Urban-Scale Sensing for Science

Phillip Sitbon, Nirupama Bulusu, Wu-chi Feng

Portland State University
{*sitbon, nbulusu, wuchi*}*@cs.pdx.edu*

1 Vision

With the continued miniaturization and increased capability of computing devices, the ability to accomplish *urban-scale sensing* is becoming possible from a device perspective. Such large-scale sensing could be enabled by a mixture of both static and dynamic entities capable of capturing a variety of complementary data to enable the understanding of such systems. An example of this includes the evolution of intelligent transportation systems. Currently, many of these systems are enabled via static inductance loops on major highways or on a limited scale by companies that place sensors in vehicles such as buses or delivery trucks. Such *trip planning* ideas are just the starting point of trying to understand large-scale phenomena.

The challenge in enabling such a dynamic infrastructure lies in creating large-scale mobile sensing systems. The goal, of course, is to enable a scientist to gain understanding of time-varying phenomena in a dynamic system. The properties of such a system are that it must be relatively reliable, adaptive to new sensing tasks, and scalable enough to support many moving components.

2 Massively Scalable Mobile Sensing Systems

To support urban-scale sensing, communication capabilities in mobile sensor systems must be massively scalable.

In broad terms, there are two extremes to collecting and reporting data in a mobile system. First, mobile agents can collect data and report to a central location only when a destination is reached. In this case, network efficiency is high because data is never duplicated, but latency between sensing and reporting is also high due to limited connectivity. On the other hand, mobile agents can additionally report data to each other as often as possible. Data duplication is high in this case, thus lowering network efficiency; however, latency is much lower because data is reported sooner. The challenge in protocol design is to achieve the best of both high network efficiency and low reporting latency.

To overcome this challenge, the concept of delay-tolerant networking is a powerful tool. In a delay-tolerant network (DTN), there is no assumption of a complete path between two nodes at any given point in time. The goal of routing in a DTN is *eventual* (often *opportunistic*) delivery of data to its destination. Routing in a DTN is similar to some existing methods of communication, such as e-mail or postal mail, which use intermediate storage and are not highly time-sensitive (but can be). This diverges from the goal of most general mobile ad-hoc network (MANET) routing solutions, which is to immediately establish a complete route; such a route may never exist in a DTN due to its ephemeral nature.

We believe delay-tolerant networks accurately describe intelligent transportation networks. Generally speaking, not much is known about a vehicle's mobility and forming an immediate path for data is difficult. Providing best-effort forwarding prioritized either to reduce latency or increase reliability has the added advantage of low network complexity because there is no demand for establishing an immediate end-to-end route.

2.1 Dynamic Mobile Sensing

Data demands may change depending on parameters such as geographic location or traffic conditions. Effective mobile sensing systems should be able to adapt to such parameters or respond to queries and conditions. Emergency response scenarios, for example, might require a higher-resolution view of traffic data near an emergency location in order to minimize response time. Dynamic sensor reprogramming and *iterative querying*, or "zooming in" on data, is a simple task in static sensor networks, but even state-of-the-art protocols do not support mobility.

Enabling dynamic behavior is critical to the scalability of mobile sensing systems. Without it, data will be too sparse or too dense for the wide variety of requirements placed on such systems.

3 Tools

Understanding the dynamics of mobility and networking at a massive scale is an important step toward effective large-scale deployment of vehicle-to-vehicle communication systems. Gaining such understanding through deployment and testing is prohibitive in terms of cost and evaluation. Detailed simulations, on the other hand, are useful tools to gain such understanding at a low cost. Here, we outline tools for high fidelity simulation of vehicles and wireless networking together.

The cross-discipline challenge of bringing together accurate transportation and wireless network simulations can be viewed as having three dimensions: (i) vehicle mobility, (ii) the wireless medium, and (iii) communication algorithms. Many approaches to this challenge vary the scale and fidelity of (1) and (2). We believe highly accurate representation and understanding of both will lead to improved design and effectiveness of communication algorithms. For these reasons, we are exploring the use of high-fidelity, large-scale vehicular and wireless network simulations and applying delay-tolerant networking concepts to them.

Achieving high accuracy in vehicle mobility models is important when network simulation and analysis needs to be done on a per-vehicle basis. Events such as rush hour traffic jams, missed turns, and lane changes are likely to have a large impact on a vehicle's local wireless communication and the subsequent outcome of the entire simulation. Choosing a purely synthetic vehicle population and road network is not desirable for this reason.

3.1 Vehicle Simulation

For our work, we chose to generate vehicle mobility data with the Transportation Analysis and Simulation System (TRANSIMS), an integrated system of travel forecasting tools[2]. The goal of TRANSIMS is to provide detailed spatial and temporal simulation of travel conditions to facilitate the effective analysis of transportation systems. In 2002, the Federal Highway Administration (FHWA) sponsored an implementation of TRANSIMS for the Portland, Oregon metro area. The result of the project was a collection of data generated from census, survey, and road information sources used to simulate realistic population activities and nearly 5 million vehicle trips throughout an average weekday.

Because TRANSIMS is flexible in terms of reporting results, we were able to tune the simulation output to our needs. Rather than collecting summary data specifically relevant to transportation planning and analysis, we configured the traffic microsimulator to output a snapshot of the location of every vehicle in the system for every second of the simulation. Although this volume of data might seem unmanageable, we were able to import it into a relational database and efficiently index over 2 billion rows of data for an 8-hour period, all on a modestly-powered desktop computer. Our method of storage and indexing provides us with extremely fast access to individual trips, all vehicles and locations for a given time step, and most importantly proximity and relative velocity calculations.

3.2 Network Simulation

Simulating generalized large-scale networks can be challenging with limited computation requirements; however, certain aspects of vehicle simulations can be exploited in order to lower these requirements. Although a day's worth of simulation may contain millions of trips, only vehicles actually traveling at a certain time need to be considered during a network simulator step. After a trip is finished, it can be completely removed from the simulation, thus freeing associated memory and computation. In the context of the Portland study, this amounts to at most hundreds of thousands of vehicles rather than millions. Additionally, the exact mobility of each vehicle can be ignored as long as something is known about the proximity and relative velocity between vehicles at every time step. Through careful database design, we are able to determine such information readily.

For our initial exploration of network simulations, we chose the Georgia Tech Network Simulator (GTNetS)[3]. It suits our needs as a fast and lightweight discrete-event simulator and, with careful modification, we hope it can scale to our required volumes. GTNetS does not directly support removing vehicles from a running simulation, but we were able to modify it to do so without affecting its behavior. We are also working on our own simulator that mimics the behavior of GTNetS specifically for wireless networks in order to maximize scalability and quantify the effects of network simulation fidelity on large-scale network behavior.

Network simulation speed can be negatively impacted by the complexity of the chosen communication algorithm, creating the need for a simple and efficient way to perform V2V communication. We believe delay-tolerant networking algorithms, which are not routing-intensive, will help make networking manageable for both simulations and real deployments.

References

[1] OTREC: Oregon transportation research and education consortium, 2008. `http://otrec.us/`.

[2] LANL. TRANSIMS: Transportation analysis and simulation system, 2008. `http://transims-opensource.net/`.

[3] Riley, G. F. The Georgia Tech network simulator. In *MoMeTools '03: Proceedings of the ACM SIGCOMM workshop on Models, methods and tools for reproducible network research* (Karlsruhe, Germany, Aug. 2003), ACM, pp. 5–12.

Biographies

Phillip Sitbon is a Ph.D. student and Research Assistant at Portland State University, where he studies hybrid sensor systems and vehicular ad-hoc networks. He received his M.S. in Computer Science in 2007 and his B.S. in Mathematics in 2004.

Nirupama Bulusu is an Assistant Professor of Computer Science at Portland State University, where she leads an active research group in sensor networks. She received her PhD from UCLA in 2002, and her B. Tech degree from IIT Madras in 1997, both in computer science. She is a recipient of the NSF CAREER award.

Wu-chi Feng is Professor and Chair of the Computer Science Department at Portland State University. He received his Ph.D. from the University of Michigan in 1996. His research interests include video sensor systems, multi-modal and mobility-based sensing, and video streaming. He received an NSF CAREER award in 1998 for his work in video-on-demand systems.