Title
Preparing the Way for Vehicle-Infrastructure Integration

Author
Shladover, Steven E.

Publication Date
2005-11-01

CALIFORNIA PATH PROGRAM
INSTITUTE OF TRANSPORTATION STUDIES
UNIVERSITY OF CALIFORNIA, BERKELEY

Preparing the Way for Vehicle-Infrastructure Integration

Steven E. Shladover

California PATH Research Report
UCB-ITS-PRR-2005-31

This work was performed as part of the California PATH Program of the University of California, in cooperation with the State of California Business, Transportation, and Housing Agency, Department of Transportation, and the United States Department of Transportation, Federal Highway Administration.

Final Report for Task Order 5213

November 2005

ISSN 1055-1425

Preparing the Way for Vehicle- Infrastructure Integration

Steven E. Shladover
California PATH Program
Institute of Transportation Studies
University of California, Berkeley

Abstract

Vehicle-Infrastructure Integration (VII) is one of the most important new national programs within the field of ITS, but until now it has been treated at the federal level as a deployment-oriented activity, without significant research elements. This report identifies some key research issues that need to be investigated in support of VII, in order to ensure that the VII concept and designs are as efficient and effective as possible and in order to produce the essential knowledge base that public agencies throughout the country can rely upon to make well-informed decisions about participating in the deployment of VII.

1. **Introduction**

The essence of an Intelligent Transportation System (ITS) is the ability to combine information gathered from both vehicle and infrastructure sources, to fuse it efficiently, and to use it to improve the operational capabilities of both vehicle and infrastructure systems. This concept of Vehicle-Infrastructure Integration (VII) has recently attracted strong interest among the automotive industry and state DOTs, represented by AASHTO, together with the USDOT, who are all now discussing how to work together to implement it. In order for the VII deployment to be effective, it will be necessary to address a wide variety of research issues which are not currently receiving much attention. This report identifies the most important of these issues, both technical and institutional, and outlines how those issues can be addressed in future research that could be conducted under state and/or federal sponsorship. This represents a rare opportunity to define a substantial research agenda that can support fundamental improvements in the operation of the transportation system, and that can leverage future state research funds with substantial federal resources as well.

The national VII effort has been focused on deployment issues but has not yet addressed the many research challenges that need attention. There are both technical and institutional issues that need to be studied and resolved before VII implementation can proceed, and they are intimately linked to each other because the technical system architecture and the deployment business model need to be compatible with each other. If California is interested in playing a prominent role in the national VII program and making sure that VII is designed in ways that meet California's transportation needs, it should seize the initiative in defining the research agenda that will enable VII. This report suggests the basic elements of that research agenda, to prepare the way for subsequent research to be conducted at both the state and national levels. Because of the prominence of VII in USDOT thinking, this should become an excellent opportunity to leverage an up-front state investment with more substantial federal resources in the coming years.

Until now, ITS has not realized its full potential, in large part because it has generally been separated into disjoint elements rather than being truly integrated. The large majority of ITS projects until now have been treated as EITHER infrastructure OR vehicle-based, or they have been segregated by mode: automobile OR bus OR truck. Very little has been done to integrate the vehicle and infrastructure elements or the different modal elements, yet this is the real advantage that ITS can offer to the transportation system. With wireless communication and computational technologies advancing rapidly now, data can be collected from any of these elements, and can be synthesized into useful information that can help all of these elements operate more efficiently and safely. However, in order for this to be realized, a variety of technical and institutional issues need to be resolved.

There are obvious philosophical differences between the infrastructure and vehicle stakeholder communities, but there are also likely to be significant differences of opinion and priorities within these communities. A solid research foundation needs to be established so that these differences can be resolved on the basis of the best available understanding of the issues.

If effective vehicle-infrastructure integration can be achieved, there are many advantages for the transportation system, such as:

- continuous monitoring of the entire roadway infrastructure for problems ranging from pavement condition to visibility, weather, incidents, obstacles in the road, etc.;
- advance warnings to drivers of hazards beyond their line of sight;
- instant notification of emergency responders;
- continuous and ubiquitous traffic condition information (travel times, speeds, incidents, etc.);
- accurate and comprehensive origin-destination travel data for planning purposes;
- safer operation of critical roadway network nodes such as intersections and merge points;
- most cost-effective collection of information on the operation of the transportation system, optimizing the choice of infrastructure and vehicle-based sensing.

California has several natural advantages that should make it possible for it to play a leading role on the national scene in the development of VII:

- largest state, with most congested urban regions, most in need of ITS relief;
- home of Silicon Valley innovations, with automotive industry leaders positioned here to gain maximum benefits from the latest developments;

- technologically sophisticated population, inclined to be early adopters of innovations;
- Caltrans' heritage as the most advanced and forward-looking state DOT;
- University of California PATH Program as the first and most advanced ITS research program in the U.S., with expertise in both vehicle and infrastructure issues;
- site of ITS World Congress in November 2005, providing a platform to show the world California's leadership in ITS and to attract industrial participation.

There are opportunities for major contributions to new knowledge in this area, primarily because relatively little has been done until now. While the general concept of using vehicles as data probes in transportation has been considered for some time, including in the early stages of the ITS program, nobody has yet done the type of in-depth research needed to support the business case (or illuminate the value proposition) for major investments by the vehicle and infrastructure industries. Some of the most exciting opportunities for discovery are associated with these aspects of VII:

- Vehicle probe data provide an unprecedented level of detailed information about transportation operations – how can that best be used?
- The data from vehicle probes can be fused at a variety of different levels for different purposes. Much needs to be learned about the best methods for fusing that data and about the most effective architecture for allocating responsibilities for data fusion.
- The integration between vehicles and infrastructure needs to be accomplished across wireless data communication links. The new 5.9 GHz DSRC service will be one obviously important element in this, but it will not be available everywhere, so much understanding needs to be gained about how it can be complemented or augmented with other wireless services such as cellular radio and satellite communication systems.
- The overall architecture for VII needs to be defined for compatible operations throughout the country, on both the vehicle and infrastructure sides. In many ways this is a more daunting assignment than the definition of the existing National ITS Architecture, because the VII architecture needs to be specified in greater depth and detail in order to meet its goal of full national interoperability.
- A new business model for transportation operations needs to be defined in order to finance the deployment and operation of the VII infrastructure elements. This model will have to be based on allocations of costs that are properly aligned with the benefits gained by the vehicle and infrastructure operation stakeholders, because each will need to make its own economic case for investment. This is likely to need more sophisticated assessments of transportation operations benefits (time savings, travel time reliability, safety, energy consumption, emissions, quality of life, etc.) than have been available heretofore.
- The technical and institutional aspects of VII are so intimately coupled that multi-disciplinary research approaches will be needed to address the major issues. In particular, the information system architecture and business model for deployment and operation are inextricably coupled with each other, so viability needs to be established in parallel on both dimensions, because a serious problem on either side would preclude success.

These issues make VII a remarkably rich and challenging topic for research. Considering the short time frame in which the automotive industry would like to be able to start VII deployment, an early start on these issues is even more important.

2. Key Open Research Issues and Recommended Approach to Each

For purposes of the discussion here, research is defined in broad terms, so that it includes any endeavors to expand knowledge and understanding of how VII can be better designed, implemented and applied (addressing technical as well as non-technical issues).

Inputs to this identification of research issues have come from the TRB-sponsored national workshop on research needs for vehicle-infrastructure cooperation in Troy, MI during the summer of 2004, from the discussions at the VII national meeting in San Francisco in February 2005 and the ITS America Annual Meeting in Phoenix in May 2005 and from the information posted on the VII program web site, especially the report addressing the system architecture issues. These have provided a limited view of what is likely to be done at the national level, and implicitly (by omission) what will still need to be done in parallel with the national effort.

California cannot sponsor or perform all of the research needed to support national deployment of VII, but it can use its participation in the national program as a medium for exposing the other states and the U.S. DOT to research needs. Hence, this identification of research needs is not restricted to those that can or should be met entirely within California, but is intended to be national in scope.

Because of the breadth of the subject matter, the definitions of research needs are clustered in broad programmatic groupings rather than being defined at the level of individual projects. The two most urgent categories of needs involve the definition of methods for efficiently sorting and aggregating the vast quantities of data that will be provided by probe vehicles and the development of methods for public agencies to estimate their costs and benefits associated with VII deployment and operation so that they can make the necessary informed deployment decisions. Other research needs involve development of the rest of the knowledge base needed to support VII development and deployment in areas such as wireless communications, backhaul system architecture, development of new transportation system management approaches that can capitalize on the enhanced data available from VII, development of economic (business) models that can effectively match benefits and costs with the relevant stakeholders, and identification of criteria for prioritizing locations for installation of the VII infrastructure elements (transceivers).

These clusters of research issues are addressed in the subsequent sub-sections below.

1. Development of efficient data sorting and aggregation approaches
2. Development of public sector cost and benefit estimates
3. Definition of backhaul system architecture
4. Development and standardization of needed wireless communication capabilities

5. Development of new transportation system management approaches building on VII data
6. Definition of economic models to facilitate deployment
7. Definition of criteria for prioritizing infrastructure installations

2.1 Development of efficient data sorting and aggregation approaches

The quantity of data that will be communicated from vehicles to the roadside in a mature VII deployment will be truly staggering. This is fundamental to the promise of VII, but is also its greatest technical vulnerability. If the data flow is managed intelligently, VII has the potential to produce major benefits to the transportation system. However, if the data flow is managed poorly, it will not be possible to turn it into useful information. The VII initiative as currently conceived does not appear to be addressing the management of this data flow.

The first step in addressing the data flow challenge is to identify how the data could be aggregated. Data aggregation is an essential function because of the sheer volume of data that needs to be considered. There appears to be some basic agreement that the communications between vehicle on-board units (OBUs) and the local roadside equipment installations (RSEs) should occur at an update interval as short as 100 ms in order to support some of the most time-critical safety services. The minimum message set that has been identified for vehicle-vehicle communications in the ongoing SAE standardization process contains approximately 50 bytes of data, and this probably represents a lower bound for the volume of data that would need to be exchanged with an RSE. This means that each vehicle would be transmitting something in the range of 500 bytes of data per second.

There are approximately 200 million registered vehicles in the U.S., but not all of them are on the road at the same time. Assuming that in the peak travel hours of the day only 25% of these vehicles are on the road, this would still represent 50 million vehicles. When VII is fully implemented, a substantial fraction of these (particularly those operating in high-density urban and suburban areas) should be within range of an RSE. If half of these vehicles were transmitting their 500 bytes of data per second, this would represent a total of 12.5 gigabytes per second. Out of that torrent of data, each application would need to extract the small subset it needs. However, it does not make sense for every application developer or owner to have to independently filter out all the data it does not need in order to access the tiny subset that it does need.

There are many levels at which the VII data could be aggregated, and each of these has significant architectural implications for VII and the users of the data. The examples suggested below are based on the assumption that has been stated by VII Initiative participants that the full build-out of VII would include approximately 500,000 RSEs and that all road vehicles in the U.S. would be equipped with OBUs.

The number of vehicles that would be communicating with an individual RSE could vary widely, depending on the local density of traffic. In lightly traveled locations, there may be no vehicles within range of an RSE for extended periods of time. In the most heavily traveled urban areas, at large intersections or on crowded freeways, as many as 500 to 600 vehicles could be within range of a single RSE at one time. At the rate of 500 bytes per vehicle per second, this represents a total of 25 to 30 Kbytes/second, which should be readily manageable at the local RSE level.

However, the total quantity of data generated regionally or nationally cannot be processed efficiently if it is *only* defined in terms of individual vehicle measurements, without intermediate layers of aggregation, as described in the current VII national architecture. Aggregation in this context means that many, and probably most, measurements obtained from vehicles would be combined *at a local level* to produce *composite* estimates of the operating conditions of sections of roadway (e.g., speed, surface condition), or neighborhoods (e.g., local weather conditions). The measurements of individual vehicle performance that are needed for commercial applications could still be accessible at the national level. It is particularly urgent that a layered architectural and data aggregation approach be developed early in the VII program so that the national system can be built on a technically solid foundation. In the absence of an efficient architecture, there is a serious risk that the VII data will be so difficult for users to access that it will not be used.

The work that is needed here includes:

- definition of relevant use cases that need to be supported with data of specific geographic scope and latency;
- definition of a hierarchical architecture to support those use cases, with estimates of the data traffic within and between each layer for a mature nationwide deployment;
- identification of most effective ways of labeling data for efficient sorting and access;
- definition of data aggregation logic needed to support the relevant use cases;
- development of prototype software to implement the data sorting and aggregation;
- testing of prototype software in simulation to identify its capabilities and limitations and needs for improvement.

This is a technically challenging activity that is likely to need three years of work, at an average level of effort of 15 to 20 full-time people, for a total of 45 to 60 labor years. It is vital that this be completed prior to the design and implementation of a large-scale VII Field Operational Test (FOT), so that the knowledge gained from this work can be incorporated into the FOT.

2.2 Development of public sector benefit and cost estimates

In order to make the informed deployment decision hypothesized by the VII program, public agencies will need to have well-supported analyses of their benefits and costs associated with VII deployment and operation. This will require transportation system modeling that exceeds the capabilities of existing tools and approaches.

The cost estimates should be tractable once the VII system architecture and technologies have been determined in the national program. Specific assumptions will have to be made about the pace of deployment so that the number of production units per year can be estimated, but from that point on the exercise should be relatively straightforward. It is important to make sure that it includes operating and maintenance costs as well as capital costs, since these are likely to be more significant fractions of the total life cycle costs than they are for typical civil infrastructure for transportation.

The serious challenge is in estimating the benefits of VII to the transportation system, to society as a whole, and to the individual agencies that will have to pay for the system deployment and/or operation. This issue remains important regardless of whether the VII infrastructure is deployed by public agencies or by the private sector, because in the latter case it will still be necessary for the public agencies to determine how much they are willing to pay the private sector for the data that they need.

The benefits need to be estimated relative to an appropriate base case, which could be either "do nothing" (rely only on systems already in place), deploy other more conventional ITS alternatives, or depend on conventional non-ITS alternatives. The same base case that is chosen for benefits evaluation needs to be applied for the cost comparisons as well. Since VII data should enable transportation operational strategies that are not possible using the other alternatives, it will be important to be able to model the effects of these strategies and estimate their benefits. This is likely to need new transportation modeling capabilities, beyond the current state of the art.

Transportation needs and operations vary greatly from location to location, depending on the density of development and existing transportation system infrastructure and capabilities. It will therefore be necessary to estimate benefits for a wide range of conditions in order to represent the nation as a whole, which indicates the need for site-specific case studies in multiple locations that span the full range of conditions in the U.S. (from highest density urban to lowest density rural).

The work that is needed here is expected to include:

- Specify the applications that will be built on the VII data, based on consultation with the stakeholder organizations that would be the deployers and beneficiaries of those applications (state, regional, county, and local government agencies).

- Define VII deployment staging scenarios, representing the annual rates of installation of roadside units over a full deployment planning horizon.

- Based on VII program definition of architecture and technologies, estimate costs of deployment and operations by year.

- Identify representative deployment case study sites that capture the diversity of conditions across the country, and where state and local agencies are interested in participating in VII evaluation and deployment.

- Define traffic management functions and strategies that could be implemented with widespread availability of microscopic data about vehicle movements on all categories of roads, and develop models to represent their effectiveness in improving traffic flow and/or safety.

- Apply a combination of existing and newly developed models to estimate the traffic improvements that could be gained from use of VII-based applications (Use Cases) compared to the other relevant alternatives for implementation of these capabilities.

- Based on citations to the best available data on safety improvements, estimate the safety benefits that could be gained from VII (recognizing that there are large uncertainties here, it is likely that this evaluation will have to be done parametrically, for a variety of assumed levels of effectiveness).

- Compare the incremental benefits and costs of the VII-based applications to the conventional ones, combining as many applications as appear to be realistic to apply to each operating scenario. It is important to address the combinations here, because the costs are likely to be nearly fixed and largely independent of the applications, while the benefits will be the sum of the benefits associated with each application.

This is also a challenging activity in several dimensions, because of the needs to develop new transportation models to capture the new effects of VII-based information services, to identify suitable case study sites and implement the models there and to develop credible results in an environment with much uncertainty. It is expected to need about 2.5 years of work, at an average level of effort of 10 people, for a total of 25 labor years of work.

2.3 Definition of backhaul system architecture

The current VII program activities are addressing the applicability of a variety of communications technologies to provide the backhaul links from RSEs. However, they do not appear to be addressing the broader issue of the most appropriate backhaul system architecture from the perspectives of efficient network management and data aggregation. This is in some sense an expansion of one of the bullet items already identified in Section 2.1, because it is closely related to the data aggregation issue. The focal point of this topic area is identifying how to most efficiently connect the half million RSEs around the country to the owners of the applications that need to exchange data with them. The local safety applications are simple to address from this perspective because they only involve a single RSE. However, the issues are much more complicated for the applications that are based on communicating specific packets of information between an individual vehicle and a service provider (such as electronic toll collection or other forms of electronic payment, or specific information requests from a vehicle operator) and the applications that depend on aggregation of data from multiple vehicles.

The issues that must be addressed here involve definition of the number of layers of information hierarchy to include in the system and the ratios between the numbers of entities that will exist at each layer. These are likely to be different for high-density urban regions and low-density rural regions, particularly considering the differences in cost and performance of backhaul communication technologies for differing distances and volumes of information.

The work that is needed here is expected to include:

- Definition of a full range of representative application environments, with different distances between RSEs and different densities of vehicle traffic at each RSE;

- Definition of logical structure for layering the VII architecture, to be consistent with the levels of data aggregation that should be needed to support typical user services. This should be based on consideration of the volume of data that needs to be passed up the hierarchy between individual RSEs and the top level, the costs of aggregating and communicating those data at each step; and the other available information infrastructures that may be combined with VII at each step in order to provide economies of scale.

- Modeling the logical structure for each application environment and then using the models to evaluate system design trade-offs;

- Recommending a backhaul architecture with intermediate layers, to augment the architecture currently under development in the national VII program.

This work is urgently needed in order to contribute to the definition of the national VII architecture, so it needs to be squeezed into the shortest feasible period of performance. It should be possible to accomplish this within the space of about one and half years, at a level of effort of about ten full time people, for a total of 15 labor years of effort.

2.4 <u>Development and standardization of needed wireless communication capabilities</u>

Wireless communications between vehicle and roadside infrastructure and among vehicles has been the strongest technical focus area within the national VII program until now, and has received substantial investments of attention and resources. It is difficult to keep track of exactly what is and is not being addressed within that work, but there have been some indications that the national program is not addressing the full range of important issues. It would be nice to learn that this turns out to be an unnecessary expression of concern, however in the event it is not, it will be important to make sure that the following issues are given adequate attention:

- Design and analysis, followed by testing, of wireless protocols to ensure that low-latency safety-critical messages can be successfully received under worst-case traffic and channel loading conditions (which could include up to 600 vehicles trying to share the safety channel at one major intersection RSE). This is likely to include work on MAC layer protocols to adjust wireless broadcast power based on vehicle speed and local traffic density and to adjust wireless message update rates based on assessment of possible hazardous conditions.

- Design and testing of wireless communications for close vehicle-vehicle cooperation (with message updates potentially faster than 10 Hz, but at relatively short range), as well as vehicle-infrastructure cooperation;

- Evaluation of communication technologies that complement DSRC at 5.9 GHz, especially for use in lower-density locations, so that VII can rely on the most appropriate mixture of technologies to meet its needs, rather than being tied exclusively to DSRC;

- Definition of minimum performance requirements for the wireless communication links that support VII services, in terms such as minimum acceptable availability, maximum permissible length of intervals of unavailability, maximum permissible packet loss rates under various operating conditions, etc.

The specific research questions and resource needs will depend directly on the extent to which these issues are being addressed within the national VII program, which makes it difficult to estimate specific time and resources needed here. Assuming that these are not being addressed at all in the national program, this additional research is likely to need two to three years of work by about ten people, for a total of 20 to 30 labor years of effort.

2.5 Development of new transportation system management approaches building on VII data

Transportation system management has historically been limited in the breadth and depth of its actions by the limitations of the available data (especially real-time data). VII raises the potential for an unprecedented quantity and quality of data to characterize transportation system operations, so that the traditional limitations would be overcome. This calls for a fundamental re-examination of transportation system management strategies, to be able to make best use of the newly available data, but also to ensure in advance through proper planning that the data that are being collected and processed have a real opportunity to lead to major improvements in transportation system management.

This topic was discussed in a roundtable workshop at the ITS America Annual Meeting on May 3, 2005, and the summary of that workshop is posted on the Web at www.ntoctalks.com/icdn/vii_roundtable_itsa05.php That discussion was kicked off with a presentation by Jim Wright of MnDOT, who provided a summary of a variety of ways in which VII could influence transportation system management, and this was followed by group discussion. Key opportunities raised in that presentation and meeting, based on the information available on the Web, included:

(a) Planning applications, including improved modeling, forecasting and O/D studies, even if the VII architecture prohibits direct collection of O/D data on individual travelers;
(b) Traffic control improvements for detection, arterial signals, ramp metering and work zones, based on second-to-second changes in conditions and much higher fidelity information about what is really happening, including vehicle speeds and accelerations. This would probably lead to the need for new traffic control algorithms to replace the current generation, which are based on more aggregate estimates of traffic volume and occupancy. It also raises the possibilities for a more decentralized traffic control system architecture than currently in use.
(c) New traffic management strategies to address ramp metering, incidents, special events and work zones, based on better knowledge of the real-time condition of the system and based on the possibility of integration with real-time traveler information. This opens the possibility of diverting vehicles in specific numbers onto specific alternate routes, of intelligently prioritizing the allocation of emergency response crews to optimize safety and traffic impacts, and of automating the control of both traffic flows and individual vehicles.
(d) High-fidelity traveler information available in real time for dissemination over public media outlets, the Internet and changeable message signs. This could include customized 511 information for travelers, as well as more effective Amber Alerts, and could make it possible to dynamically assign vehicles to the fastest routes. That raises important technical and policy issues about which travelers should be assigned to which routes, in which political jurisdictions. It also opens the possibility, when all vehicles are equipped, that some functions that are currently performed by infrastructure elements (such as both static and dynamic signs) could be performed by in–vehicle displays, customized to the needs of the individual drivers.
(e) Asset management and maintenance functions of state and local agencies could be significantly facilitated, providing real-time inventories of signs and pavement conditions, including snow/ice control and identifying potholes as soon as they appear. These could produce labor savings, as well as changes in the way these functions are managed.
(f) Performance management of the transportation system could be raised to a completely new level with the availability of comprehensive real-time information about its operation. Transportation system managers could be made accountable for the performance of their system based on authoritative measurements of how well it is performing, and maintenance and construction

contractors could similarly be held accountable for the impacts that their work has on accurately measured transportation system delays.

These opportunities are beyond the current state of the art because VII-level data have never been available before. Research is needed to explore the implications of each of these opportunities, as well as to seek other opportunities based on the advent of VII. The work that is needed here is expected to include:

- Building on the topics identified above, convene a national workshop to identify additional new opportunities arising from the availability of VII data, and to prioritize all the opportunities, bringing together experts from throughout the transportation community.

- Define the planning applications that could be implemented, assuming that confidentiality of travelers' information is well protected within the VII architecture. In particular, determine how effectively O/D information could be synthesized from VII data when individual vehicle identifiers are changed at each RSE, and identify the vehicle data elements that are needed from a VII system in order to support transportation planning needs. Develop the algorithms needed to synthesize the important aggregate planning information from the data that would be communicated from the individual vehicles. This is likely to require two years of work, at a level of effort of about five people, for a total of ten labor years of effort.

- Define and test new traffic control approaches suitable for isolated intersections, corridors of intersections and two-dimensional regions of intersections, based on availability of VII data from various percentages of the vehicles on the road. Define deployment staging strategies to advance from early stages with small vehicle market penetration, through intermediate and then high levels of market penetration of equipped vehicles. Identify the VII data elements that would be needed to implement these and the acceptable latency for each element. Define an appropriate information architecture for implementing such control, considering which functions need to be centralized and which can be distributed in the field. This is likely to require three to four years of work, at a level of effort averaging ten people, for a total of 30 to 40 labor years of effort. This would lead up to the stage that a large-scale field operational test of several of these approaches could be performed to verify their effectiveness.

- Define new traffic and incident management strategies to build on VII data availability. Develop the concept of freeway ramp metering at the level of merge assistance to individual vehicle drivers, helping them find specific gaps in the mainline traffic flow, and then assess its implications for safety and traffic flow capacity and congestion. Develop new incident management strategies based on assuming precise and timely knowledge of incidents and their severity and assess their effectiveness compared to today's alternatives. Define strategies for distributing traffic flows throughout a regional road network in order to balance or minimize traffic congestion, and evaluate their likely effects on both local and regional congestion using traffic simulations. For each of these strategies, design and conduct limited-scale local tests of their effectiveness in order to develop the knowledge needed to support a subsequent large-scale field operational test. This work is likely to need a team of about ten people for a period of three or four years, for a total investment of the order of 30 to 40 labor years.

- Develop the concept and methods for real-time dynamic traffic assignment, so that specific vehicles can be assigned to the "best" routes, and these can be continually updated as traffic conditions change. Evaluate its effectiveness on both individual and aggregate regional travel delays by means of a site-specific case study using simulation, at a sufficient level of detail that it could serve as the basis for a future large-scale field operational test. This is likely to be a three-year effort, at an approximate level of effort of ten people, for a total of 30 labor years.

- Identify the VII information that transportation agencies need to support their asset management and maintenance functions, and then define algorithms to extract it from the anticipated streams of real-time data received from vehicles. This is likely to be a two-year effort, at an approximate level of effort of four people, for a total of 8 labor years.

The combination of all of these activities could take up to four years and require something in the range of 108 to 128 labor years of effort. Some of these issues may be addressed in the "Mobility Services for VII" program if it is promoted from "Tier Two" to "Tier One" status in the national ITS program.

2.6 Definition of economic models to facilitate deployment

One of the most sensitive elements in the national discussions about VII has been the "business model" for system deployment and operation, because this is where the economic interests of the different categories of stakeholders are likely to diverge. The private sector participants in the program have well-established approaches for assessing the economic viability of their investments, but the public sector participants face a new challenge because of the mixed public-private character of VII. Their investment decisions are not necessarily as straightforward as they would be for more traditional public-only projects. Research on relevant economic models can help the public sector identify the most appropriate ways for them to participate in VII. Suggested activities in this area include:

- Review prior experience of public-private partnerships and other mixed forms of investment, with an eye to identifying their strengths, weaknesses and potential applicability to the VII setting. This should include other transportation projects such as recent toll road developments, as well as projects from other fields of activity such as the Tennessee Valley Authority.

- Identify legal limitations on the ability of public agencies to participate in different types of institutional arrangements, such as consortia and non-profit corporations, and the changes that would be needed in order to remove limitations if one of these were determined to be otherwise the most desirable.

- Identify the categories of public benefits that would merit public investments in VII development, deployment and operation. Consider the differences between direct benefits to public agency budgets, such as savings in current capital or operating costs for other systems within their responsibility, and indirect societal benefits such as improvements in public safety, traffic congestion, environment and quality of life. Once these benefits were to be estimated (for example, through the work described in section 2.2 above), how would those benefit estimates be converted into estimates of justifiable public investments for VII capital and operating costs? This is likely to require interaction with representative public agency stakeholders, perhaps in a workshop format.

- Define the process(es) that will be most likely to succeed in obtaining widespread national participation in VII by the needed state and local stakeholder agencies, based on developing answers to questions such as the following: What level of estimation or demonstration of benefits will be needed to be sufficiently convincing for enough agencies? What incentives or transfer payments are going to be needed to convince the less-interested agencies to participate in order to enable the national deployment to reach the necessary critical mass? What kinds of consultation among stakeholder agencies and their respective national associations will be needed in order to build consensus on the value of national VII deployment and operation? Will substantial federal funding through the next transportation reauthorization legislation be sufficient motivation, or will more be needed?

These activities are likely to require about one and half years, at a level of effort of about two full-time people, for a total resource need of about 3 labor years.

2.7 Definition of criteria for prioritizing infrastructure installations

It will not be possible to install the entire national VII infrastructure at one time, but that installation will need to be spread out over a period of multiple years. There are likely to be a variety of competing criteria for determining which locations receive their VII infrastructure earlier, such as:

- politically well connected to get earmarked funding
- wealthy enough to be able to afford it more easily
- technologically sophisticated population and/or agency eager to be a leader
- most severe need based on safety and/or traffic congestion problems
- easiest to implement in connection with other already-planned transportation improvements

- desirable showcase for an agency or local industry, especially if coupled with a special high-profile local event (such as a Super Bowl or Olympics)
- presence of a prominent champion in local government or industry, willing to make an up-front investment
- heightened concerns about terrorist attacks or natural disasters
- completing a "missing link" between other deployment sites
- large enough concentration of local interests to achieve an early critical mass.

Conversely, some factors could tend to impede local VII implementation, regardless of other advantages or needs:

- active hostility by prominent local political or industry leaders
- heightened concerns about personal privacy issues
- local agency budget or staff limitations
- lack of leadership in key agencies.

Regardless of these local circumstances, VII is intended to be a national system and needs to be national in scope in order to realize its potential benefits. In order to minimize frictions among localities (and between the public and private sectors) over the infrastructure deployment timetable or "roadmap", it would be useful to develop a uniform set of national criteria for prioritizing VII infrastructure deployments, based to the maximum possible extent on the nationwide benefits and costs. Carefully structured research could define the process for developing these criteria, as well as developing the criteria themselves. Suggested activities would include:

- Define a process for engaging the key stakeholders in consideration of the deployment prioritization, based on consultations with their national representatives. Then, implement the process to work through the issues to aim toward a durable consensus approach.

- Identify a comprehensive list of deployment prioritization factors, such as those listed above, and describe the generic advantages and disadvantages of each. Then define what kind of analysis or modeling work needs to be done to provide sufficiently accurate estimates of these advantages and disadvantages to be useful for decisionmaking at both the national and local levels.

- Implement the analysis and modeling work for a wide variety of representative case study locations throughout the country, demonstrating the local and national advantages and disadvantages of each kind of prioritization criterion. Use these results as the basis for workshop consultation with the stakeholder groups, to identify their concerns, to determine their level of comfort with the results, and to start the process of seeking the needed national consensus.

- Continue with the stakeholder consensus-building process as defined in the first bullet item activity, aiming toward a sufficiently durable consensus that the national deployment can be allowed to proceed primarily based on maximizing the national ratio of benefits to costs.

This work is expected to require three to four years of continuing interactions with the stakeholder community, although the research elements are likely to be scattered intermittently throughout this period. Estimating an average of three full time equivalent people, the resource needs would be in the range of 9 to 12 labor years.

2.8 Summary of Research Needs

Based on the foregoing discussion, it appears that almost all of the needed research could be accomplished within a period of three years. The resource estimates are summarized here for convenience and to provide perspective:

Topic	Duration (Years)	Full-time people	Labor Years
1. efficient data sorting and aggregation	3	15-20	45-60
2. public sector benefit and cost estimates	2.5	10	25
3. backhaul system architecture	1.5	10	15
4. wireless communication	2-3	10	20-30
5. new transportation system management	2-4	39	108 - 128
6. economic models to facilitate deployment	1.5	2	3
7. prioritizing infrastructure installations	3-4	3	9-12
Total			225 - 273

If this number of labor years were to be expended in a university setting, the total cost would be in the range of $25 to $30 million, while in a commercial contracting setting it could be double that cost. Considering the urgency of the work, and the need to conduct such a large research program within little more than three years, this is likely to exceed the capacity of the universities having relevant knowledge, expertise and interest. Therefore, much of it is likely to have to be done in the private sector at higher costs. Even at a cost of $50 to $60 million, it is a small fraction of the investment that will be needed to deploy VII. It would be prudent to make this up-front investment in the needed research in order to improve the likelihood that the much larger investment in VII deployment will be spent wisely.

3. **Candidate VII Information Architecture and Business Model Alternatives**

3.1 Current National VII Architecture

Relatively little information has been made public about the architectural approach proposed for the national VII program. The "VII Architecture and Functional Requirements Draft, Version 1.0" Report dated April 12, 2005 provides some indication of the current thinking in the national program, but it is not clear whether it is generally accepted within the national VII coalition. Nevertheless, in the absence of other information, that report will be considered as a starting point of reference here.

The architectural approach defined in the VII report is “flat” rather than layered, with only one intermediate layer between the individual local RSEs and the national VII Operations Center(s). That intermediate layer is described in terms of “VII Message Switches”, each of which would communicate with thousands of RSEs as well as with all its hundreds of peer Message Switches around the country. However, the message switches do not store or process any data; rather they only appear to pass the data from place to place based on a “publish and subscribe” model. The RSEs appear to have some information processing capability to handle local safety applications that require very low latency, but this processing does not appear to involve data received from elsewhere in the VII network. With that exception, it appears that all information processing (including data aggregation) is left to the responsibility of the individual applications that would subscribe to receive the raw VII data.

The VII architecture appears to make special provisions for transactions between the automotive OEMs and their customers’ vehicles, but there is no indication whether similar private transactions between travelers and independent information service providers (ISPs) or special authorities such as toll road operators would be accommodated similarly. Public agencies appear to have a lower level of access to data, ostensibly for privacy reasons. The priority for RSE processing of messages was specified as:

1. safety applications
2. information *from* state DOT *to* vehicle
3. automotive OEM vehicle support messages
4. financial services, including tolling, parking and purchases (somewhat surprising considering the time criticality of tolling transactions when passing a toll collection location).

One place where the VII architecture defines a hierarchical structure is in the categorization of “Network Users”, however these are considered to be outside interfaces to the VII architecture rather than an integral part of it. The categories of network users were defined as:

- national
- state TOC
- regional TOC
- local TOC
- local safety systems
- local controllers and processors
- other vehicles (for vehicle-to-vehicle applications).

One of the more surprising aspects of the VII architecture is that messages from vehicles are categorized at the RSE into one of three classes:

(1) a local message to be sent directly to the local controller or safety application processor
(2) a message for a specific other user such as an automotive OEM, which the message switch routes to that subscriber
(3) a message containing "public sector data", which the message switch routes to subscribers to that data stream.

This segregation of messages is likely to introduce significant inefficiencies. For example, the vehicle's periodic broadcast of its location and speed is extremely important for local safety applications such as CICAS, but it is also vital public sector data for determining link traffic volumes and speeds. Indeed, the GPS coordinates of the vehicle are likely to be vital contents of messages of all three classes. If this information needs to be transmitted in separate messages for separate purposes, the message traffic volume appears to be growing significantly without any added value.

The flat architecture and lack of intermediate data aggregation are likely to make use of the VII network data extremely difficult for users who want to derive aggregate information about the state of the traffic system or the condition of the infrastructure. They will need to "subscribe" to data flows from all message switches in their region of interest, and will then be faced with the daunting challenge of extracting the small fraction of this torrent of data that is really needed for their application. This is akin to finding a needle in a very fast moving haystack.

Other important limitations include:

- "conversational" messaging involving interactions with consecutive RSEs will not be supported by the VII architecture, meaning that each transaction between a vehicle's OBU and an individual RSE generally needs to be complete and self-sufficient, and an incomplete transaction cannot be carried over to the next RSE unless special provisions are made for that within the applications built on top of VII (which would appear to be complicated to accomplish).
- Different RSEs will have differing functions, and not all will be fully capable. For example, some may only broadcast information but not receive any information back from vehicles.
- No data will be retained within the VII system, but data archiving will be the complete responsibility of subscribers to the VII data stream
- The baseline plan assumes two national control centers, one on each coast, to provide redundancy, and each center would contain all the message switches. This appears to require the raw data from every vehicle to every RSE in the country to be communicated to these two locations, which looks astonishingly inefficient and costly.
- The "VII Network Operations Entity", a national body managed by the state DOTs, USDOT and automotive OEMs, would be responsible for all hardware support,

including maintenance and debugging. This appears to eliminate opportunities for synergy with local traffic control equipment maintenance, and may even preclude the VII hardware from being located within existing traffic control cabinets, introducing significant additional deployment and O&M costs.

- The architecture does not address data flows among vehicles, but leaves those as the province of the automotive OEMs. This appears to make it difficult to capitalize on opportunities to use vehicle-to-vehicle message relaying in locations where RSEs are far apart and only infrequently within range of any individual vehicle. Such message relaying could be particularly important for public safety information in rural areas.

3.2 Additional Considerations, Beyond Current National Architecture

Certain categories of data from vehicles will need to be accessible on a nationwide basis, but an initial assessment of the possible vehicle data and applications indicates that this should be a very limited amount of data, and all of it should be reported on an event-driven basis, not repeatedly or continuously. This means that these would be individual one-time messages from a vehicle, not repeated at the 10 Hz update rate of most other messages. An initial listing of data of nationwide significance is:

- electronic payments for services (transportation or commercial)
- "mayday" messages for emergency assistance
- Reports of emergencies of national significance (natural disasters, terrorist attacks)
- Requests for information from national commercial services
- Emergency vehicle messages
- Commercial vehicle regulatory messages.

Other than these, the large majority of the remaining data could be aggregated at a variety of other levels and do not need to be accessible at the national level. Indeed, on a regional or national scale this type of raw data is akin to noise because it cannot be managed effectively. It only becomes useful when it is aggregated at levels that are usable by the "network" applications. Examples of such levels are suggested below:

Level 1 – Nationwide

Data from every vehicle in the country would flow to a common access point, where users would extract from it what they wanted. This appears to make no sense, and does not appear to have been suggested seriously at this point.

Level 2 – Regional

At an earlier stage of the VII program, there was a suggestion of about 100 regional data "caches", although this was not mentioned in the February meeting nor the April 12 report. Perhaps these have been superseded by the "VII message switches". Some of the regional boundaries would probably be metropolitan areas, or perhaps some of the largest

(New York, Los Angeles) would even have to be subdivided into more than one region. At the opposite extreme, other regions might be defined as clusters of several lightly-populated states. If we assumed 100 regions, each would have an average of 5000 RSEs, although the range could probably be anywhere from 1000 to 20,000 RSEs considering the likely diversity of the regions.

In a large urban region, a couple of million vehicles could be operating simultaneously during the peak period, transmitting 1 gigabyte per second of data. However, the vast majority of this data is not needed at the regional level, but would only be needed at lower levels. The only information that is likely to be needed at the regional level is:

- current aggregate weather conditions
- average travel times, volumes, and speeds on major trunk road links
- identification of major incident locations, severity and duration
- major emergency response activities.

Level 3 – Sub-regional

The sub-regional level would correspond to individual municipalities or districts within a large urban region or small metropolitan areas within lightly-populated states. An average region might have 50 sub-regions, and an average sub-region might have 100 RSEs, although they could probably range from 20 to 500 each, depending on local conditions. During peak periods, the sub-regions may be hearing from as many as 50,000 vehicles simultaneously.

Most local traffic management and traveler information services are likely to need information at the sub-regional level, where it would be useful to have access to data regarding:

- travel times, volumes and speeds on all roads down to the collector/distributor level
- location, severity and duration of all incidents
- safety hazards in the roadway (obstacles, stopped vehicles, potholes)
- localized road condition information (wet or icy patches, potholes needing repair, etc.)
- parking availability and reservations
- transit and other vehicle fleet operational data
- all emergency response activities.

Level 4 – Corridor

Within a municipality or a district of a large city, activities would still be concentrated within more localized corridors, where the traffic signals could be closely coordinated with each other and travelers may have a variety of choices of parking facilities before walking to their final destinations. These corridors might average ten RSEs, but could contain up to perhaps 100 RSEs, and could be receiving data from one thousand to 10,000 vehicles, depending on their geographic scale and traffic density.

The most microscopic traffic management and traveler information services are likely to be delivered at this level, where the data from individual vehicles could potentially be useful without aggregation.

Level 5 – Sub-Corridor

The largest and most congested of the urban corridors at Level 4 may be too large to efficiently manage the most microscopic level of vehicle data, at the individual vehicle level. Therefore, it is possible that these could benefit from an additional layer in the architecture, in which individual vehicle data would be collected from a handful of RSEs and used for the most local traffic management and traveler information services, and only aggregate traffic flow volume, travel time and speed data would be passed up to the corridor level.

In the interest of efficiency, data should only be passed from a lower level to a higher level if it serves a useful purpose at that higher level. If an aggregated summary of the lower level data is sufficient to meet the needs at the higher level, it should be aggregated at the lower level before being passed up to the higher level. Decisions about what aggregation should take place at each level have serious implications for the long-term capabilities, performance and cost of the system, and therefore need to be considered carefully in the design of the architecture right from the start.

The current study is only exploratory and does not have the resources to do an authoritative analysis of the trade-offs associated with different levels of data aggregation within various geographic scopes. However, examples of the *types of issues* that could be addressed by such a more thorough study are suggested here:

Individual RSE (at intersection):

- Cooperative intersection collision avoidance systems will require microscopic data about individual vehicle movements (locations, speeds, accelerations, driver intentions, etc.) at the maximum update rate (10 Hz), so these would not be aggregated.

Corridor and Sub-corridor (Levels 4 and 5):

- Trajectories of individual vehicles (locations, speeds) could be useful for coordinating traffic signals and dynamically adjusting their phase transitions, but they could probably be updated at 1 Hz rather than 10 Hz. The aggregation of this information would provide queue lengths at intersections.

- Emergency vehicle pre-emption and transit signal priority signals are needed with minimum latency, but these could be individual event-based transmittals rather than repeated continually.
- Local hazard information (wet or icy pavement, stopped vehicles, objects in road) is needed, probably at the maximum update rate.

<u>Sub-regional (Level 3):</u>

- No individual vehicle data would be needed here, but only aggregate data by network link (average speed, travel time), probably for all roads. These data could be averaged over 15 or 30 seconds, significantly reducing the update rate.
- All incidents requiring incident response would need to be reported here, including crashes involving lane blockages, injuries or fatalities, as well as events requiring immediate maintenance action (potholes in road surface, fallen tree branches, etc.), but these reports would be event-based rather than periodic.
- Updates of reports on incidents and traffic bottlenecks would be needed, with estimates of their severity and duration, with an update interval of the order of one minute (but these would be derived by off-line infrastructure processing, based on analysis of incoming link travel data).
- Estimates of local weather and road surface conditions, aggregated from vehicle data, updated on a one-minute basis.
- Transit and other fleet management information updates, with update intervals from one second to one minute depending on specific local service needs.

<u>Regional (Level 2):</u>

Most data needs are likely to be similar to sub-regional, but at a higher level of geographic aggregation:

- Traffic flow data only on major arterials and highways;
- Only major incidents of regional traffic impact, not all incidents.

<u>National (Level 1):</u>

Other than special services called out earlier, the only information likely to be needed at this level regards major incidents of national significance, such as natural disasters (earthquakes, floods, large storms) or terrorist attacks. This type of information is likely to be conveyed based on human judgment and intervention, rather than automatically, so it probably does not need to be incorporated directly in the VII information processing architecture.

Obviously, detailed studies are needed of all of the planned public-sector use cases, to identify their specific data needs and to consider the advantages and disadvantages of addressing them with different levels of data aggregation. Efficient data aggregation would be greatly facilitated by appropriate labeling of the data packets transmitted by each vehicle, so it will be important to address the data aggregation architecture early enough that the appropriate labels can be specified in the definition of message sets. More importantly, the data aggregation approach to be adopted has fundamental significance for the design of the information processing hardware and software needed at each RSE and for the communication systems needed to link them to the rest of the VII system, all the way up to the network layer system management function.

4. Concluding Remarks

This report is intended to suggest the most urgent research questions that need attention in order to support the development of the most effective VII capabilities and the best-informed public sector deployment decisions regarding VII. It is not intended to be exhaustive in scope, but it is intended to set key directions for VII-supporting research in the immediate future, so that progress toward VII deployment is not impeded by serious knowledge gaps.