarra, and P. M. A. Sloot, 1090–1097. Berlin: Springer-Verlag.
- PTV. 2009. VISSIM User Manual 5.10. Karlsruhe, Germany: PTV Planung Transport Verkehr AG.
- Whitt, W. 1982. “Approximating a Point Process by a Renewal Process, I: Two Basic Methods.” *Operations Research* 30: 125–147.
- Whitt, W. 1983. “The Queueing Network Analyzer.” *The Bell System Technical Journal* 62: 2779–2815.
- Ye, T., H. Kaur, S. Kalyanaraman, and M. Yuksel. 2008. “Large-scale Network Parameter Configuration using an On-line Simulation Framework.” *IEEE/ACM Transactions on Networking* 16:777–790.
- Zhao, F., and L. J. Guibas. 2004. *Wireless Sensor Networks: An Information Processing Approach*. Amsterdam: Morgan Kaufman, Inc.
- Zouaoui, F., and J. R. Wilson. 2004. “Accounting for Input-model and Input-parameter Uncertainties in Simulation.” *IIE Transactions* 36:1135–1151.

## AUTHOR BIOGRAPHIES

**YA-LIN HUANG** is a Ph.D. student in the College of Computing at the 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, respectively. Her current research focuses on distributed simulations. Her e-mail is [huangyl@cc.gatech.edu](mailto:huangyl@cc.gatech.edu).

**WONHO SUH** is a Ph.D. student in the School of Civil and Environmental Engineering at the Georgia Institute of Technology. Prior to joining Georgia Tech, he worked as a research engineer for five years at Yooshin Corporation in South Korea. He holds an B.S. and M.S. in Civil Engineering from Seoul National University (1998 & 2000). His current research interest is microscopic traffic simulation. His e-mail is [Wonho.Suh@gatech.edu](mailto:Wonho.Suh@gatech.edu).

**CHRISTOS ALEXOPOULOS** is an Associate Professor in the H. Milton Stewart School of Industrial and Systems Engineering at the Georgia Institute of Technology. He received his B.S. in Mathematics from the Aristotle University in Greece and his Ph.D. in Operations Research from the University of North Carolina. His research interests are in the areas of computer simulation, applied probability, and optimization of stochastic systems. He is an area editor for *ACM Transactions on Modeling and Computer Simulation* and an associate editor for *Networks*. Dr. Alexopoulos has a leadership role in the *Winter Simulation Conference*, having served as proceedings co-editor for the 1995 conference, as associate program chair for the 2006 conference, and as a member of the Board of Directors since 2008, on behalf of the INFORMS Simulation Society. His e-mail is [christos@isye.gatech.edu](mailto:christos@isye.gatech.edu).

**RICHARD FUJIMOTO** is a Regents’ Professor and the Chair of the School of Computational Science and Engineering 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 led 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 e-mail is [fujimoto@cc.gatech.edu](mailto:fujimoto@cc.gatech.edu).

**MICHAEL HUNTER** is an Associate Professor at the School of Civil and Environmental Engineering at the Georgia Institute of Technology. His primary teaching and research interests are in transportation op-

erations, design and safety, specializing in 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 analyses, etc. His e-mail is [michael.hunter@ce.gatech.edu](mailto:michael.hunter@ce.gatech.edu).