Position Paper: Coping with Integration Challenges in Large Scale Smart Transportation Systems

Authors: Aniruddha Gokhale, William Otte and Abhishek Dubey
Email: {gokhale,wotte,dabhishe}@isis.vanderbilt.edu
Institute: Vanderbilt University, Nashville, TN, USA

Though individual transport systems are becoming smarter with ongoing research in problems related to automatic inter-vehicle collision detection or blind spot detection or self-parking, a number of challenges still remain, especially related to issues of integrating smart transportation systems with smart city infrastructure.

The RFID revolution has given rise to a new distributed systems paradigm called the Internet-of-Things (IoT), which is characterized by the ultra large-scale connectivity among every conceivable object – living (e.g., humans) and non-living (e.g., machines) – that are involved in exchanging and processing valuable information to help solve a multitude of problems of grave societal import that the world faces today. These include problems in the domain of health care, energy, disaster management, and global warming. Additional technologies that are related to IoT, such as Industrial Internet[1] , which also involves M2M communications, exist. We consider all of these related technologies as IoT.

IoT envisions a shift away from traditional stovepiped vertical silos of systems, which focus inwardly within an individual application domain, to a world where significant horizontal interaction among domains will be the norm. It assumes that a large number of objects – stationary and mobile – are fitted with sensors that will emit data, which can then be used to solve one or more problems of interest.

This paradigm is especially suited to transportation systems. For example, consider the problem of traffic congestion in a major city. Traffic congestion is a problem with broad economic and environmental impact. Congestion means that workers do not make it to work on time, deliveries can be significantly delayed and vehicles that are stuck in traffic jams emit far more harmful pollutants into the atmosphere and consume far more fuel than they otherwise would if they were to be able to proceed quickly to their destination.

Now consider that roads and highways today are fitted with a large number of sensors, which, if correlated, could provide quite a lot of insight into current and developing traffic patterns. Traffic cameras, for example, often monitor highways and intersections. Automobiles, for example, have onboard computers that monitor the speed, acceleration,

---

[1] http://www.industrialinternet.com

and braking patterns. Many drivers today carry GPS-enabled and Internet connected smartphones that could be used to relay this information to a municipal traffic monitoring authority. There are a number of useful things that can be done with this data: First, the municipal authority could correlate the data from the various sensors that it has available to detect traffic congestion as it occurs and preemptively take corrective action. Such actions could include routing drivers to alternate routes (taking care to not overload any one particular alternate route). While re-routing would work well for autonomous vehicles, human drivers may require economic incentives to change their behavior, so this data could be used to influence dynamic toll prices for road usage. Moreover, changes in traffic light signaling intervals to optimize for traffic flow could be implemented by the local municipal authority. Finally, this data could be used to detect traffic accidents as they occur and preemptively dispatch both first responders and cleanup crews to ensure that lanes are re-opened as quickly as possible.

A number of complexities stand in the way of this vision, however. First, the hardware and software environment exhibits extreme heterogeneity: there exists a large number of vendors supplying elements of this network: various manufacturers for vehicles, smartphones, and the various sensors (e.g., cameras, magnetometers used to track traffic patterns) and actuators (traffic signaling systems, bulletin boards) in use today. Moreover, as a practical matter it is inconvenient or impossible to upgrade these systems regularly.

This implies a need for a mechanism that all the disparate parts of this system may efficiently use to communicate. Second, the "scope of concern", i.e., which vehicles or sensors a particular municipality may care about, may change dynamically. For example, vehicles may pass through many municipalities while driving from their starting location to their ultimate destination. Additionally, traffic congestion and the resulting mitigation efforts may involve the cooperation of several municipalities. This implies the need for collaborative dynamic service orchestration whereby traffic coordinators may discover vehicles that may be entering their scope of concern and may discover mitigation services provided by surrounding municipalities.

An essential element of the solution includes a standardized data ontology that can be used by all devices in this system to share data with each other. Defining such ontology as part of a standard, however, may be difficult from both a technical and political perspective. First, it is difficult to envision now the types and capabilities of sensors that may be deployed in the future, so the data model would have to be specified in such a way to preserve future flexibility. Unfortunately, such flexibility is difficult to achieve in such a way that preserves the interoperability advantages of a standardized data model. Second, from a political perspective, it may be difficult or impossible to convince all of the interested parties (i.e. manufacturers of cell phones, vehicles, and sensors) to agree on a common data model.

This approach is central, for example, to the portability and interoperability approach used in the Future Airborne Capability Environment (FACE)[2] standard. In FACE, the core

---

[2] https://www.opengroup.us/face/

standard specifies message formats and basic data types, but currently says very little about the *content* of those messages. The content of these messages is expected to be specified by the FACE Shared Data Model (SDM). The SDM is a shared repository of message content structures that vendors are encouraged to use when developing new products. In recognition of the fact that new products often require new data formats, the FACE also specifies governance procedures that allow interested parties to extend the vocabulary supported by the SDM.

Addressing the need for collaborative dynamic service orchestration is far more challenging. In the system outlined above, location-aware matching between data sources (vehicles, sensors, cameras) and data subscribers (municipal traffic centers, actuators used for mitigation) is essential. Not only must such matching be performed for the current physical location of the data source or subscriber, but must be performed based on predicted location in the future. To the best of our knowledge, existing data dissemination approaches do not address this physicality constraint. For example, DCPS Discovery in the OMG Data Distribution Service[3] allows one to silo participants into *domains* (which completely segregates discovery), subdivide participants into *partitions*, and finally associate particular data types to *topics*. While such a system could conceivably be used to implement physically constrained discovery (given a high level protocol for assigning domains, partitions, and topics based on physical constraints), the fully decentralized and peer-to-peer nature of this protocol would likely not scale well to the hundreds of thousands of participants in this system.

Next generation smart transportation infrastructure holds the promise to significantly increase the efficiency of severely congested roadways in today's large cities. Such infrastructure could coalesce the observations of tens- to hundreds of thousands of sensors spread throughout a city's roadways and highways – both statically deployed (such as cameras and electronic roadside signs) and dynamically located (sensors and GPS units on vehicles themselves) to preemptively detect and proactively mitigate roadway congestion and traffic jams. A significant challenge, in addition to determining how to analyze this data and react to it is the integration of the various sensors and actuators present in the system: extreme heterogeneity in the hardware motivates the need to standardize an ontology for data communications, while the extremely dynamic nature of the system requires that we develop a cooperative and dynamic service discovery and configuration service that is aware of its cyber-physical nature.

---

[3] http://portals.omg.org/dds/