CRITICAL NATIONAL NEED IDEA

Achieving ubiquity, timeliness, accuracy, and reliability in nextgen traffic information systems

University of California, Berkeley
California Center for Innovative Transportation

Alexandre Bayen (PI)[1], Roberto Horowitz[2], Raja Sengupta[3], Alexander Skabardonis[4], Pravin Varaiya[5], Tom West[6]

[1] Department of Civil and Environmental Engineering,
University of California, Berkeley, CA, 94720-1710
Tel: (510) 642-2468, Fax: (510) 643-5264
Email: bayen@berkeley.edu

[2] Department of Mechanical Engineering,
University of California, Berkeley, CA, 94720-1740
Tel: (510) 642-4675, Fax: (510) 643-5264
Email: horowitz@me.berkeley.edu

[3] Department of Civil and Environmental Engineering,
University of California, Berkeley, CA, 94720-1710
Tel: (510) 642-9540, Fax: (510) 643-8919
Email: sengupta@ce.berkeley.edu

[4] Department of Civil and Environmental Engineering,
University of California, Berkeley, CA, 94720-1710
Tel: (510) 642-9166, Fax: (510) 642-1246
Email: skabardonis@ce.berkeley.edu

[5] Department of Electrical Engineering and Computer Science,
University of California, Berkeley, CA, 94720
Tel: (510) 642-5270, Fax: (510) 642-7815
Email: varaiya@eecs.Berkeley.edu

[6] California Center for Innovative Transportation,
University of California, Berkeley, CA 94720-3830
Tel: 510-642-4522, Fax: 510-642-0910
Email: tomwest@calccit.org

Key words: Traffic monitoring, Cyber-physical systems, Sensor networks, Large-scale systems, Data assimilation

Achieving ubiquity, timeliness, accuracy, and reliability in nextgen traffic information systems

Alexandre Bayen (PI), Roberto Horowitz, Raja Sengupta, Alexander Skabardonis, Pravin Varaiya, Tom West

1 Background

1.1 Fighting economic losses due to inefficiency in the transportation network

An area of critical national need in the US. Traffic congestion in the US alone causes a 78 billion drain on its economy annually, a figure which has doubled in the period 1997-2007. In total, drivers lose 4.2 billion hours and waste 2.9 billion of fuel gallons per year because of traffic congestion (Schrank and Lomax, 2007). The current recession has only marginally improved the state of traffic congestion due to reduced activities in some regions of the nation, but long term projections by the *Energy Information Administration* (EIA) at the US Department of Energy in 2009 show a flattening - but not a significant decline - in this trend.

Prerequisites for transformational results. It is unlikely that traditional physically-centered mitigation strategies, such as building more highways, widening existing roads, and expanding or starting new transit or rail routes, will be successful by themselves. These approaches are simply not sustainable in the current economic and environmental climate. Rather, innovative paradigms are needed to marshal breakthroughs for operations of the transportation network, to transform the manner by which the traffic management issues are addressed. One key component which is missing to address this is an exhaustive *traffic information system*, which can be used at global scale. With partial information only, it is very unlikely that the management of the system could be improved significantly, and that the adoption of new transportation paradigms by the public could shift from its current state.

A societal challenge. At the heart of this problem are significant barriers which prevent operations and system use from being addressed at a national scale. In alignment with TIP's mission, we propose to tackle the *technical* high-risk high-reward issues associated with these challenges, which we believe to be centered around information gathering, integration, and distribution. To better operate the system, agencies need a *integrated* view of the full transportation system. To participate in a change in the way the transportation system is used, the public needs *location based* information, available in real time, to make more efficient decisions, which might include increased use of the transit system. With the emergence of the *mobile internet* on cellular devices and the rapid proliferation of location based services, it is the right time to launch initiatives which will result in paradigm shifts in the way the public evolves in the transportation network. In particular, we are specifically interested in finding ways in which to give incentives to the public to use the transit system. This involves several challenges, which include (*i*) providing users with location based transit information, including on their mobile phones, (*ii*) interfacing this information with existing traffic information to promote multimodal transportation, (*iii*) increasing the quality of the information available to them for transit, so they view transit as a viable option.

1.2 Challenges to be overcome for the creation of an ubiquitous traffic information and operation system

The main challenge in using traffic information is that no single source of data possesses the four following features: ubiquity, timeliness, accuracy, and reliability. These features, however, are a prerequisite of efficient use of information for future operation systems. These properties are also required to ensure the value of the information being provided to the public.

Ubiquity. A major enabling technology to offer motorists alternatives to highways is the possibility of using the secondary (arterial) network. This ability relies on information about the state of the arterial network, which in most of the US cities is not available. Agencies are thus unable to provide recommendations through channels like the radio, changeable message signs, internet, etc. Motorists are therefore reluctant to leave the highways without certainty of finding uncongested alternate routes. Because of the cost of dedicated infrastructure, equipping the secondary network with monitoring infrastructure is not an economically feasible option. Ubiquity also has a temporal meaning; high availability is a necessary feature of a real world system.

Timeliness. The timescales at which traffic congestion evolves are rapid enough that traffic information which is not received in a timely manner could prove to be useless, and even detrimental, to operations and individual choices in the transportation network. One reason why information currently gathered and available for operations is not always timely is the different layers of filtering, processing, aggregation and broadcast the data goes through. In the absence of integrated systems, data gathered from sensors travels from proprietary system to proprietary system, aging a few minutes in each step, which reduces its value when transmitted to the user. Currently measured delays in the process commonly exceed 30 minutes on some traffic information systems, which makes the data simply obsolete.

Accuracy. Accuracy of collected data is critical to proper operations of the transportation system. Accuracy refers to the pointwise error between the system's estimate of congestion and the actual congestion. Misplacement of an accident on the network because of inaccuracy in the measurements can lead to wrong estimates of traffic, which in turn lead to inefficient operations, sometimes with dramatic consequences for users (for example, a single wrong choice in routing through the Bay Area can lead to 200% in commute time difference, based on the bridge used). Misplacement of the edge of extent of congestion can lead to significant differences in the estimate for *total traveled time* (TTT) because of saturation of the mainline of freeways when off ramps would have provided a better alternative.

Reliability. One of the main issues in adoption of transit systems, or even in use of travel information systems for more efficient planning, is the reliability of the information. In addition to the accuracy of the estimate, one also needs a confidence interval for the estimate. Ideally, one would want to be able to use the data to make statements of the type "travel time using this mode of transportation along this route is within 10% of this value with 90% confidence". This type of question is critical. For example, schedule adherence in the San Francisco area is estimated to be 70% (San Francisco Municipal Transportation Authority, 2009). For the past two years, official *Metropolitan Transportation Authority* (MTA) numbers for New York show a system reliability of around 80% for subways and 66% for buses. From a user's perspective, incorrect information from the static transit schedule database can lead to a commute plan that appears feasible, yet is impossible due to missed connections. Similarly, low quality highway traffic information systems are not giving the traveling public incentives to use them.

Unless these four challenges are overcome, the creation of efficient transportation information systems for agency operations and users is problematic. Yet, without such information systems, the dramatic impacts of transportation on economic efficiency cannot be faced by increased operation efficiency. Finally, it does not seem possible to induce paradigm changes in the way the public travels without interfacing traffic information systems with transit information systems properly. In order to provide alternate choices to established patterns, it is key not only to integrate transit more deeply into the transportation network, but also to enable the public to make informed decisions about its use, including multimodal travel.

2 Proposed work

2.1 Scientific and technology contributions

Executive summary. This team has accumulated significant expertise in the last decade on the development of single real-time monitoring and control systems, which span almost all aspects of the surface transportation network (highways, secondary road networks, transit, urban, suburban, rural). We propose the creation of a large scale integration of all the efforts of the team into a single traffic information system which can handle all available data in a unified manner. The future of operations in urban and suburban areas cannot continue to be performed as disconnected partial systems. The problems of ubiquity, timeliness, accuracy, and reliability can only be solved if global solutions are envisioned. But adding systems to each other will not provide the solution. Integration of systems has to rely on the proper handling of the data, and the appropriate architecture of the system.

Our team can make a difference in the development of theory, algorithms, software, and a prototype system to integrate multiple sources of data, which is a pathway to the four features outlined above. The problems which the team proposes to tackle encompass a range of disciplines wider than commonly accepted in the peer-review processes, due to the very practical application of the proposed research. Therefore, we believe that the TIP is the appropriate program for this research. This section presents our envisioned contributions. In the subsequent sections of this document, we will explain how the previous work of the team can be leveraged to provide a compelling implementation of the concepts presented herein.

Data processing and fusion. No single organization has the ability to collect all the data necessary for an ubiquitous traffic information system with the features outlined above. The data, by its nature, comes from various sources, which include public agencies' dedicated infrastructure, industry entities in the business of data collection, academic institutions which have developed monitoring infrastructures, and consumer data available from participatory sensing. If integrated properly, this data has the potential of achieving the features outlined earlier. This integration faces several challenges.
- *Data reconciliation.* Because of measurement noise, the data obtained from various sources (and often even from a single source) contradicts itself. Proper algorithms need to be developed to handle this problem.
- *Data fusion.* Data describes different attributes (speed, density, travel time, etc.), it comes in different formats, at different time intervals, with different levels of noise and accuracy. Efficient algorithms need to be developed to integrate these different data sources in ways which can lead to efficient handling of the data.

- *Fault detection.* Numerous sensors are faulty and fail on a daily basis. For example, in California, 40% of the inductive loops have daily malfunctions. Data processing methods need to be developed to isolate faulty, but realistic looking, data (i.e. sensors which provide plausible but incorrect information) from quality data, which is a major problem. This is in addition to the standard problems of noisy data (which needs to be filtered).

Data assimilation. Massive data collection by itself is not the answer to the problems above. In fact, massive data collection creates problems of its own. The currently available traffic data (static sensors, probe sensors, video, radar, etc.) is too large to be stored or even used in real time. For example, many of the probe vehicles currently capable of sending data in the transportation network can transmit their position and speed every second. This data might not only be impossible to process in real time, but it is not even clear if this level of granularity can help at a global scale. Creating a system capable of handling all the collected data means creating a system capable of extracting the proper subset of the data of use for the application. The proper approach to create a system such as the one proposed in this document is to create the proper *data representation*, which only saves the features of the data useful for the application of interest. This is achieved as follows.

- *Spatially and temporally aware sampling.* For each of the data sources available in practice, proper sampling strategies need to be developed. Determining the proper rate at which the data should be kept provides a viable data processing framework. Similar considerations need to be tackled with all data sources, so that integration of multiple data sources becomes possible. Spatial as well as temporal sampling strategies will be investigated.
- *Efficient modeling of traffic.* In order to incorporate measurements into the system, it is necessary to be able to efficiently model the way the transportation network evolves. The physics of traffic flow on highways requires a specific sampling approach; arterials require another paradigm because of the higher complexity of their traffic features (stop signs, traffic lights, intersections, etc.). Each domain requires the development of an appropriate model to enable the proper definition of the required sampling (previous item). The modeling techniques will require contributions in traffic flow theory and statistical modeling.
- *Data assimilation.* The process of integrating measurements into a model (statistical or flow based) is called *data assimilation* or *inverse modeling.* It is key to the success of any estimation engine (for *nowcast* or real-time assessment of traffic). Because of the sparsity of measurements, of noise issues, and of sensor malfunctions, among others, measurements are not enough to provide an accurate picture of traffic: mathematical models are required to specifically address the absence of data at given locations of the network, and proper assessments of the value of the data where it is measured but seems erroneous. Algorithms will be derived which incorporate methods from *variational data assimilation, statistical learning* and *optimal filtering.*

Distributed modeling. Increased operational performance is made possible by the development of efficient simulation tools, capable of handling traffic at the scale of numerous counties and metropolitan areas in parallel, which can be challenging for dense urban areas such as Los Angeles. The key to better operations of the network relies on the development of real-time models (i.e. models which can run on computers faster than the actual phenomena they model), so that real-time actuation becomes possible. For this, the following developments are required.

- *Abstraction of traffic for multiscale applications.* In order to run active control schemes on an urban area such as Los Angeles (for example, actuation of metering lights, or rerouting some traffic in the urban network), it is necessary to develop models which are tractable. In particular, microsimulation models which model the behavior of cars individually are not a sustainable option for operations at this scale. For this, traffic needs to be modeled at an aggregate scale precise enough to handle small size traffic features (shocks, queues, etc.), but large enough so that

handling a complete city (for example San Francisco) is not a problem. The development of these models and their implementation will be investigated.

- *Adaptation of existing models to super-large scale environments.* Specific computational tools need to be developed to handle super-large problems. To be able to do real time computation for the entire Bay Area, for example, parallel computing algorithms must be developed to be able to decompose the geographic transportation network in computational sub-modules which can all be handled separately. The team will develop the proper mathematical tools to parallelize all computational tasks to be implemented as part of this project.
- *Incorporation of dynamic maps in models.* Numerous modeling items are required for the accurate performance of such models, in particular the knowledge of road geometry (number of lanes, expressways), of the presence of traffic infrastructure (stops, traffic lights), the knowledge of operational features (duration of the traffic cycles, changeable traffic rules, etc.).

Active distributed control. Improvements in operational efficiency, and in user utilization of the network, can be achieved by large scale control schemes, which rely on the ability to use the aforementioned models efficiently. Active control (or analysis) of the transportation network can be achieved if the corresponding models are cast in a format which can be handled by the appropriate mathematical tools. We propose the following tasks.

- *Control theoretic model of highway and arterial flows.* Based on earlier traffic flow models, we will adapt existing models to integrate the capabilities of controlling and optimizing traffic flow in large scale urban areas. For this, the following frameworks will be used: *partial differential equations*, *cell transmission models*, and *dynamical systems*, to answer each of the subproblems specifically (highways, arterials, urban, suburban, etc.). Particular attention will be given to the modeling of actuation, i.e. the various ways in which traffic operations can be run. Capabilities to be developed as part of this effort include (*i*) ramp metering, (*ii*) dynamic speed limits, (*iii*) actuation on the traffic light cycles, (*iv*) flow diversion (based on information available on changeable message signs, (*v*) transit advisories.
- *Optimal control and optimization formulation of improved traffic operations.* Leveraging the models developed earlier, we will incorporate the traffic information into the estimation mechanism underlying these control schemes (i.e. the traffic estimation engine), and we will conduct studies on how to handle super-large scale traffic in an integrated manner. This involves contributions in the field of *distributed parameter systems, model predictive control, nonlinear control, adjoint based optimization* and *hybrid systems*.
- *Improved real-time transit planners.* The transit information available from the various sources in this work will be interfaced with the transit information available will be interfaced with the traffic information, to provide the public with increased availability of choices for travel in the network. Optimization based solutions will be computed to provide the public with intermodal transport information, which is the only way to create information based incentives for alternate modes. Particular emphasis will be put on *multimodal transport, casual carpooling,* and *location based services for the routing of pedestrians in the transit network.*

Cyberphysical infrastructure systems architecture development. In addition to the theoretical and technological challenges tackled by the plan above, a system such as the one proposed in this plan requires a specific contribution to the field of *cyberphysical infrastructure systems.* Large scale cyberphysical infrastructure systems are systems which integrate computational and physical processes, using embedded computers and networks to monitor and control the physical processes, usually with feedback loops where physical processes affect computations and vice versa (Lee 2008). Using the motion of humans and goods in the built environment, the notion of cyberphysical systems can be extended to infrastructure such as transportation networks, water distribution networks, the power grid, instrumented bridges, etc. Large scale cyber-physical infrastructure systems thus

encompass physical processes related to the infrastructure, such as the modeling of motion of people ("physical") and the corresponding information gathering, communication and computing system ("cyber"). This proposal deals with the specific case of mobile sensing in large scale systems, which are spatially distributed. Cyberphysical systems require the creation of a specific architecture which this team proposes to design as part of this project.

- *Modular database.* As appeared with the scope of data of interest for this work, the core of a cyberphysical system of this nature is the design of the appropriate database capable of handling massive amounts of properly sampled data (i.e. not the raw data, but the processed data). Designing this database so that it can run in real time will be preliminary to the implementation of the system of interest.
- *Real-time i/o feeds.* Key to the implementation of the proposed system is the set of input/output feeds leading to the database. Since the sources of data are varied, one needs to design the architecture of the database in a way that it can handle the streams of data without altering the performance of the database.
- *Filtering infrastructure.* While the data arrives raw to the system, it must be filtered before being archived in the database, and therefore processes have to run downstream from the input feeds before the data can be added to the database. These processes must run in such a way that the amounts of data arriving into the computational system do not alter the performance of the database.
- *Real-time control architecture.* Because the real time estimation and control processes running on the system will make extensive use of the database, it needs to be designed in a way that it can handle both real-time feeds, and massive queries from the computational processes.

While these architectural features are not challenging for a single system, the main issue faced in solving the four problems above is that all these components integrate subsystems which were not initially designed to work together. All the systems which we propose to integrate were developed independently, and did not include specifications to be integrated to other systems by live feeds. For example, it is not obvious how the corresponding database could be queried simultaneously, and how the feeds could be integrated, without the proper architectural design.

2.2 Expected outcomes and capabilities

The expected outcomes of the proposed work include the following items.

Extensive analysis of available traffic information sources. As part of the work on the data, the team will assemble a portfolio of data sources which can be used for potential breakthroughs in ubiquitous traffic monitoring systems. The challenge is not in exhaustive use of the data sources (statics sensors, mobile sensors, video, radar, toll tag readers, busses, fleets, taxis, etc.), but in the proper use of the data. The team will come up with conclusions on how to select the data to use, since the challenge is not in integrating all the data (which is too big), but on the proper selection of the portion of the data which should be used appropriately.

Real-time monitoring and estimation system. The team will construct a prototype system to demonstrate the capabilities of integration of the data into a single engine capable of assembling the streams into traffic models. Existing systems, such as Google Traffic, Bay Area 511, and Traffic.com, are in their infancy; they only incorporate some of the data, and are faced with the challenges of developing efficient models to handle the data. We will produce a prototype system capable of handling all the data, and produce meaningful information out of it.

Case study and demonstration for real-time operations. We will demonstrate the system with a case study in California covering half of the state (the choice between the Bay Area and the Los Angeles area is still open), developing a prototype which could serve as an example for the rest of the nation. The prototype will be used to demonstrate the real-time control and monitoring capabilities to be used for active control of cities. Through our partners in California, we will demonstrate the use of this system by incorporating them with the real time operations of the various agencies we are already working with.

2.3 Path to achieving the goals

Prerequisites to achieving these goals: The goals of this program are extremely ambitious. Very few teams in the US have accumulated enough knowledge about large scale monitoring systems for transportation to be able to achieve the goals we are proposing here. The envisioned system cannot be created without previous experience and accomplishments in the field. The system *has to leverage existing infrastructure* by its very definition. UC Berkeley in the last decade has pursued all components of this work separately, and is now in a unique leadership position to demonstrate how the *next step* in traffic information systems can greatly benefit from the integration of these components. This section summarizes the achievements of each of the PIs of the project, which gives the team the proper credibility to achieve the proposed goals. Each of the PIs has been involved with the creation and the operation of *at least one testbed of significant size* to be integrated in the present system. Most of the corresponding achievements have been covered in the media, and provided significant breakthroughs in the their respective areas. For the first time in history, technological developments will enable this team to integrate these tools and platforms in a system which, if combined, could provide *ubiquity, accuracy, timeliness, and reliability* to the full transportation network.

- *Dedicated wired sensing data sources (lead: Skabardonis).* Numerous Departments of Transportation have developed their own dedicated traffic monitoring infrastructure systems, which generally mainly consists of loop detectors, radars, and video cameras. In California, systems such as the *Performance Measurement Systems* (PeMS) system have already enabled the integration of all loop detectors available for highway sensing in a data aggregation and analysis tool with real time feed capabilities. A large portion of the expertise in the PeMS development resides at UC Berkeley, which historically led the effort in a partnership with the California DOT. Most of the DOTs in the US have video feeds, potentially accessible directly, or through Traffic.com (part of Navteq), which is now a division of Nokia, a prime partner of Bayen's effort.
- *Dedicated wireless sensing data sources (lead: Varaiya).* Breakthroughs in wireless technology have led to the creation of new generations of wireless sensor systems (for example the *Sensys Systems* sensors), which will augment the pavement inductive loop detector infrastructure. This technology is progressively penetrating the transportation network, and provides a cost effective data source solution for congested arterials where accurate knowledge of the flows is needed for more efficient operations. Radar data is also available from equipment developed by Navteq in urban and suburban areas around 27 cities in the US, in a partnership with the US DOT.
- *Probe data (lead: Bayen).* With the emergence of GPS, and the rapid penetration of this technology in the cell phone fleet, cellular phones are progressively evolving into a source of traffic data with the potential of spanning locations which are not equipped with dedicated traffic monitoring infrastructure. With alr eady more than 5,000 downloads, *Mobile Millennium*, which is a field operational test resulting from a partnership between UC Berkeley, the California DOT, the US DOT, Nokia and Navteq, aims to integrate data from tens of thousands of users in California using free software running on GPS equipped smartphones. In addition, numerous dedicated

fleets have GPS equipment available for traffic monitoring (taxis, delivery fleets such as FedEx, etc.). This source of probe data is one way to obtain traffic data at locations where no dedicated infrastructure is currently deployed to monitor traffic.

- *Transit data (lead: Sengupta).* Most transit agencies in the US are now equipping their fleets with GPS, to improve published transit planners' travel times, which are very often inaccurate because of the severe impact of congestion on the schedule (static schedules are usually not accurate and not reliable). In addition to the added benefit provided to the users (i.e. the knowledge of estimated arrival time), this data also provides significant route based information, which can be used for traffic estimation and to integrate knowledge of scheduled stops into traffic sensing, provided the proper processing of the data is done. The *Networked Traveler* project has built an architecture capable of integrating this data at a regional scale, based on publicly available databases (such as NextBus.com) into a transit planner engine which could also be used for traffic monitoring. This data is dynamic (real-time), not static.
- *Regional integration for active control (lead: Horowitz).* The main challenge of integrating such various sources of data is the capabilities of merging the data into an integrated model capable of handling traffic flow estimations (and potentially control) at the regional scale. In the case of California, the relevant scale for such integration is at the size of Los Angeles – San Diego, or the entire Bay Area. The California DOT – UC Berkeley project *Tools for Operation Planning* (TOPL) integrates a data assimilation engine with a simulation engine capable of modeling the transportation network at such scales, and capable of using it for active control (i.e. answering questions such as "how to create active control schemes to reroute a given percentage of the Los Angeles traffic in order to alleviate choking at given bottleneck congestion points").

California: a case study hosted at CCIT. The *California Center for Innovative Transportation* (CCIT) will host the work proposed by this effort. Led by Tom West, a former executive at the California Department of Transportation, CCIT's mission is to provide the proper institutional framework for the deployment of new technology for transportation systems. This team proposes to integrate existing systems into a new large scale environment, which falls directly under the mission of CCIT. CCIT already has extensive experience in the rapid deployment of new technology, in particular with the recent example of *Mobile Millennium* (2008). CCIT will assist institutionally to achieve the goals outlined above. California will serve as a case study, since the different team members have already significant contributions and systems working in California, and have strong institutional links with the state. The proposed system will provide sources of data which can be used by the various participating agencies in an integrated manner. Just for California, based on previous work, and combining all the above sources of data available to the team, the proposed system will have access to sources which include *billions* of data points every day, but will only archive and use the *millions* which are determined to be useful and effective for the proposed system to start operating.

2.4   Role of Government in cyberinfrastructure development

Achieving the proper institutional framework. Because of the variety of the data sources involved, and because of the number of institutional partners which can contribute to this work, this very ambitious goal needs to be achieved under a *Federal* partnership. Previous successes include the *Safe Trip 21* project, under the umbrella of the US DOT, which is a partnership between the California DOT, UC Berkeley, Nokia and Navteq. One of the achievements of *Safe Trip 21* is the creation of *Mobile Millennium*, which was mentioned earlier. *Mobile Millennium*'s success was partly due to the fact that a *Federal* agency was able to assemble government, academic and

industry partners together under the same goal. For the present project, this aspect is even more prevalent, since the variety of data sources and institutions is even bigger. Government can help by federating the different players. The involvement of TIP will thus greatly help with leveraging millions of dollars of public and private funding, and with leveraging private / public partnerships that are already in place. This is critical at a time when separated and disconnected systems are limited in the answers they can provide for global problems.

Other benefits to Government: disaster relief, evacuation and unanticipated events. While the proposed work is specifically focused on the creation of a new traffic information system of an unprecedented scale for operations, it also provides the technological framework for what could appear as a side application, but is also critical to the Nation's infrastructure management. The proposed data collection, data processing and control engine which we will build encompasses all the capabilities desired to handle unanticipated events.

- *Natural disasters.* Major disasters such as hurricanes or earthquakes cause partial destruction of critical infrastructure. The variety of sources available to our proposed system mean that it will still be reconfigurable. We believe that this system will possess such levels of redundancy in the data sources and streams that it will be the ultimate response tool in the case of major disruptions. While an academic team can only construct a prototype, it is clear that with the proper sources of funding, an operational tool of this nature could be taken over by several branches in Government, not only DOTs, but also FEMA, DOE, DOD, etc.

- *Management of special events.* Urban and suburban environments can be significantly disturbed by special events. The most common are major sporting events, which release tens of thousands of additional vehicles into the transportation network in less than one hour, creating massive disruptions of traffic. The most severe are evacuations (as seen with hurricanes). The modeling and control engines developed as part of this work can be used to handle this type of situation in an integrated and more optimal manner.

- *Planning and design for increased robustness.* The above two problems raise the question of increased robustness of the transportation network. Analysis of this network (through simulations) can lead to more efficient and robust redesign of the system. The proposed achievements will thus also serve as a tool for improved planning.

3   References

1. E. Lee. Cyber physical systems: Design challenges. In *International Symposium on Object / Component / Service-Oriented Real-Time Distributed Computing (ISORC)*, Orlando, FL, May 2008.

2. Schrank, D., Lomax, T. (2007) The 2007 *Urban Mobility Report*. Texas Transportation Institute, The Texas A&M University System, September.

3. San Francisco Municipal Transportation Authority, *Service standards report*. Technical Report Q3 FY08, July 2008.

4. *Mobile Millennium:* http://traffic.berkeley.edu.

5. *TOPL:* http://path.berkeley.edu/topl/software.html.

6. *PeMS:* https://pems.eecs.berkeley.edu/.

7. *Networked Traveller:* http://www.viicalifornia.org/.