The Rise of the Mobile Internet: What does it mean for Transportation

Raja Sengupta, Yaser P. Fallah, CEE:Systems, University of California, Berkeley
sengupta@ce.berkeley.edu, yaserpf@berkeley.edu

1.0 Introduction

We are motivated by the rise of the mobile Internet. Smartphone sales reached 32.2 million units worldwide and growth is outpacing the sale of laptops. USDOT is working on multiple fronts to get wireless into cars through its national Vehicle Infrastructure Integration program, its promotion of prototyping programs such as VSCC and CICAS with the automotive industry, and its acquisition of spectrum for DSRC with FCC.

The grand cyber-physical opportunity is sensing on a massive scale for the actuation of control devices. Possibilities today already include rich streams of GPS data on car motion, queues at traffic lights, how people walk, cross streets, cut through alleys, illumination on roadways, precipitation, temperature, road condition, emissions, and so on. Great socio-economic value could be derived by connecting this data to control actuators such as traffic lights, next bus signs, CMS, or cars and buses themselves. In the near term, actuation of cars or buses would be indirect, i.e., via advice, alerts, or warnings to the driver.

Current control systems in road transportation do not have the architecture to leverage these massive streams of data. Nor do the devices in road transportation have the interfaces to utilize such data. The following section explains why.

2.0 Three Fundamental Limitations

We think of current road transportation control systems as closed, as opposed to Internet based systems, which require a more open style. For example, in a traffic signal control system the sensors, actuators and controllers between the two, are owned and operated by the roadway operator, i.e., a department of transportation. We are using the term 'closed' to describe a system building style where an institution integrates the sensors and actuators to deliver the system. In this sense, a car or bus is also built in the closed style. The automotive OEM delivers the system.

By contrast, leveraging the Internet requires a more open style of system building. Systems emerge from the interaction of components over protocols such as http, TCP, IP, web services, etc., representing the Internet. The components may be built by institutions in the closed style. But the "system" itself is delivered by no single corporate body. For example, if the control of traffic lights is to be based on vehicles approaching the light sending GPS data, the sensor will no longer be owned or controlled by the department of transportation. The signal control system encompassing sensor and actuator emerges from the interaction of a GPS source owned by consumers with an actuator owned by the roadway authority.

We believe the massive rivers of sensor data discussed earlier can only be enabled by the mobile Internet. It is the best known way of deriving value from bits contributed by millions of people or cars. Thus the roadway operator needs to build control systems relying on sensors not owned by the operator. Likewise, an in-vehicle warning system may be a combination of sensors built by the OEM moving data to interfaces on the driver's smartPhone built by the phone company. This kind of open-system is the Internet-style and this style is massively scalable. To leverage this massive scale we have to learn to do transportation control in the Internet style. This style comes at a price. The sensor data arrival process would be erratic and its precision highly variable. If motion data is coming from GPS in smartphones, one would expect a wide range of precision. The price of Internet-style scalability has always been QoS.

Thus we see at least three fundamental limitations, challenges, and requirements in today's road transportation control systems:

- Current road transportation systems are not architectured to emerge over the Internet. The challenge is to get beyond closed systems and their centralized architectures. This requires control with open systems, i.e., systems emerging from the interaction through Internet protocols of sensors, actuators, and controllers with different owners.
- Current systems are not designed to control with an erratic data arrival process. We require control with erratic, intermittent data. Warnings to drivers, or the switching of signals are real-time and safety critical.
- They are not designed to control with data of highly variable precision. This requires being able to extract information in the massive, i.e., useful in large sample.

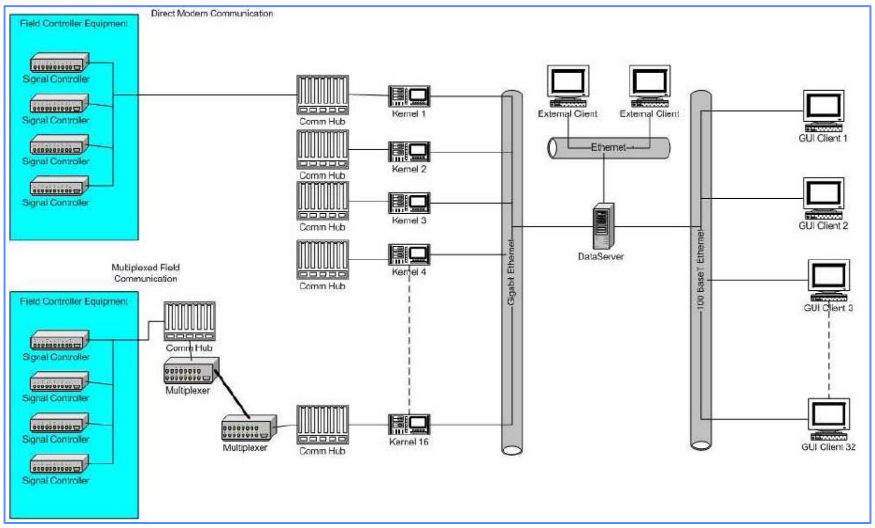

Figure 1: Architecture of LADOT's ATCS

Figure 1 shows the architecture of LADOT's ATCS (Adaptive Traffic Control System). Every 1 second a computing cluster receives sensor data from loop detectors distributed throughout Los Angeles, computes and sends out actuation commands to approximately 4000 signals. This system is an example of the closed style. To make it work, they even own the communication network! This is not the right way to harness the mobile Internet. Likewise the architecture envisaged by the national Vehicle Infrastructure Integration (VII) program for the collection of GPS data from cars, proposes to have all the data flow into two message switches, i.e., one on the west coast and the other on the east coast. From there, the data would fan out to the millions of little control systems that would put it to use. It is hard to imagine all the sensor data from all the smartPhones or cars in the United States, flowing into two national databases!

3.0 Innovations and Abstractions

To utilize this rich emerging universe of data, we need to enable targeted control spaces at multiple spatial scales. When an actuator needs a piece of the big universe of sensor data, there should be an easy way, a control space, to connect the sensors producing the slice with the actuators using it. In general, the set of sensors required to be present in a space is context-sensitive. It depends at least on location and time. Thus the control space and the control system it induces have variable structure. Figure 2 shows you our initial experimentation with this concept. We are creating a website, networked-traveler.org, which one can use to specify information to be put in and information to be taken out. For example, the My Safety tab in the figure allows a pedestrian to specify that she wants to send out a message when crossing the road by checking the box next to “Watch” out for me. Likewise a driver can elect to receive such messages by checking the “Awareness of Nearby Pedestrians”. This configures a control space, i.e., the pedestrian safety control space. This is easy because it is done using a browser at a web server. Now, at many points in time and space, when pedestrian and driver become proximate, in the semantics of the pedestrian safety control space, the space must autonomously instantiate itself to affect real-time information flow between pedestrian and driver. At this stage, we only partially understand how to make the intelligence for autonomous instantiation ubiquitous in the infrastructure.

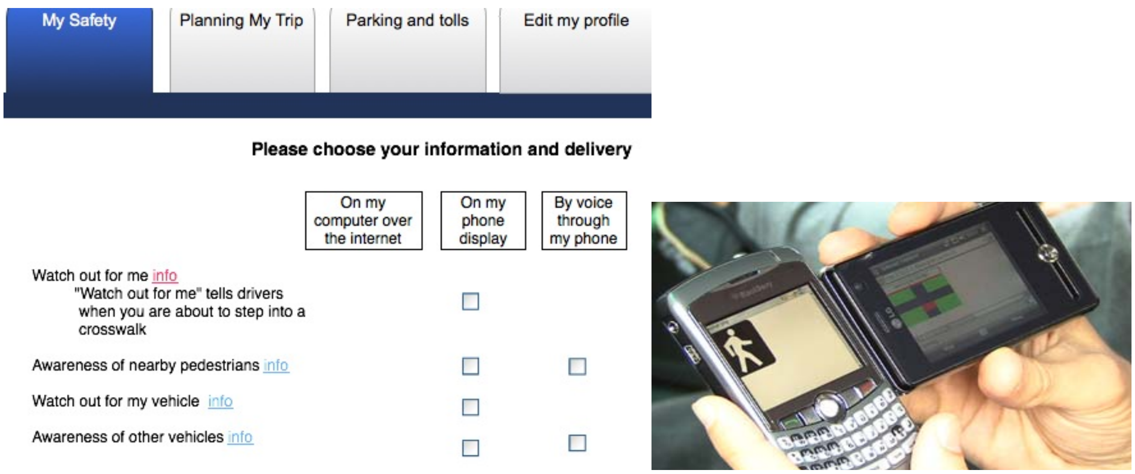

Figure 2: Creating control spaces via networked-traveler.org The picture on the right shows the pedestrian and driver smartPhones.

Finally, there are technical challenges associated with interfacing with the physical world. For example, current traffic lights can only be remotely controlled over a serial interface and modem. This is incompatible with the emergence of open systems. As a first step, we have developed the ATCP protocol enabling the control of traffic lights via XML messages via TCP/UDP ports. This works on the Siemens and Econolite 2070 controllers. They can now interact in control spaces.

Bio: Dr. Raja Sengupta is Associate Professor in the Systems Engineering program within Civil and Environmental Engineering, University of California at Berkeley. He received his Ph.d from the EECS department, University of Michigan at Ann Arbor. His research interests are in vehicular ad-hoc networks and unmanned air vehicles. He was Associate Editor of the IEEE Control Systems magazine, Program Chair of the IEEE Conference on Autonomous Intelligent Networked Systems 2003, Co-General Chair of the first ACM MOBICOM Workshop on Vehicular Ad-hoc Networks 2004 and Co-Program Chair of the second workshop. In 2008 he chaired the First International Symposium on Vehicular Computing Systems.
Dr. Yaser P. Fallah is a Postdoctoral Fellow at the Institute of Transportation Studies, University of California Berkeley. He received his Ph.D. in 2007 from the ECE department, University of British Columbia, Canada. His research interests are in the area of wireless communications and networking, with emphasis on vehicular transportation applications. Dr. Fallah’s is the recipient of several prestigious Canadian national awards and fellowships.