Ad Hoc Distributed Simulations

Richard Fujimoto

Computational Science and Engineering Division

Georgia Institute of Technology

Acknowledgements

Computational Science and Engineering — *Richard Fujimoto*
Prashant Chari, Ya-Lin Huang, Ling Liu, Vishu Palagummi, Mahesh Palekar, Jason Sirichoke, Hao Wu

Civil & Environmental Engineering — *Michael Hunter, Randall Guensler, Frank Southworth*
Hoe-Kyoung Kim, Wonho Suh, Jaesup Lee, Joonho Ko, Hainan Li

Industrial & Systems Engineering — *Christos Alexopoulos*
Carlos Valencia

Computer Science — *Karsten Schwan*
Bala Seshasayee

Research supported by NSF grants EFRI-0735991, CNS-0540160, EIA-0219976, ECS-0225447 and Intel Corporation

Outline

- A Motivating Application
- Ad Hoc Distributed Simulations
  - Algorithms
    - Space-Time Memory
    - Synchronization algorithm
  - Example Simulations
- Communications
  - Vehicle-to-Vehicle Network Simulation
  - Field Measurement

The Costs of Mobility

- **Safety**: 6 Million crashes, 41,000 fatalities in U.S. per year ($150 Billion)
- **Congestion**: 3.5 B hours delay, 5.7 B gal. wasted fuel per year in U.S. ($65 Billion)
- **Pollution**: > 50% hazardous air pollutants in U.S., up to 90% of the carbon monoxide in urban air

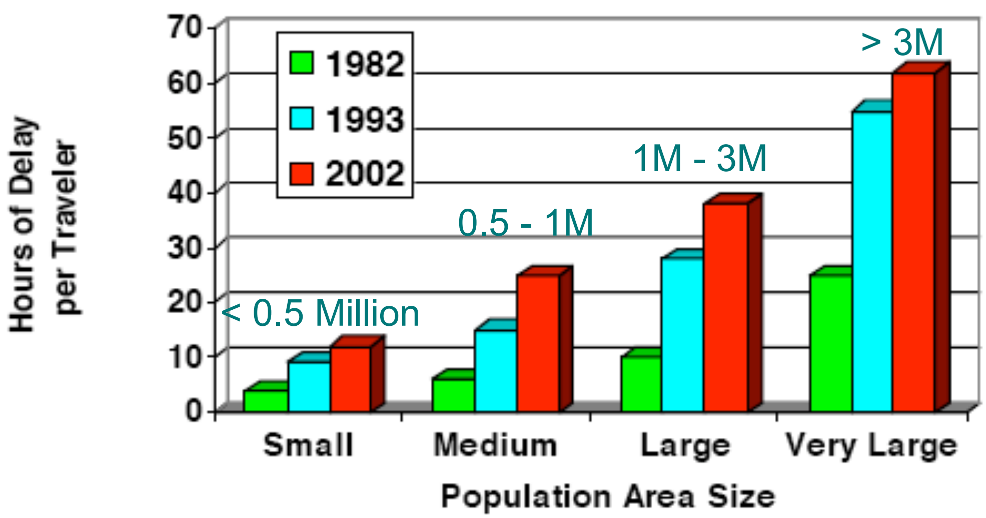

- Disproportionate increase in car ownership relative to population growth in China, India

Source: 2005 Annual Urban Mobility Report (http://mobility.tamu.edu)
Texas Natural Resource Conservation Commission (http://www.tnrcc.state.tx.us/air)

Intelligent Transportation Systems

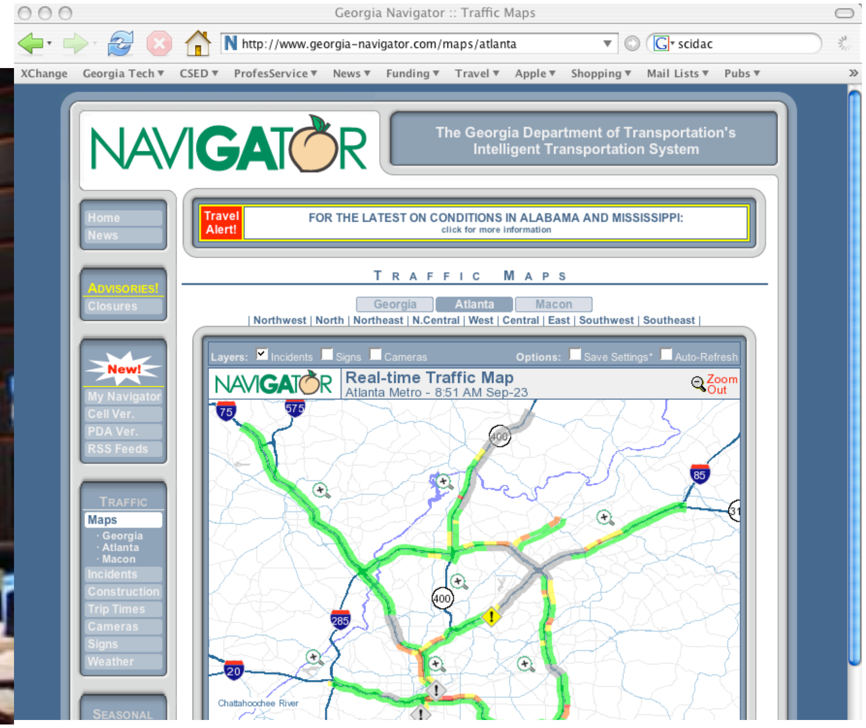

- ITS deployments: Traffic Management Centers (TMC)
  - Roadside cameras, sensors, communicate to TMC via private network
  - Disseminate information (web, road signs), dispatch emergency vehicles
- Infrastructure heavy
  - Expensive to deploy and maintain; limited coverage area
  - Limited traveler information
  - Limited ability to customize services for individual travelers

Current Trends

Smart Vehicles

- On-board GPS, digital maps
- Vehicle, environment sensors
- Significant computation, storage, communication capability
- Not power constrained

Dedicated Short Range Communications (DSRC)

- 5.850-5.925 GHz
- V2V, V2R communication
- 802.11p protocol
- 7 channels, dedicated safety channel
- 6- 27 Mbps
- Up to 1000 m range

U.S. DOT Vehicle Infrastructure Integration (VII) Initiative

- Public/private partnership
- "Establishment of vehicle-to-vehicle and vehicle-to-roadside communication capability nationwide"
- Improve safety, reduce congestion

The Future …

Applications

- Collision warning/ avoidance
- Traffic monitoring
- Emergency vehicle warning
- Internet Access
- Traveler & Tourist Assistance
- Entertainment

Ubiquitous Transportation Simulations

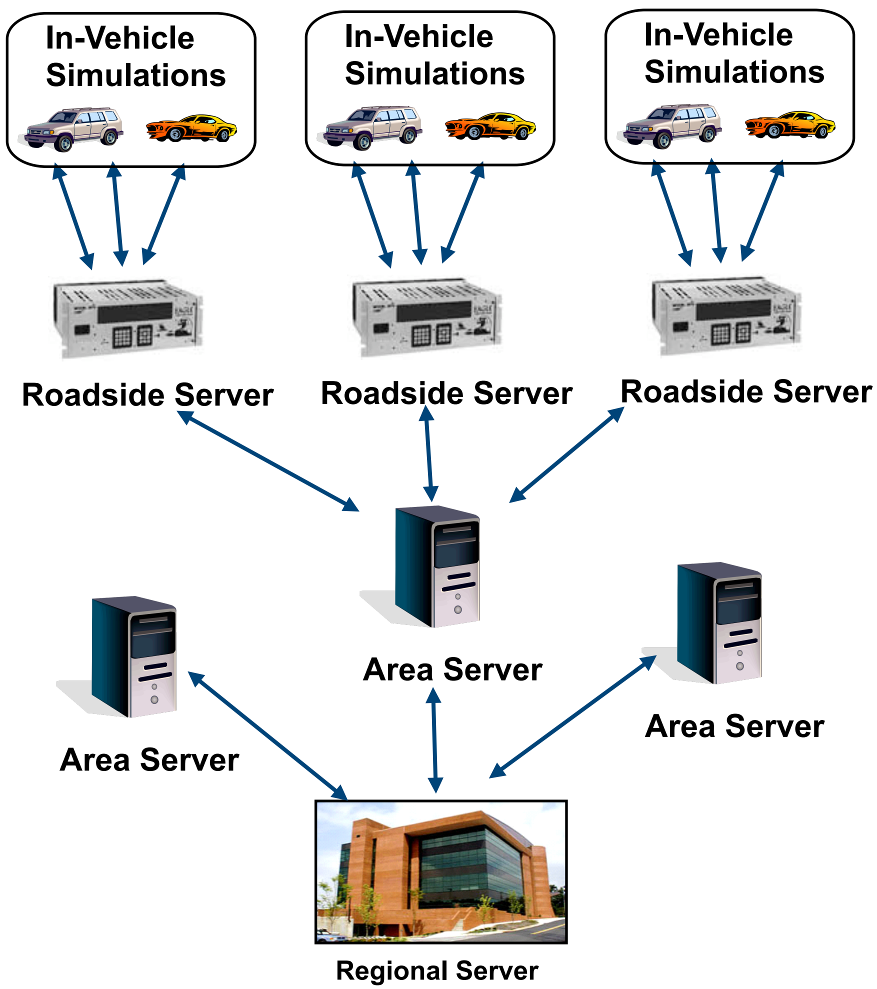

- New uses of simulation in transportation system management
  - Roadside, in-vehicle computing augments traffic management centers
  - Vehicle networks augment transportation communication infrastructure
- Operating simulations close to data sources offers several advantages
  - Potentially high fidelity by utilizing local, detailed real-time data
  - Opportunity for short response time for local decision making
  - More robust, resilient to failures relative to centralized approaches
- Public/private infrastructure, reduces public sector deployment and maintenance cost

On-Line Distributed Simulation

- On-line simulation (aka Symbiotic Simulations, Dynamic Data-Driven Application Systems [DDDAS])
  - Collect sensor data from environment
  - Construct current system state from sensed data
  - Compute future states via simulation
  - Optimize system to steer toward desired system states
- Example applications
  - Manufacturing, Business Processes (NTU)
  - Telecommunications (UCLA, GT, UCB)
  - Preparation for Inclement Weather (Univ. of Oklahoma, Indiana, …)
  - Crisis Management (Purdue, …)
  - Defense (SAIC, …)

Ad Hoc Distributed Simulations

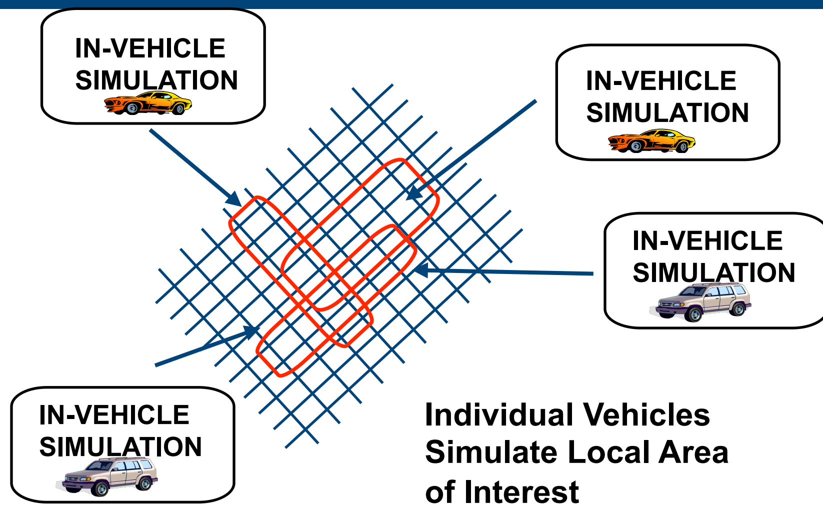

- An Ad Hoc distributed simulation is a composition of autonomous on-line simulators, each modeling its own “area of interest” independent of other simulators
  - Simulators may be stationary or mobile
  - Area of interest may vary over time
- Not a clean partitioning of physical system
  - Areas modeled by different simulators may overlap
  - Some areas may not be modeled at all

Conventional vs. Ad Hoc Distributed Simulation

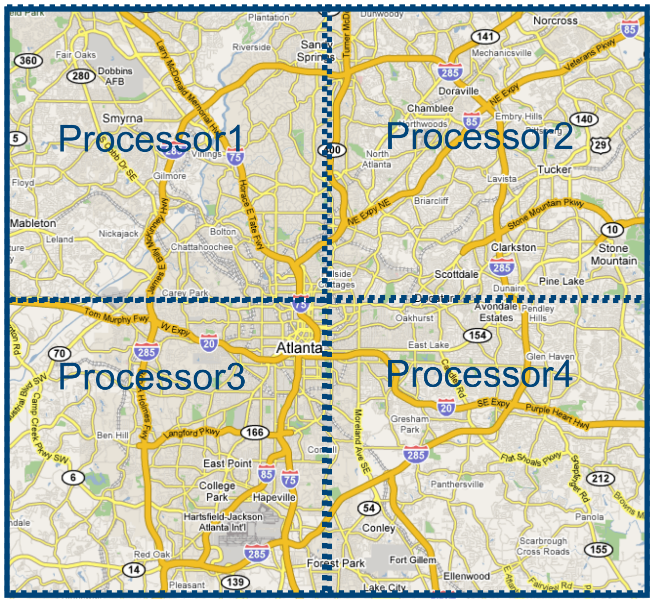

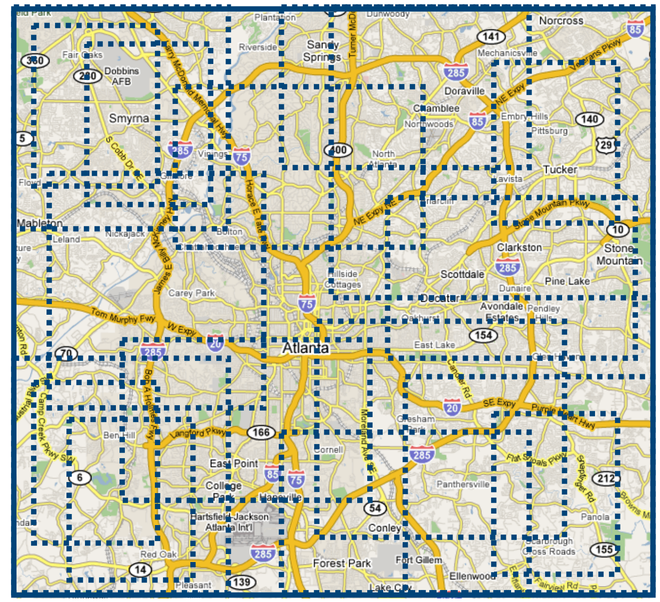

Conventional

- Top-Down construction
- Clean partition of state space; static partition
- Produce same results as a single run

Ad Hoc

- Bottom-Up construction
- Ad Hoc partition of state space; dynamic partition
- Produce same statistical results as replicated runs

Relationship to Other Simulation Approaches

- Conventional distributed simulations
  - In ad hoc simulations, multiple simulators compute the value of state variables
  - Synchronization algorithm needed to coordinate simulators, but based on aggregated state estimates rather than “correct” values of system state
- Replicated Trials
  - Replications model subsets of the entire physical system
  - Replicated simulations interact via the synchronization algorithm
  - Outlier simulations are not rolled back

State Prediction Questions

State prediction problems:

- Can a collection of localized simulations provide accurate predictions of the overall system state?
- Static prediction: Given a current snapshot of the state of the system, what is the predicted, future state?
- Dynamic predication: Given a *new* snapshot of the state of the system, what is the (revised) prediction of future system state?

Execution Mechanism

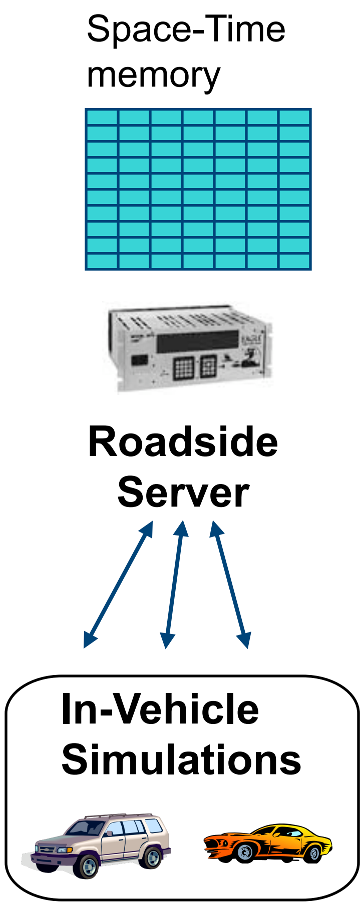

(other levels of hierarchy, e.g., regional, traffic management center not shown)

- System State: Space-Time Memory
  - Time stamp addressed memory
  - Stores current, predicted system state
- Autonomous simulators
  - Read current, predicted state from STM
  - Compute future state predictions
  - Provide updates to STM
- Optimistic synchronization (Rollback)
  - Prediction errors arise when
    - Sensor readings do not match predictions
    - Predictions from other simulators change
  - If error sufficiently large, roll back simulator and re-compute new projection

Space-Time Memory

- Collection of global state variables G_1, G_2, … G_N
  - May be read, updated by any simulator (logical process)
- Timestamped addressed memory
- Write operation
  - Write(G_i, v, t_1, t_2): estimate G_i's value to be v during the time interval [t_1 to t_2)
- Read operation
  - Multiple estimates to a single state variable for a given time
  - Read(G_i, t): returns $G_j(t)$ = C(G_j, t), where
  - Composite function C(G_j, t) computes an estimate of G_j at time t based on estimated values provided by other simulators
    - Simple average (possibly weighted)
    - Random value drawn from empirical distribution

STM Example

Writes to G_1: v@[t_1,t_2)

90@[5,8); 100@[8,15); 120@[10,15); 125@[12,20)

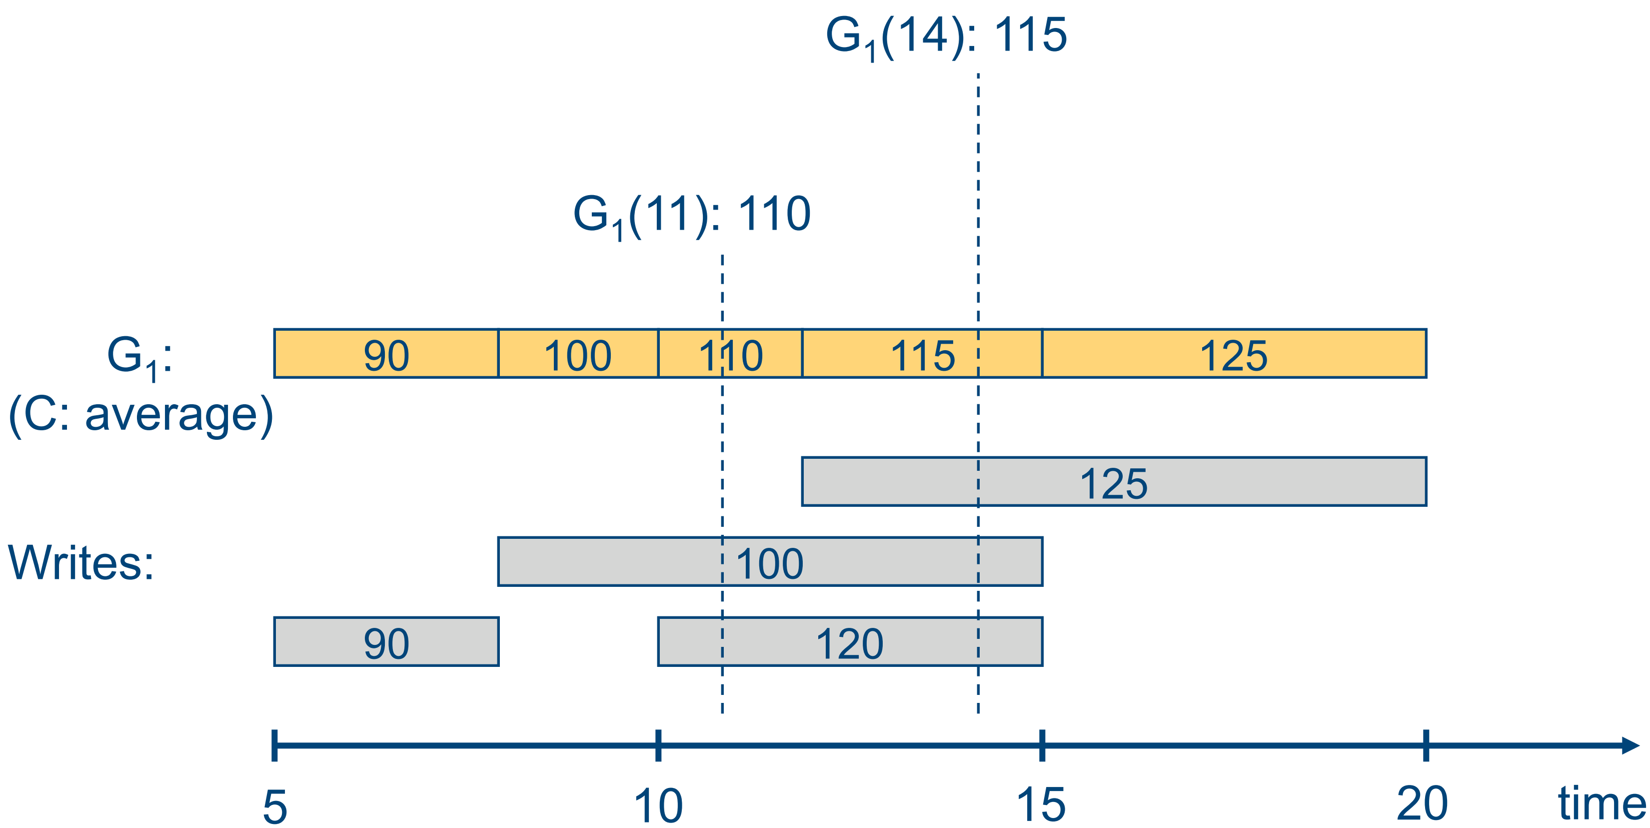

Synchronization

- Simulators predict future state of system based on on-line measurement
- These predictions may be wrong due to unexpected events (e.g., accidents)
- If prediction does not match measured state, roll back simulation, and re-compute new future state based on measured data
- If new predicted state very different from previously projected state, may trigger additional rollbacks (cascaded rollbacks)

Automated Update via Optimistic Synchronization

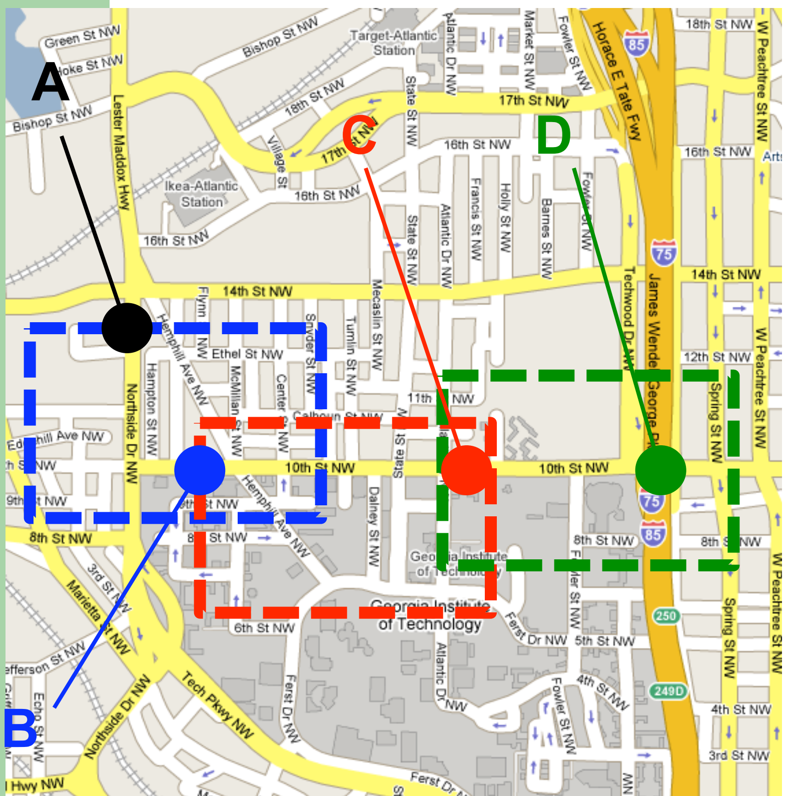

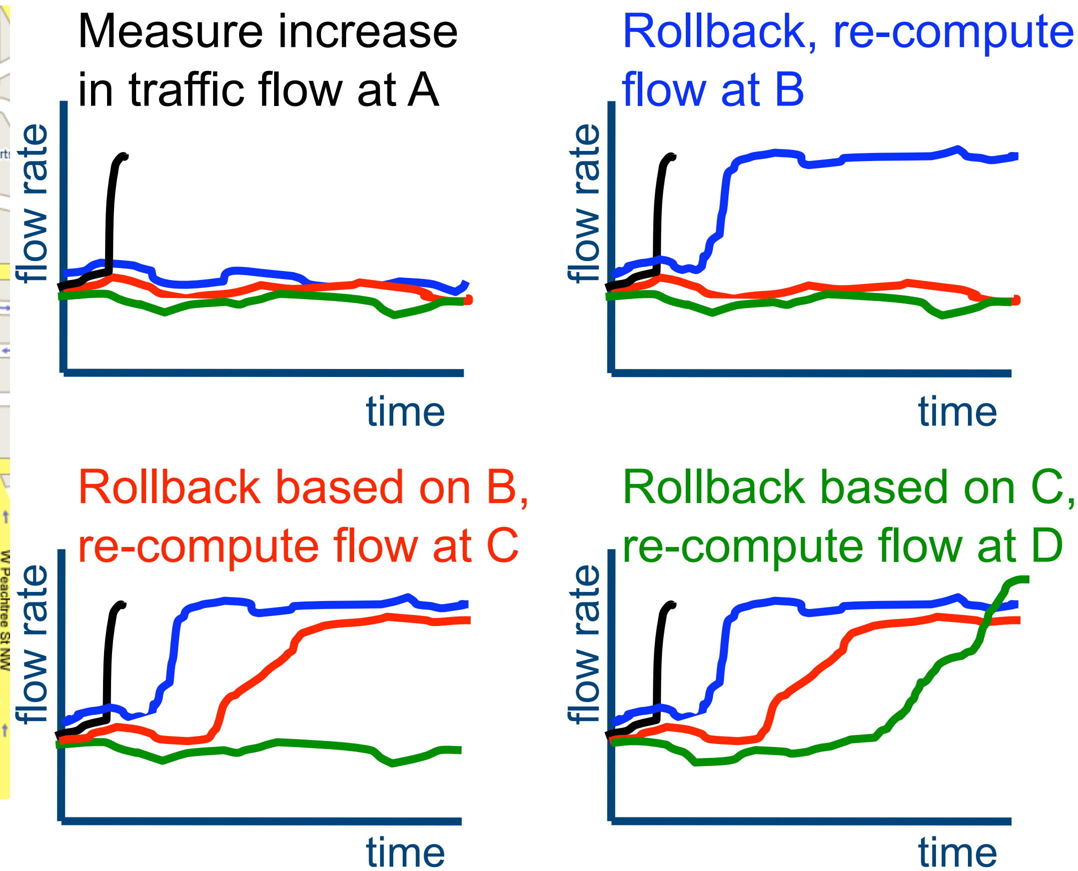

Roll back simulator when

- Prediction and measurement disagree
- Predictions from other simulators change

Synchronization Algorithm

- Collection of logical processes (LPs): LP_1, LP_2, LP_3, …
- Rollback function $R(G_i, v, t)$
  - $R(G_i, v, t) = |G_i(t)-v| > H$ for some threshold H
- LP_i: Read(G_j, t)
  - Return $v = G_j(t) = C(G_j, t)$
  - Log (LP_i, v, t) in STM
- LP_i: Write(G_j, v, t_1, t_2)
  - For each LP_k that read a value v' from G_j at time $t \in [t_1, t_2)$, if $R(G_j, v', t)$ then roll back LP_k to time t
    - Restore state of LP_k to that at time t
    - Generate anti-writes for writes performed by LP_k at times > t
- LP_k: Anti-Write(G_j, v, t_1, t_2)
  - Delete value v from STM, update composite value
  - For each LP_m that read the value v from G_j at time $t \in [t_1, t_2)$, if $R(G_j, v, t)$ then roll back LP_m to time t
    - Restore state of LP_m to that at time t
    - Generate anti-writes for writes performed by LP_m at times > t

Prototype Implementation

- Simulators
  - Custom traffic simulator
    - Cellular automata
    - Custom designed for ad hoc execution mechanism
    - Simplified models
  - Commercial simulator
    - VISSIM
    - Detailed, “industrial strength” microscopic traffic simulator
- Simulation infrastructure
  - Built over HLA RTI software (FDK package)

Cellular Automata Simulator

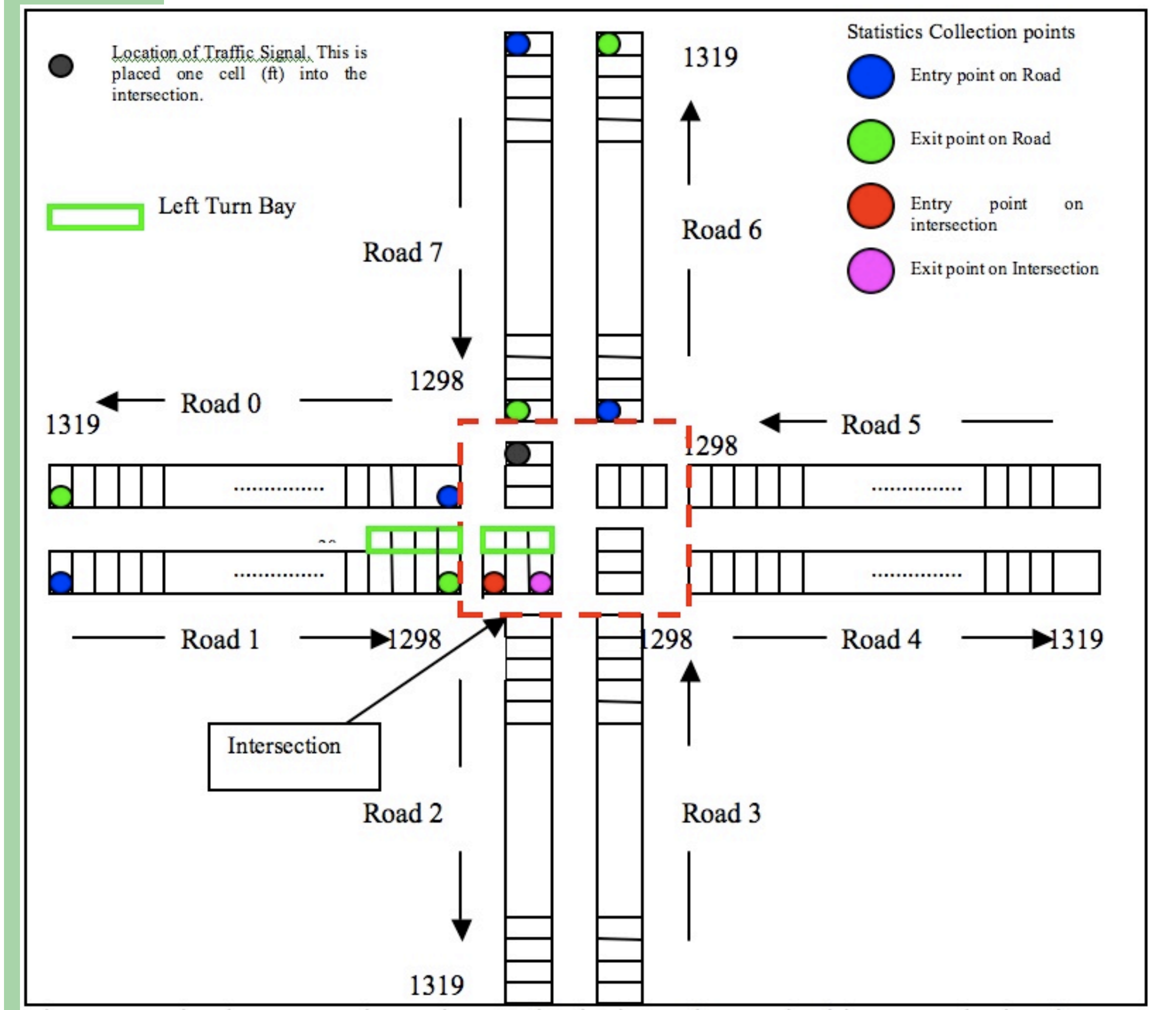

- Vehicle rules
  - Acceleration
  - Deceleration
  - Randomized speed change
  - Car motion
- Straight, turn probabilities (0.95, 0.02 left, 0.03 right)
- Signal timing
  - 120 second cycle
  - Left turn signals
- Timestep execution
- Global state: vehicle flow rate

Initial Test Network

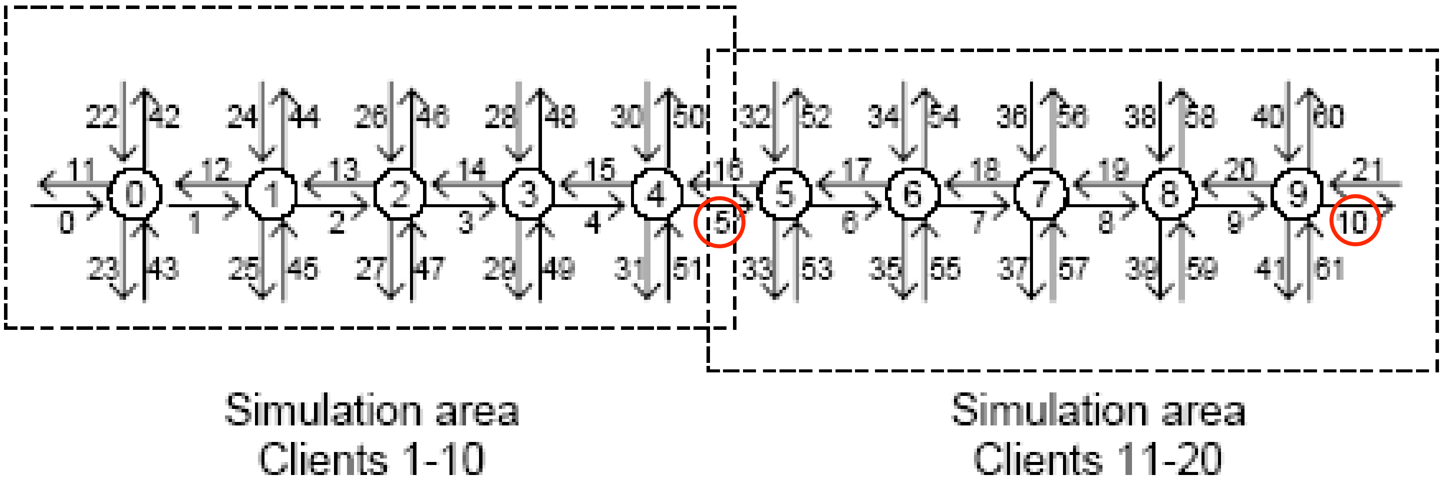

- Test Configuration
  - 20 in-vehicle simulators, each simulates half of the network
  - 1 server (space-time memory)
  - Intel Xeon processors (2.0 to 3.2 GHz), 1 GB memory, running Redhat Enterprise Linux 4 OS, 2.6.9-22.0.1 kernel; LAN interconnect
- Test scenarios
  - Sudden influx of eastbound traffic at western most link
    - Clients 1-10 roll back due to sensor data
    - Clients 11-20 roll back due to change in predictions of clients 1-10
  - Compare ad hoc distributed simulation against replicated simulation experiment of entire network

Steady State, Exit Link

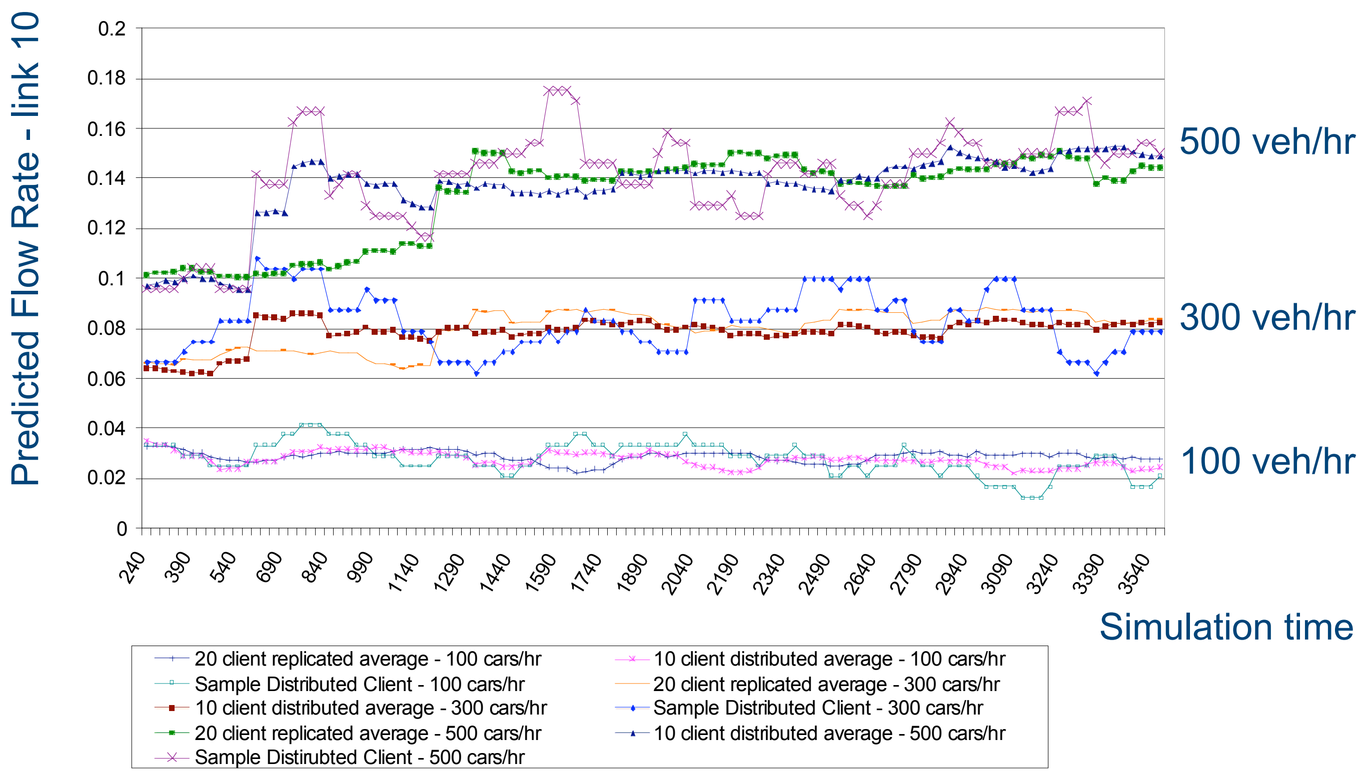

- Constant input rate at edge of network throughout experiment
- Measure flow rate on rightmost link at edge of network
- Compare average (replicated trial), client average, single client

Change in Input Rate

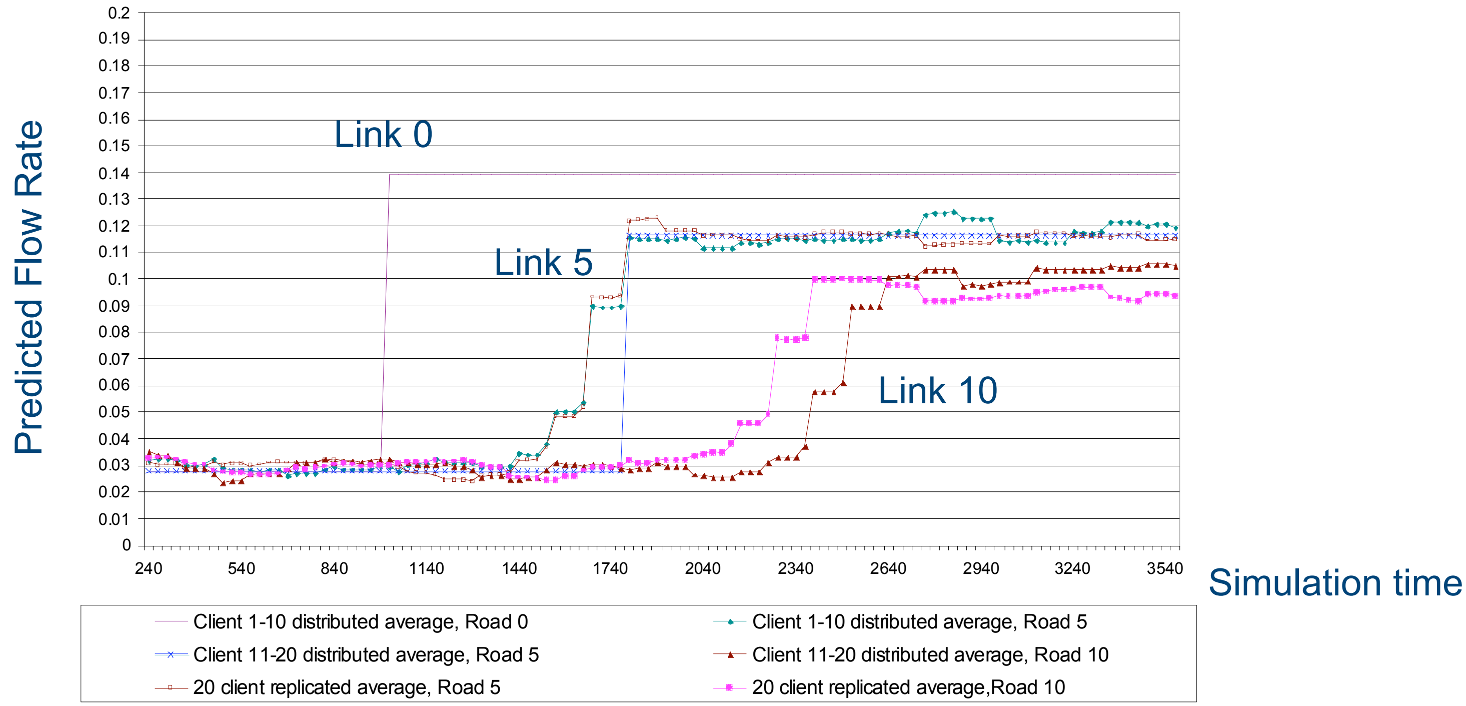

- Initial input rate of 100 veh/hr
- At time 1000, increase to 500 veh/hr
- Clients 11-20 roll back when change occurs
- If the simulators not coupled, clients 11-20 would not predict increase in flow until higher traffic volume reached link 5

Commercial Simulator: VISSIM

- Demonstrate applicability to ad hoc simulations to a commercial simulation tool
- Widely used commercial transportation simulation tool for transportation system analysis
- Discrete, stochastic, time-stepped microscopic traffic simulator
- Rich set of features for modeling traffic control mechanisms, vehicle types, driver types, etc.
- VISSIM version 4.10 used in these experiments
- VISSIM COM interface provides access to objects, methods, properties
- Rollback added using state save / restore capability

Study Area

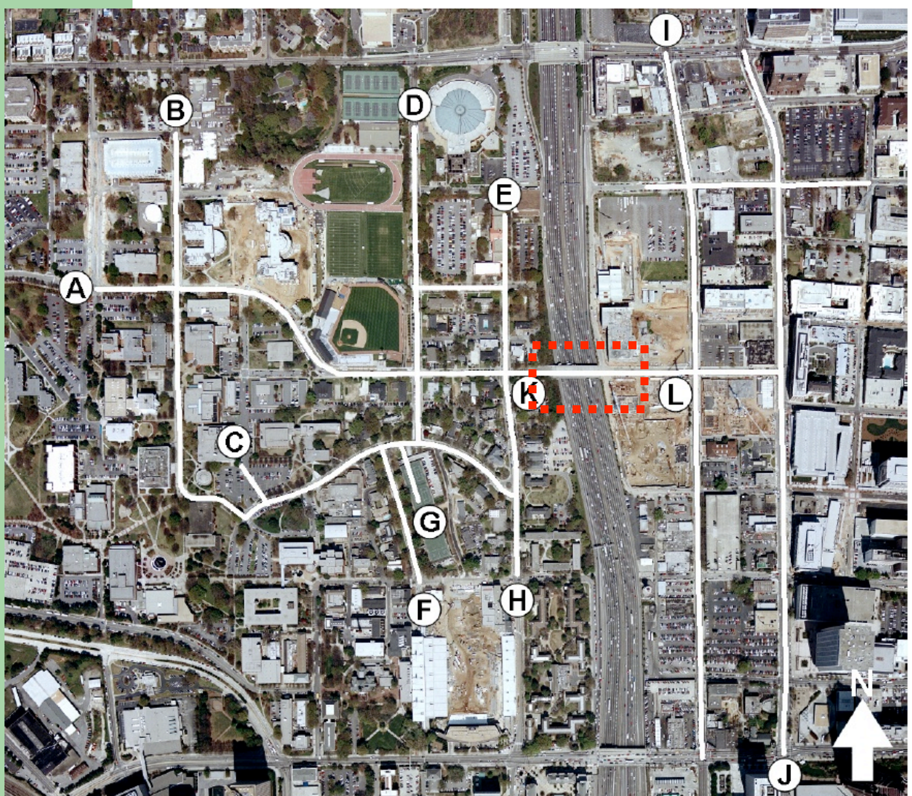

- Scenario: evacuation of Georgia Tech campus
- Normal traffic demand at points A-J
- Traffic at point A increases from 100 to 600 veh/hr 1800 seconds into scenario
- Indicated link is bottleneck (highway overpass)

Study Area

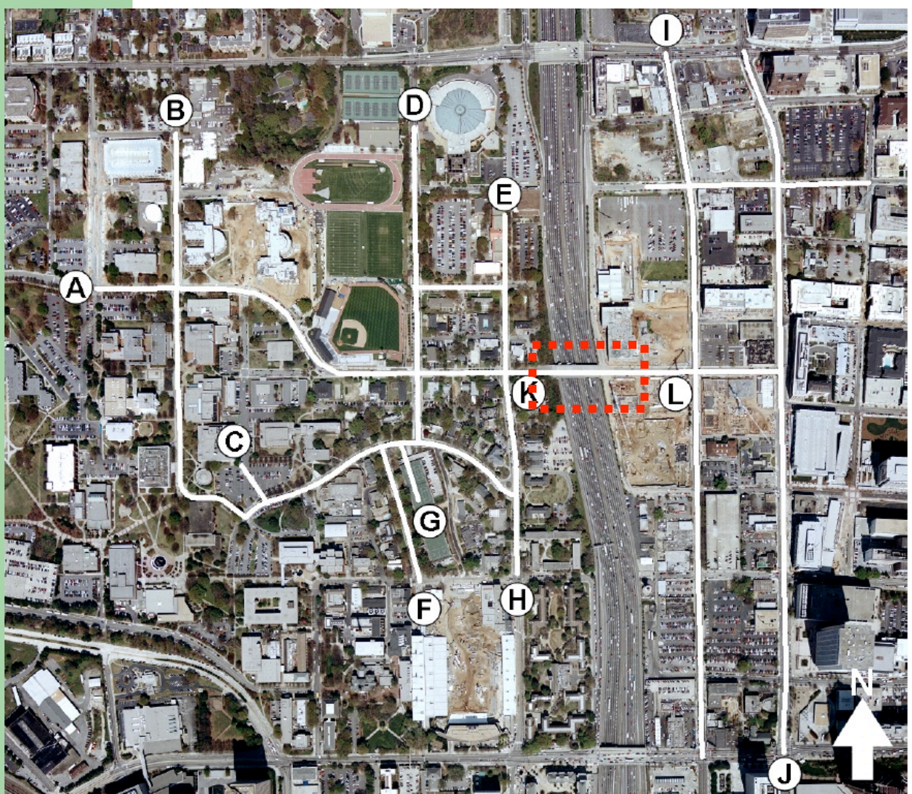

- Scenario: evacuation of Georgia Tech campus
- Normal traffic demand at points A-J
- Traffic at point A increases from 100 to 600 veh/hr 1800 seconds into scenario
- Indicated link is bottleneck (highway overpass)

Traffic Flow (Point A)

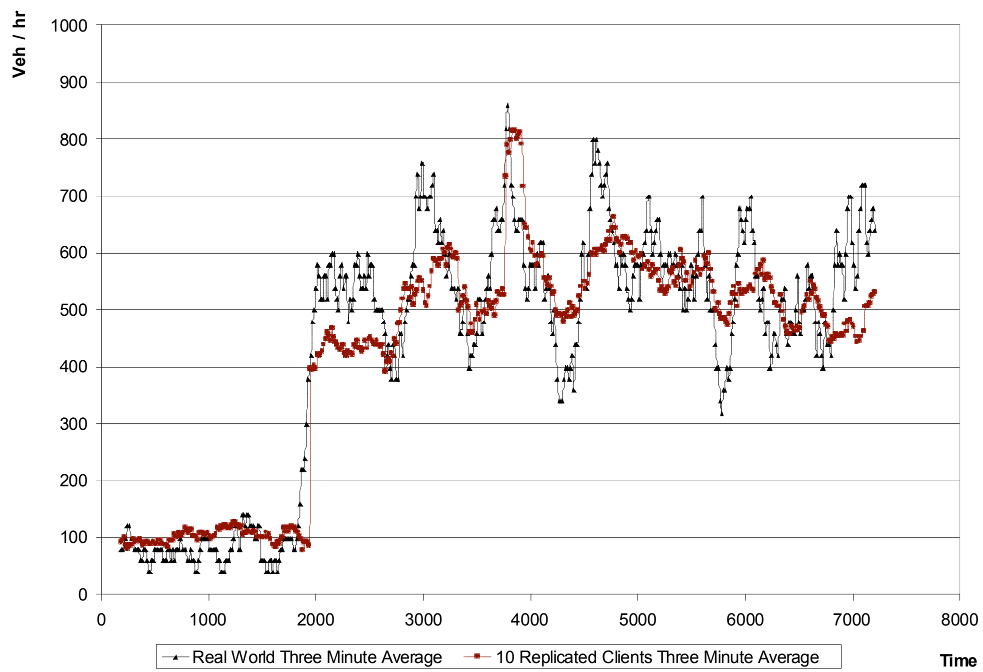

- One VISSIM simulator serves as “real world” generate traffic updates to server
- Ten VISSIM clients (on-line simulators)
- If VISSIM client input rate (point ‘A’) differs from real world by more than threshold, roll back VISSIM client
- VISSIM rolls back, incorporates new “real world” data

Traffic Flow (Overpass)

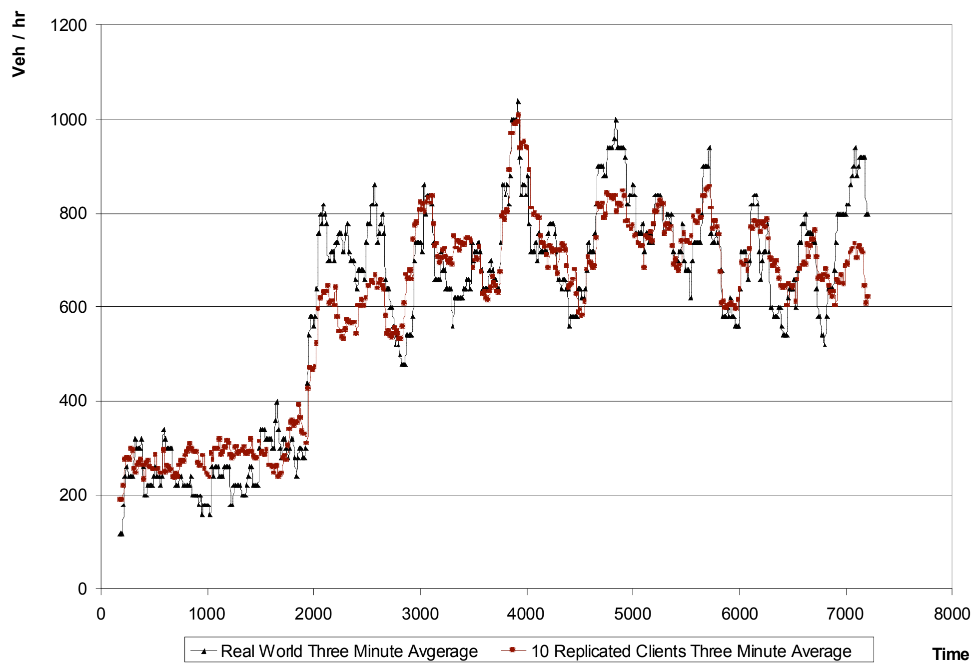

- Predicted flow on overpass link tracks that of “real world”
- VISSIM predictions at bottleneck link track “real world” data

Other Research

- Distributed Simulation Tools
  - Traffic Simulations
  - Wireless Network Simulations
- Vehicle-to-Vehicle Networks
  - Data dissemination algorithms
  - Data propagation analysis
- Network Performance Measurement
- I-75 Corridor Study

Traffic Corridor Study Area

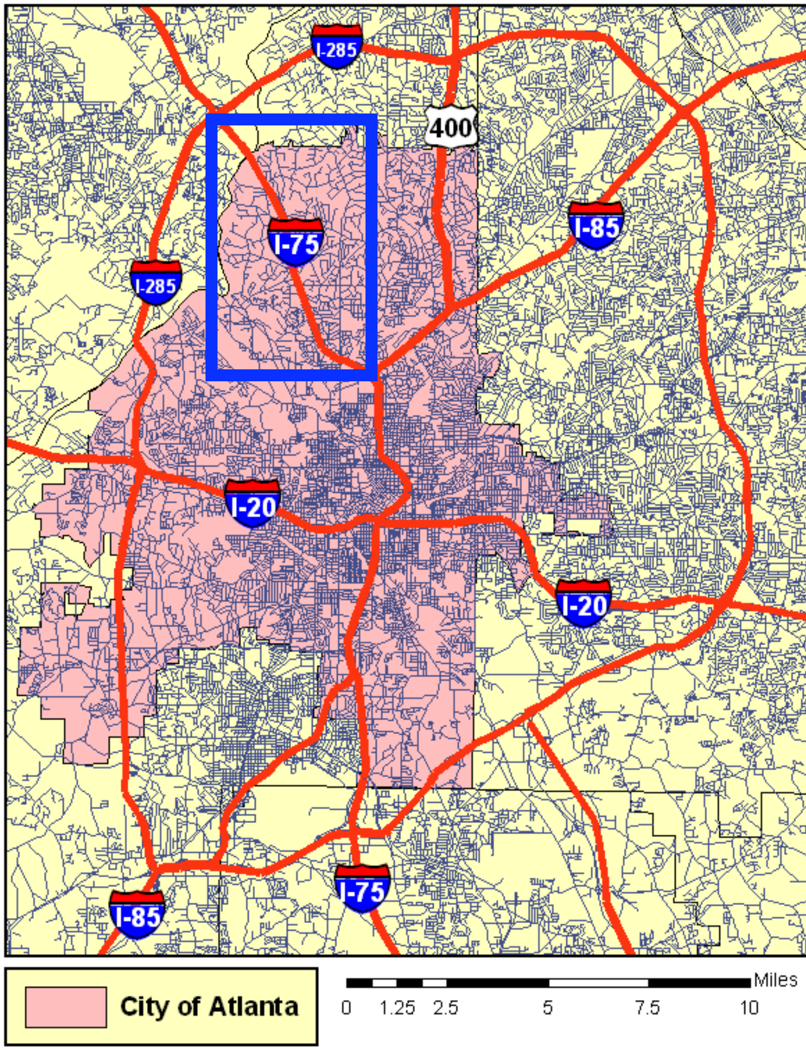

- I-75 and surrounding arterials in NW Atlanta
- 189 nodes (117 arterial, 72 freeway)
- 45 signalized nodes
- 365 one-way links (295 arterial, 70 freeway)
- 101.4 arterial miles
- 16.3 freeway miles (13.6 mainline, 2.7 ramp)

Integrated Distributed Simulations

- Microscopic traffic simulation
- Vehicle-to-vehicle and vehicle-to-infrastructure wireless communication
- Distributed simulation over LANs and WANs

CORSIM

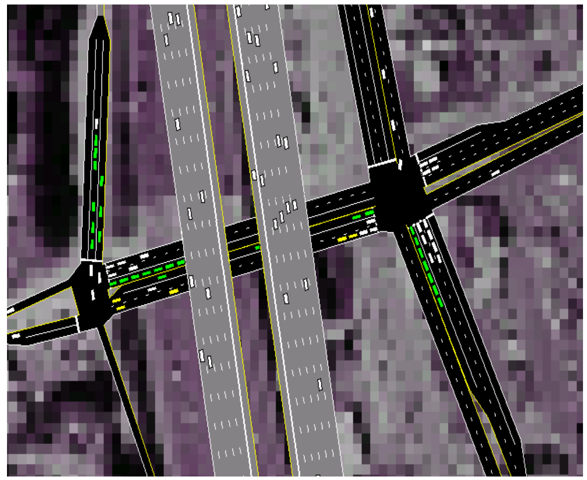

Traffic Simulator

QualNet

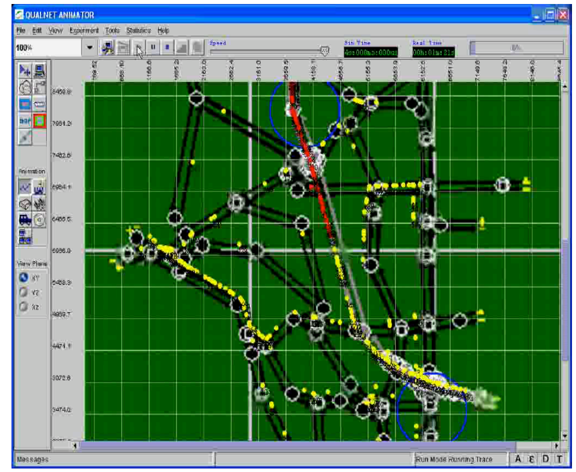

Comm. Simulator

Run Time Infrastructure Software

Federation management | Pub/Sub Communication | Synchronization (Time Management)

LAN/Internet

Traffic Simulation Model

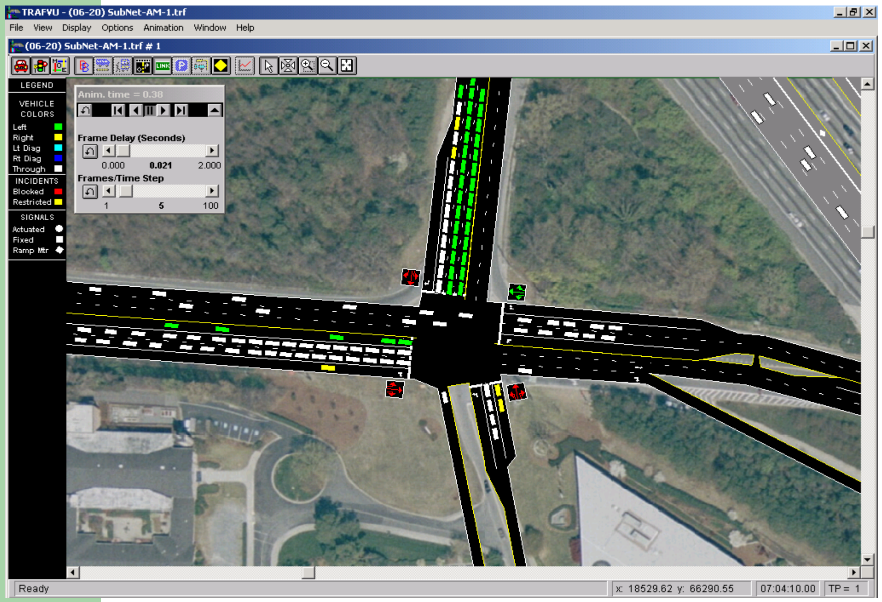

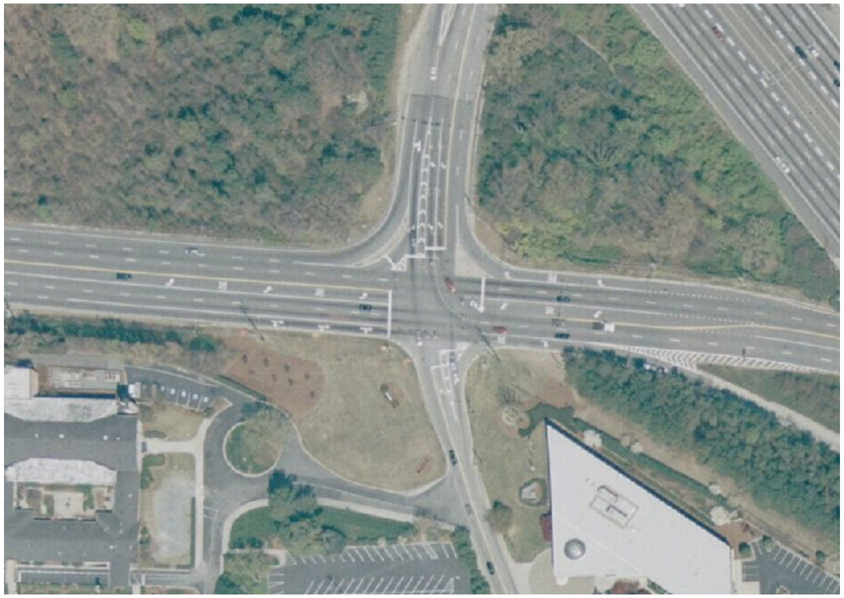

- One-foot resolution United States Geological Survey (USGS) orthoimagery aerial photos used to code lanes, turn bay configurations, and turn bay lengths for each intersection
- Traffic volumes, signal control plans, geometric data, speed limits, etc., obtained from local transportation agencies

Model Calibration & Validation

H. Wu, J. Lee, M. Hunter, R. M. Fujimoto, R. L. Guensler, J. Ko, "Simulated Vehicle-to-Vehicle Message Propagation Efficiency on Atlanta's I-75 Corridor," *Journal of the Transportation Research Board,* 2005.

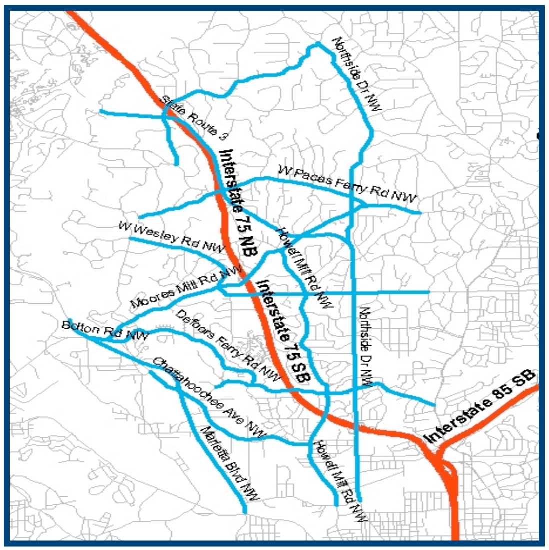

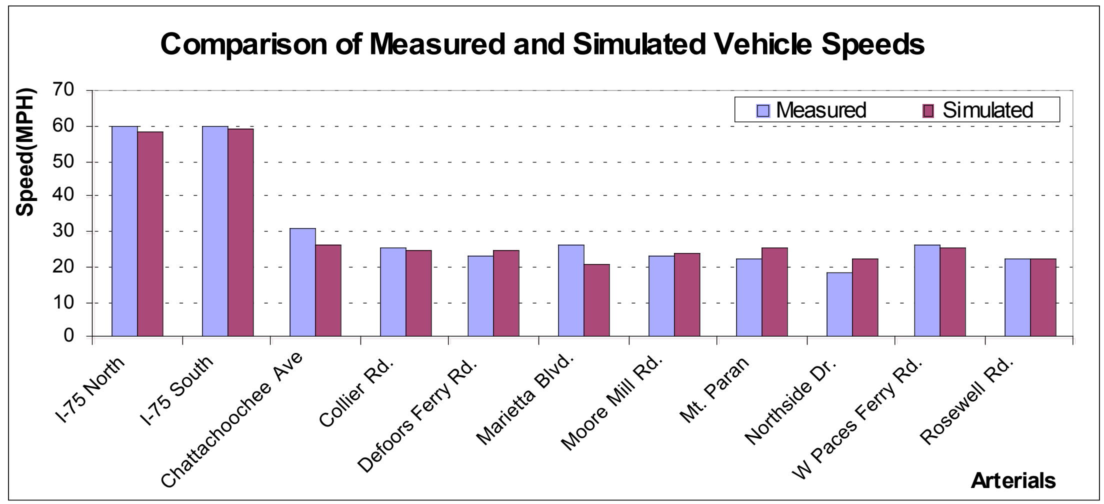

- Anomalous (simulated) delays observed at some locations
  - Field surveys completed at six intersections to calibrate model
- Validation using instrumented vehicle fleet collecting second-by-second speed and acceleration data
  - GPS data from 7 AM to 8 AM peak used
  - 591 weekday highway trips (Feb.-May 2003)
  - 601 weekday highway trips (July-Sept. 2003)

Spatial Propagation Problem

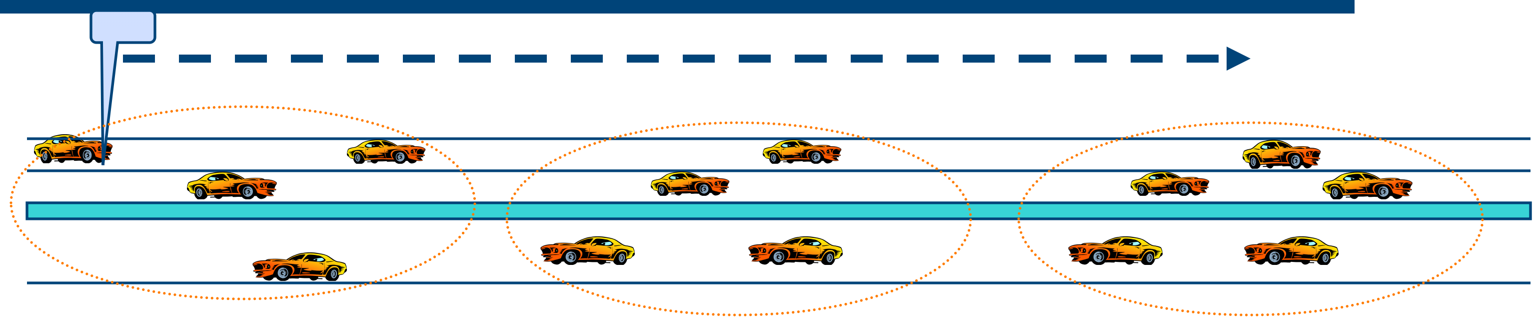

Spatial Propagation Problem:
How fast can information propagate with vehicle forwarding?

Focus on V2V ad hoc networks (802.11) in order to understand the limitations of message forwarding

Observations
- *One dimensional partitioned network*
- *Vehicle movement helps propagate information*

Vehicle Ad Hoc Networks

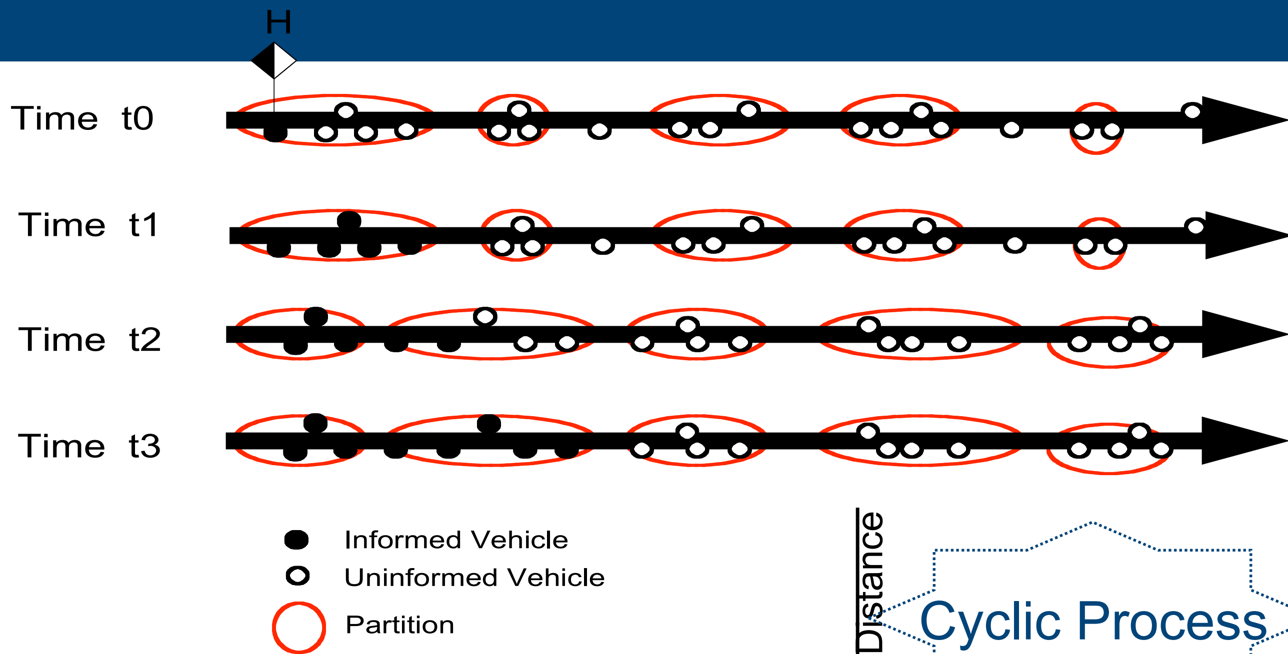

- **Partitioned Network**
- **Forward mode**
  - Message forwarding within a partition
- **Catch-up mode**
  - Vehicle movement allows message propagation between partitions

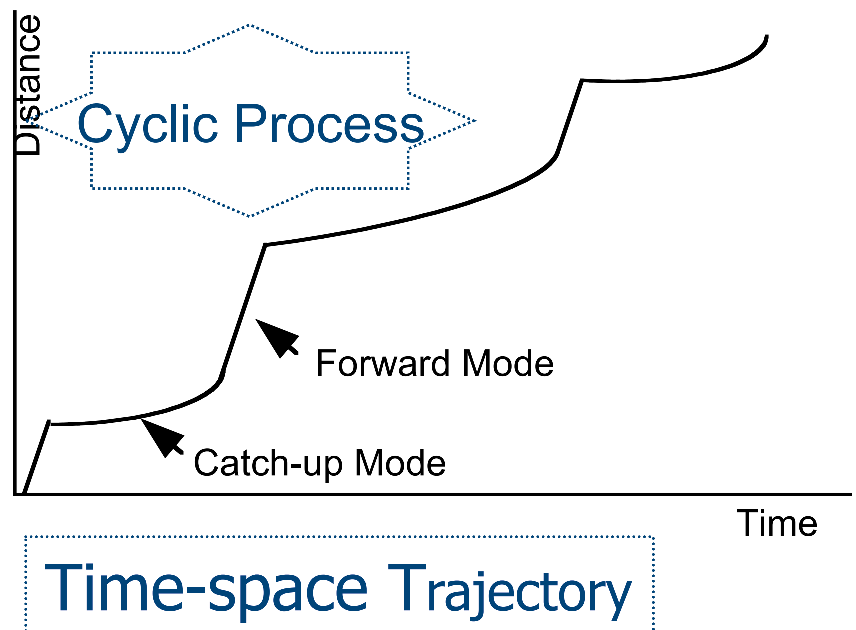

Distributed Simulation Demo
Vehicle Ad Hoc Network

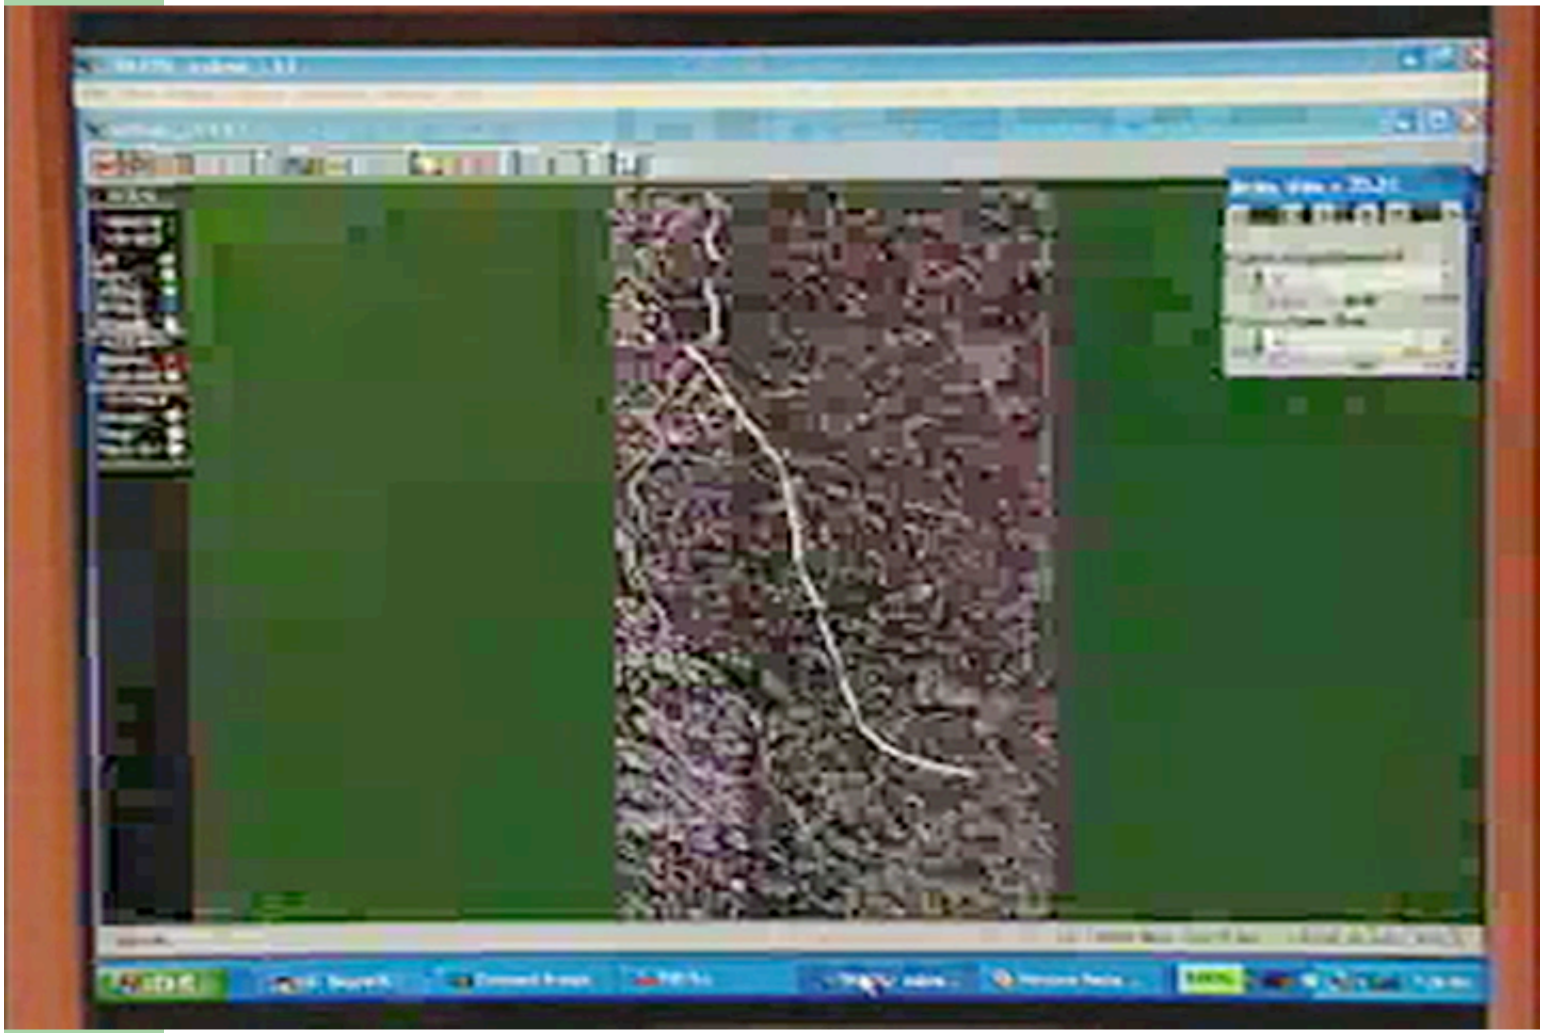

Communication Simulation (QualNet)

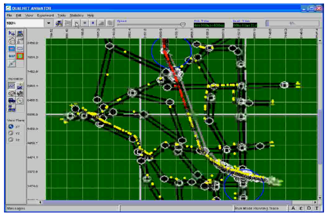

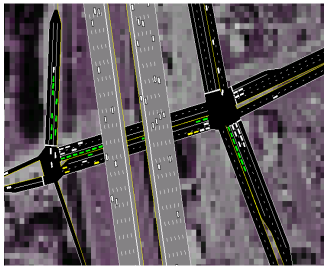

Traffic Simulation (CORSIM)

Data Dissemination using vehicle-to-vehicle communication (802.11)

End-to-End Delay Distribution

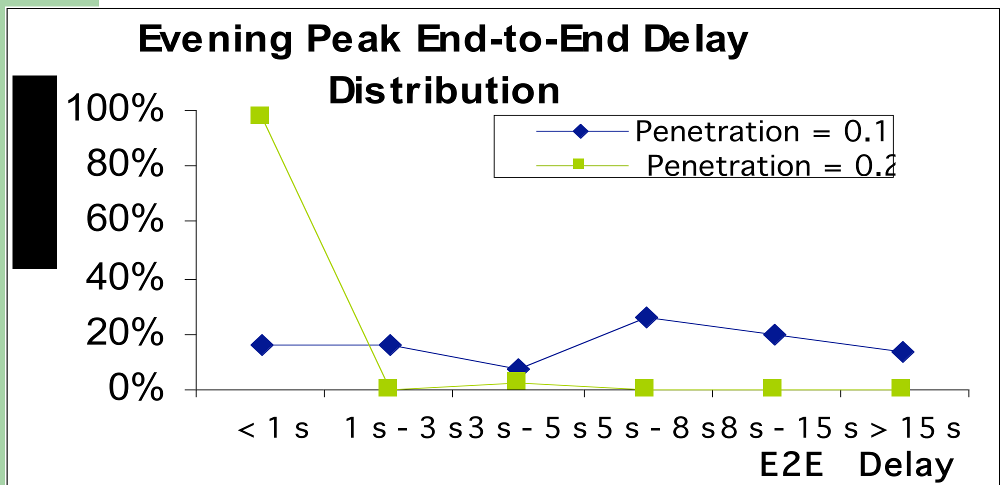

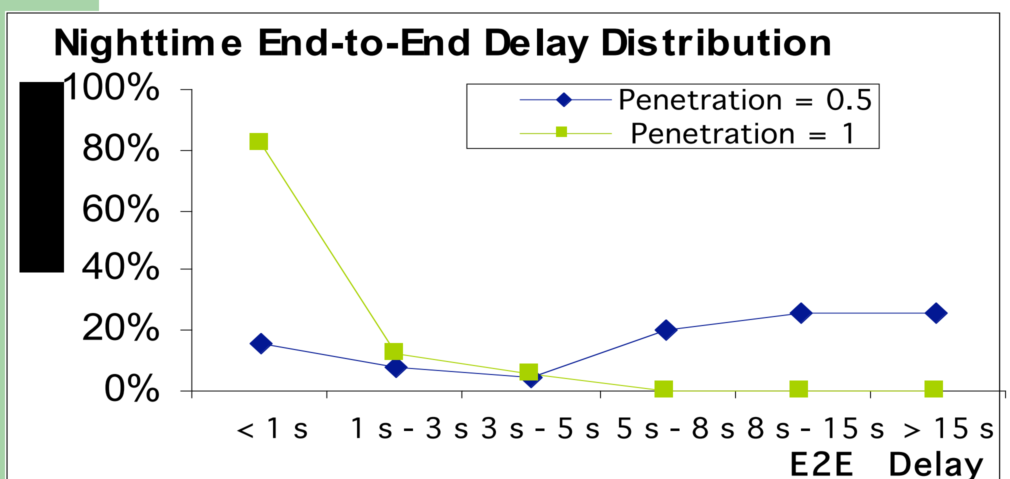

- Delay to propagate message 6 miles along I-75 (southbound)
- Heavy (evening peak) and light (nighttime) traffic
- Penetration ratio: fraction of instrumented vehicles
- Significant fraction of messages experience a large delay

Wireless Communication: Performance Measurement

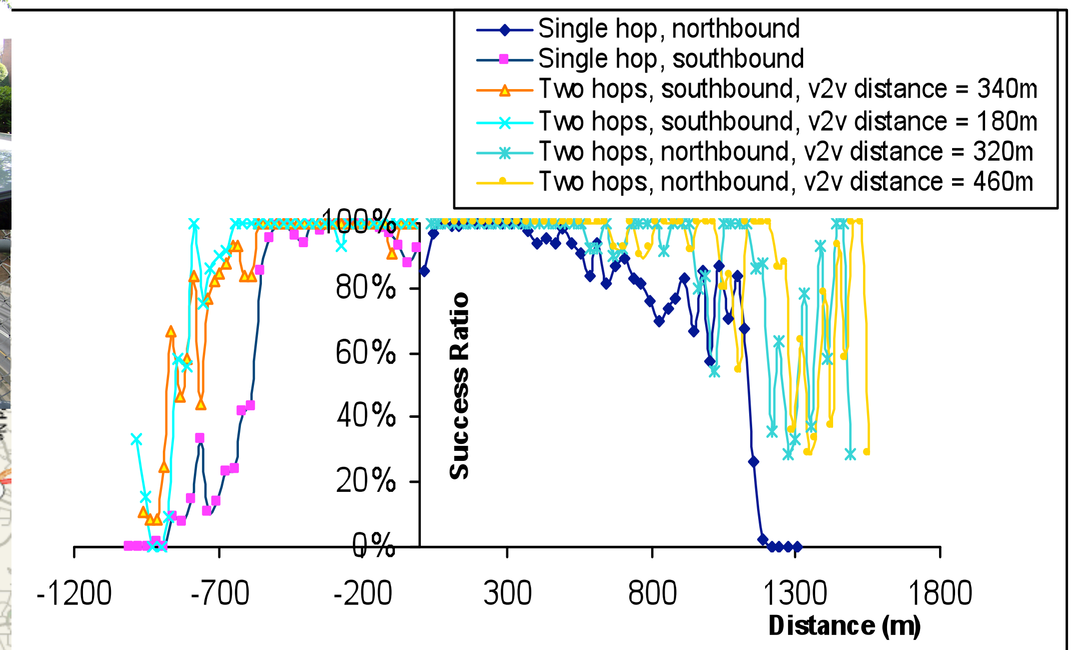

- In-vehicle system
  - Laptop, GPS receiver, 802.11b wireless card, external antenna
- Software
  - Iperf w/ GPS readings; data forwarding module
- Location
  - Northwest sector of Atlanta, GA, along I-75 between Exit 250 and Exit 255
- Un-congested traffic

Conclusions

- Ad hoc approach presents a new class of distributed simulations
  - Networked on-line simulators
  - Constructed in bottom-up fashion
- Combine elements of conventional distributed simulations and replicated trials
  - Optimistic synchronization protocol
  - Multiple independent updates to common portions of system state
- Applicable to other types of systems involving distributed sensing and the need for predictive simulations

Current & Future Research

- Runtime Infrastructure Software
- P2P (server-less) Architectures
- Operation over Unreliable Transport
- Statistical Output Analysis
- Deployment, Field Experiments
- Beyond Transportation System Simulation
  - Parallel Simulation
  - Other Applications

Thank You!