Understanding Key Tradeoffs for Cost-Effective Deployment of Surveillance to Support Advanced Traveler Information Systems (ATIS)

April 2004

Alan Toppen
Dr. Karl Wunderlich

Contract Sponsor: Federal Highway Administration
Contract No.: DTFH61-00-C-00001
Project No.: 900610-D1
Department: J190

MTS™ Mitretek Systems
Falls Church, Virginia

DEPARTMENT OF TRANSPORTATION UNITED STATES OF AMERICA

Acknowledgements

The authors wish to acknowledge numerous colleagues for their contributions to this study:

In the Federal Highway Administration (FHWA): Joe Peters, the sponsor of this work at the ITS Joint Program Office.

At the University of California at Berkeley: Chao Chen, Karl Petty, and Hamed Benour for access to and use of the Los Angeles freeway detector data.

Colleagues at Mitretek who contributed to the study: Soojung Jung, James Larkin and Don Roberts.

Table of Contents

List of Figures

List of Tables

Abstract

In this paper, we present our understanding of a key tradeoff in ATIS investment planning: Investment in expanding surveillance coverage to additional miles of roadway vs. improving the accuracy of the information provided on roadways already covered by ATIS. Building on previous work using the Heuristic On-Line Web-Linked Arrival Time Estimation (HOWLATE) evaluation methodology, we explore the relevant cost and benefit relationships to provide guidance to high level decision-makers implementing ATIS systems. The actual critical decision points will vary from region to region, but we contend that the fundamental relationships are the same. While the cost of adding surveillance is roughly the same regardless of the amount of existing surveillance in place, the user benefits resulting from initial miles of surveillance in a region are greater than for surveillance miles added to an existing system. At some point, the marginal benefit to ATIS users fails to exceed the marginal cost of additional surveillance. We estimate this point may arise in the vicinity of 50-60% of full network coverage (defined as all major freeways and arterial roadways.

For ATIS systems that report point-to-point travel times, the gains in user benefits from improving the accuracy of an already very accurate system are less than for a system that is very inaccurate. As a guide, a minimum average error of 20% is required for the majority of users to garner mobility benefits from ATIS use. When travel time estimation error below 5% can be achieved, the costs of maintaining detectors in such a state of good repair may exceed users' mobility benefits.

Finally, we present an updated nomograph for decision-makers to consult regarding these relationships based on the current state of their regional deployments, or for initial planning for new deployments. In future work, we plan to deepen our understanding of what contributes to travel time estimation error and the costs associated with maintaining high levels of accuracy.

1. Introduction

Advanced Traveler Information Systems (ATIS) are widespread, largely due to significant public sector investment. As of 2002, 26 U.S. metropolitan areas provided automated telephone services to distribute freeway travel times and over 35 metropolitan areas provide freeway travel times or speeds via the internet (1). More recently, under the federal mandate for a national traveler information number (2), public transportation agencies continue their regional ATIS efforts with plans for larger, more comprehensive deployments. However, deployments are typically completed in phases and what may be considered "full deployment" is often the culmination of numerous intermediate steps. To date, little work has been done to identify how these incremental ATIS investments may be selected in the most cost-effective way.

We propose that an optimal ATIS investment strategy for a metropolitan area is one that produces the greatest user benefit at every step on the way to full deployment (we will define full deployment later). To that end, we identify a fundamental tradeoff between increasing surveillance coverage (i.e., the number of roadway miles for which traveler information is available) and improving the accuracy of the information for the roadway miles with surveillance. Intuition suggests improving each of these would have a positive effect on user benefit. Not only would they boost the benefits for current users of ATIS, they would likely attract new users—users for whom information is suddenly available on the routes they commonly use and users who are more sensitive to accuracy. Under certain conditions it is preferable to improve accuracy, such as when coverage is extensive and accuracy is low. At other times, however, it may be better to focus resources on expanding coverage, such as when accuracy is high and coverage is low. Between these two extremes, the most effective type of improvement is not always clear.

Figure 1 presents an earlier attempt (3) to illustrate the aforementioned trade-off in planning cost-effective ATIS deployments. Decision-makers considering investments in ATIS systems with poor accuracy (the "don't deploy" region) should not consider deployment unless they can implement a system with sufficiently lower error in reporting of ATIS. Past research has shown that commuters are better off without ATIS if the accuracy of the ATIS information is low relative to the day-to-day travel time variability they experience (3). Clearly, systems currently in this region should have improving accuracy as their primary objective. Conversely, sufficiently accurate ATIS services (the "add coverage" region) should consider geographic expansion rather than further refinement of ATIS accuracy. At higher levels of coverage or ATIS information error, decision-makers again find their ATIS systems in the "get better sensors" region and may need to consider further improvements in ATIS accuracy that may come at the expense of better or more reliable sensors. Continued investment in ATIS service coverage or accuracy reaches a natural end point in the "stand pat" region where both geographic coverage and accuracy of the ATIS service are at a level where the marginal benefits from additional improvements do not warrant the cost of such improvements. While we postulate that ATIS is far from this point in every metropolitan area in which it is deployed, based on the rate at which ITS is expanding, many regions will get to this point in the near future.

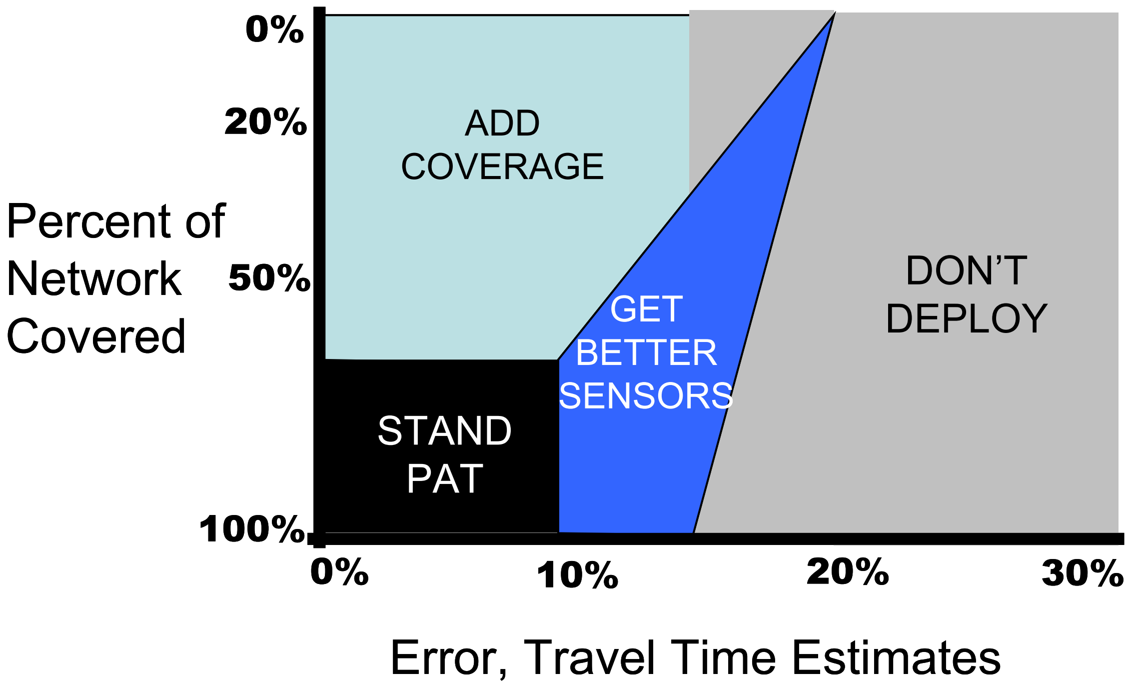

Figure 1. Notional Nomograph of Potential Decision-Making Regimes

In this paper, we aim to more robustly define the boundaries of the regions depicted in Figure 1. In previous work, we studied how user benefit is affected by varying accuracy and coverage levels (3). In this paper, we present cost models that complement the benefit models already developed and allow us to develop cost-benefit relationships. For example, the "stand pat" region is reached when the marginal costs of geographic expansion and accuracy improvements begin to exceed the marginal user benefits. The boundary between the "add coverage" region and the "get better sensors" region is the point where the benefit-cost ratio of adding coverage is identical to that of improving accuracy, i.e., the same benefit can be expected from an additional dollar investment in either. On either side of the line, the return on that dollar yields greater benefit for one versus the other.

Of course, these two investment decisions are not mutually exclusive—they often are implemented complementarily. We suggest that at every deployment state, there is a "trajectory" of maximum benefit-to-cost ratio. In addition, the boundary points in the notional nomograph are dependent on other factors, most importantly regional day-to-day travel time variability. As a result, the values derived for one city may not be applicable to another. Furthermore, parallel traffic management activities aim to reduce the variability caused by incidents, construction, weather, etc. Thus, these boundary points may change over time within a region depending on the success of these efforts. In this paper, we will take another step toward defining a set of guidelines for ITS decision-makers by which they might evaluate their various investment options regarding surveillance to support ATIS.

The following sections present a discussion of refinements to our understanding of the following relationships:

- Travel Time Estimation Accuracy and ATIS User Benefit

- Travel Time Estimation Accuracy and Costs to Deploy and Maintain ATIS
- Percent Deployment and ATIS User Benefit
- Percent Deployment and Costs to Deploy and Maintain the system

The depth of our understanding on each of these critical dimensions varies. However, what we do know can be utilized to update the 2002 Deployment Nomograph. In sections 3 and 4, we will present simplified models that generalize each of the four key dimensions of the ATIS investment decision space.

Based on these simplified relationships, section 5 presents two nomographs for decision-makers considering ATIS deployment or ATIS expansion. The nomographs attempt to bring together the complex interactions between cost and benefits for the purpose of developing cost-effective ATIS investment plans. We split the 2002 nomograph into two to resolve ambiguities between planning an initial ("green field") ATIS deployment plan and decisions about making incremental investments to an existing ATIS deployment. Section 7 concludes with a discussion of the implications of each nomograph as well as a summary of potential next steps.

2. The Costs and Benefits of Improving Travel Time Estimation Accuracy

This section discusses the relationship between travel time accuracy and detector maintenance costs. For a surveillance network already in place, accuracy is determined primarily by detector reliability. Figure 2 shows these two relationships graphically. When accuracy is low, benefit is low and cost is low. We expect that when accuracy is low, improvements in accuracy would garner significant benefit at a low cost. When accuracy is high, benefit is high and cost is high. However, most of the possible benefit has already been achieved and further improvements in accuracy result in small increases in benefit. We also expect that improvements in accuracy when travel time estimates are already very accurate are expensive as they may rely on new, complex or exceptionally costly detection technologies. From (4), we only need to extrapolate user benefits to regional benefits to derive the benefit curve in Figure 2Figure 2. "Benefit on a unitary trip basis can be calculated from (4) and aggregated based on regional demand patterns to generate a benefit curve as illustrated in Figure 2Figure 2." To generate the cost curve in Figure 2, we need to establish estimates of the cost of providing information of varying levels of accuracy. This will be discussed in detail in a subsequent section.

There are two points in Figure 2 that are of interest to us. The first is the breakeven point—denoted by the dotted vertical line—which is where the benefit and cost curves cross. To the right of this point—the shaded area—the cost of maintaining the high level of accuracy is greater than the total regional user benefit. In this area, more is being invested in ATIS compared to the amount of public mobility benefits garnered in return.

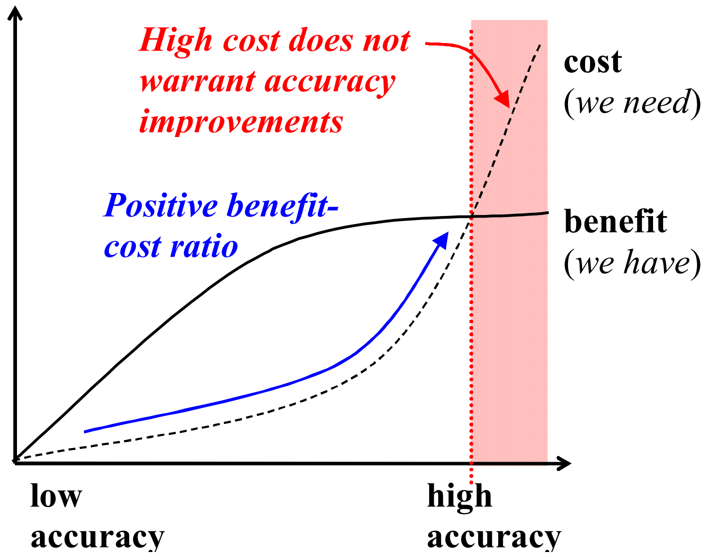

Figure 2. Cost and Benefits of ATIS as Functions of Accuracy

The second point of interest is the point at which the marginal cost exceeds the marginal benefit (solid vertical line). This is where the slope of the cost curve becomes steeper than the slope of the benefit curve. At this point it is no longer cost-effective to spend money to improve accuracy because the cost of improving accuracy by one percent is greater than the public benefit garnered by that investment. For example, if an additional 1% of travel time accuracy cost $10,000 annually and the resulting user benefit was $10,000 per year, we would have reached this point. This point is more important than the breakeven point for two reasons. First, it is the point of diminishing returns where the cost-benefit ratio is at a maximum. Secondly, there may not be a breakeven point. The return from ATIS may never exceed the investment. This is not necessarily a problem since ATIS does not exist in a vacuum. The surveillance infrastructure that supports ATIS is used by other ITS applications such as ramp metering, incident management, performance monitoring and others. Therefore, there are multiple ways in which ITS surveillance infrastructure benefits the public, of which ATIS is only one. What we are presenting here is merely meant to provide sound guidance for ATIS investment planners as they think about the potential return on their investment in ATIS for future deployment and maintenance decisions.

2.1. Travel Time Estimation Accuracy and User Benefit

In (4), we established a relationship between ATIS travel time accuracy and user benefit. This is shown in generalized form in Figure 3. We found that ATIS travel time error needed to be below 20% for the average user to consistently benefit from using the service in order to improve on-time reliability. Furthermore, we found there is little to be gained by reducing error below 5%, a level below which may be practically impossible given current technology. In our case study of Los Angeles, assuming 20% error—an accuracy that is probably close to what most ATIS providers are achieving— 37% of simulated trips benefited and among the trips that benefited, the average benefit was $0.96 per trip based on the utility function we employed (4).

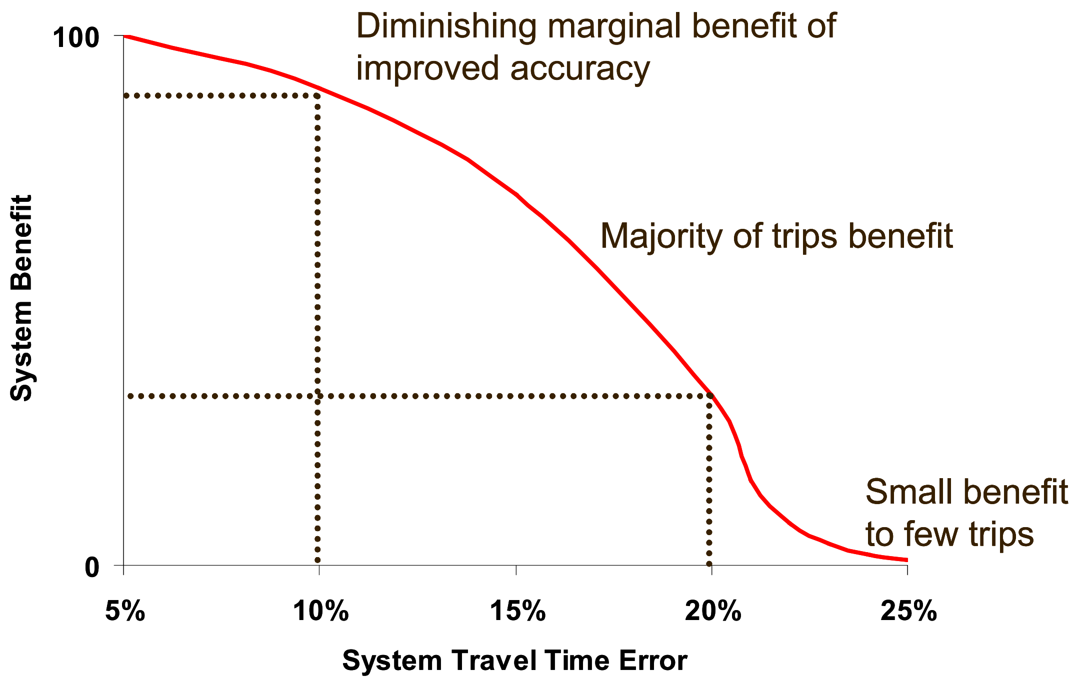

Figure 3. ATIS Benefit vs. Accuracy Relationship for Los Angeles

Improved accuracy has a two-fold effect. First, non-ATIS commuters for whom a travel time service was not valuable at higher error levels now see positive benefit from ATIS and become regular users. Second, existing users of the ATIS service realize a rapidly improving benefit per trip because of the improved travel time estimation.

As accuracy improves from a 20% error level to a 10% error level, the number of trips that show positive benefits for regular ATIS use in their commutes increases from 37% to 75%. The relationship show in Figure 3 is concave on the interval 20-25% error level and convex below 20% reflecting a diminishing return given the highest-value trips have already become ATIS users. There is a steady decline in the marginal benefit of improving accuracy from 10% to around 5%. In our experiments, we do not see any additional practical benefit for ATIS users from accuracy improvements beyond the 5% error level.

Note that this curve is independent of the costs of improving accuracy and assumes a stable deployment (no expansion). In general, the switchover point from convex to concave shape occurs between 10-20% travel time estimation error depending on the level of congestion and travel time variability in the metropolitan area.

2.2. Travel Time Estimation Accuracy and Costs to Deploy and Maintain

A survey of metropolitan areas that provide ATIS travel time estimates shows that loop detectors are still the predominate source of traffic measurement technology, though in some places non-intrusive detectors such as microwave or infrared are becoming more common. With the exception of Houston and New York, which have infrastructure to directly measure point-to-point travel times of vehicles equipped with toll tags, all other cities measure speeds at spot locations and extrapolate those speeds over homogeneous distances to estimate segment traversal times. Therefore, our analysis of costs will focus on the costs of loop detectors. In this section we will focus strictly on the cost of maintaining loop detectors; in the section on cost-benefit analyses for coverage we will look at the cost of building out surveillance on additional freeway miles.

The relationship between cost and ATIS travel time accuracy is complex. Numerous factors influence accuracy including:

- Type of detector (single vs. dual loops [indirect vs. direct speed measurement])
- Detector spacing (⅓-mile, ½-mile, etc.)
- Detector placement (influence of speed variability in merging areas)
- Detector reliability (percent off-line at any given time)
- Method of diagnosing and filling in missing data for off-line detectors
- For single loops, method of conversion from volume and occupancy to speed
- Method of conversion from point speeds along a segment to travel time
- Inability to account for interchange delays
- Time lag between measurement and reporting
- Forecasting method

While all of these factors can contribute to error, the primary determinant of accuracy is the reliability of the detectors themselves. Therefore, we will focus on the costs of maintaining a certain level of detector reliability. To this end, we will define reliability as the expected percentage of detectors that are down at any given time.

There is no way to directly relate cost and accuracy. Rather we must arrive at it from the following pair of relationships:

- Travel time accuracy and detector reliability, and
- Detector reliability and cost (maintenance budget)

To relate travel time accuracy and detector reliability, we draw upon data available from other cities. For instance, we have the notional relationship between detector reliability and travel time estimation error shown in Figure 4. From separate studies, we found that in San Antonio, approximately 11% of all detectors are down at any given time (5) and its ATIS system estimates travel times with an error of approximately 17% in peak periods (6). From other studies, the best possible prediction of travel times five minutes into the future has a 9% error (7,8,9). Beyond these two figures, the relationship is merely notional. We expect that single loops would have more error than dual loops because of the additional step of estimating speed from volume and occupancy measurements. As part of future work, we will study the effects of detector reliability and spacing in more detail using single loop data from Minneapolis/St. Paul, Minnesota.

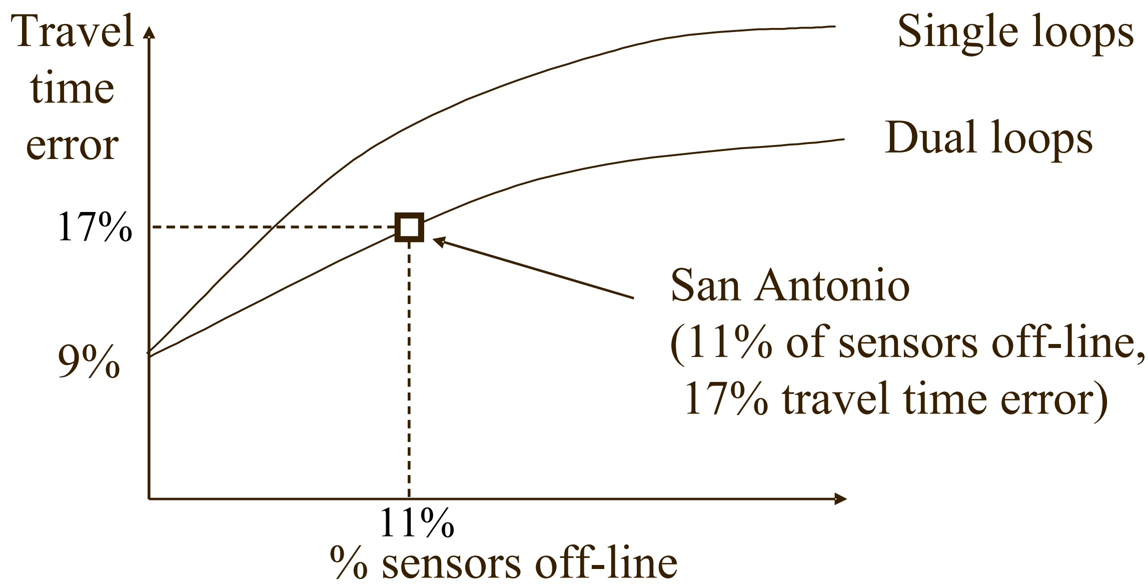

Figure 4. Proposed Relationship between Travel Time Error and Detector Reliability

To relate detector reliability and cost, we need to know something about the cost of maintaining a particular percentage of loop detectors in operation at any given time. This is elusive. The best we may be able to do is get a rough estimate of what cities pay for maintenance and what percentage of their detectors are down at any given time. Anecdotally:

- San Antonio spends $25,000 per year on detector maintenance for a system of 1550 loops, and 11% of them are down at any given time.
- Los Angeles spends about $500,000 per year on a system with over 5000 loops. Approximately 33% are down at any given time. The Los Angeles system is much older than the detector system in San Antonio (10).
- Loop detector replacement costs were estimated at $240 per loop per year based on data collected for the New York/ New Jersey TRANSMIT system based on an expected lifespan of 5 years (11,12)

Currently, we have the least amount of information to bolster our relationship between travel time estimation accuracy and the costs to deploy and maintain travel time estimation infrastructures that support these levels of accuracy. We do know that loop detector surveillance systems with short spacing (e.g. < 0.5 miles) can produce travel time estimates with as little as 10-15% error when properly maintained. For a common technology like loop detectors, we expect that when additional resources are directed to improved loop maintenance and sensor density we can expect steady improvements at relatively low cost. However, at some point the capability of loop detectors to provide improvement below a 10-15% level diminishes. To improve travel time accuracy below this threshold, more advanced technologies are required. We hypothesize that these more advanced technologies such as Automatic Vehicle Identification system (AVI) will have costs that rise rapidly as more and more accurate levels of travel time estimation are required.

The relationship shown in Figure 5 is somewhat speculative in nature and only notionally addresses a single technology (loop detectors) and a leap to some other more accurate and presumably more expensive travel time estimation technology. It is certainly possible that other technologies (current or future) have different cost functions relative to travel time estimation accuracy. Note that the relationship shown is independent of ATIS benefit and assumes a stable level of deployment. In future work, we hope to better understand the relationship between detector reliability and travel time estimation accuracy using detector data from Minneapolis and St. Paul, Minnesota.

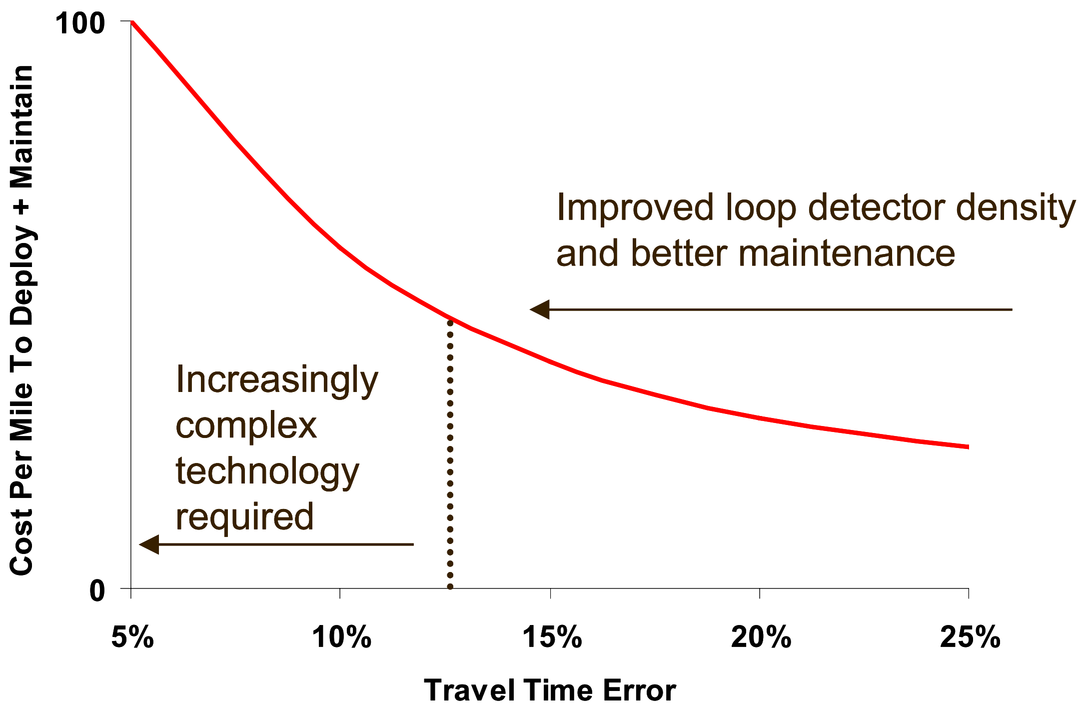

Figure 5. ATIS Accuracy vs. Cost per Mile

3. The Costs and Benefits of ATIS Surveillance Coverage

In this section, we explore the cost and benefit relationships of expanding surveillance coverage. These relationships are more easily developed than the accuracy relationships. Here we present what we know about how user benefit varies with extent of surveillance coverage as well as some sample costs of deployment. As we calculated in the appendix, the average per day benefit for ATIS is in the range of \$6.9-\$8.1 million per year in the Los Angeles area. Therefore, we can use this as a ballpark figure for benefit under full coverage, assuming a reasonable level of accuracy (10%-20%).

3.1. Percent Deployment and ATIS User Benefits

In (13), a logarithmic function is presented that describes a near-optimal incremental deployment strategy for links in a roadway network (assuming the accuracy of the sensors deployed is constant). A simpler regression model based on annual average daily traffic and geographical information approximates the log function. A generalization of these functions and their expected system benefit is presented in Figure

6. This generalization addressed two alternative deployment strategies (logarithmic or regression model) as well as differences in network shape, size and average congestion level seen in the metro roadway networks utilized in this study as test cases.

The relationship shown in Figure 6 has three linear components, corresponding to three breakpoints consistently observed in the percent benefit/deployment relationship. Initially deploying travel time surveillance on links with high utilization and high variability, up to the first 20% of the full network, yields roughly a 2% increase in system benefit for every 1% deployed. This can be considered the expected return from an initial deployment targeting key bottlenecks and other trouble spots. In the "build-out" phase between 20% and 80% of full deployment, ATIS system benefits flatten out to a roughly 0.9% increase in marginal benefit for every 1% of the full network that is instrumented to support ATIS. This phase can be characterized as the expansion of ATIS coverage to include all major facilities across the region. The final "completion" phase (between 80% and full deployment) fills in connector links and other roadways to provide end-to-end travel information for users. Here, the marginal improvement drops to 0.4% for every 1% of the network instrumented to support ATIS.

Note that the relationship here is independent of cost and also does not consider marketing. The slope of the relationship in each of the three phases relate to a relative return in terms of total ATIS system benefit, not a ratio of benefits to cost. Figure 6 does not imply that expansion of ATIS coverage should halt when the slope of the function drops below 1.0. The cost-benefit of deploying instrumentation to cover another 1% of the roadway network depends on the accuracy of the travel times generated, the cost per mile of the sensors, and the level of congestion on the roadways.

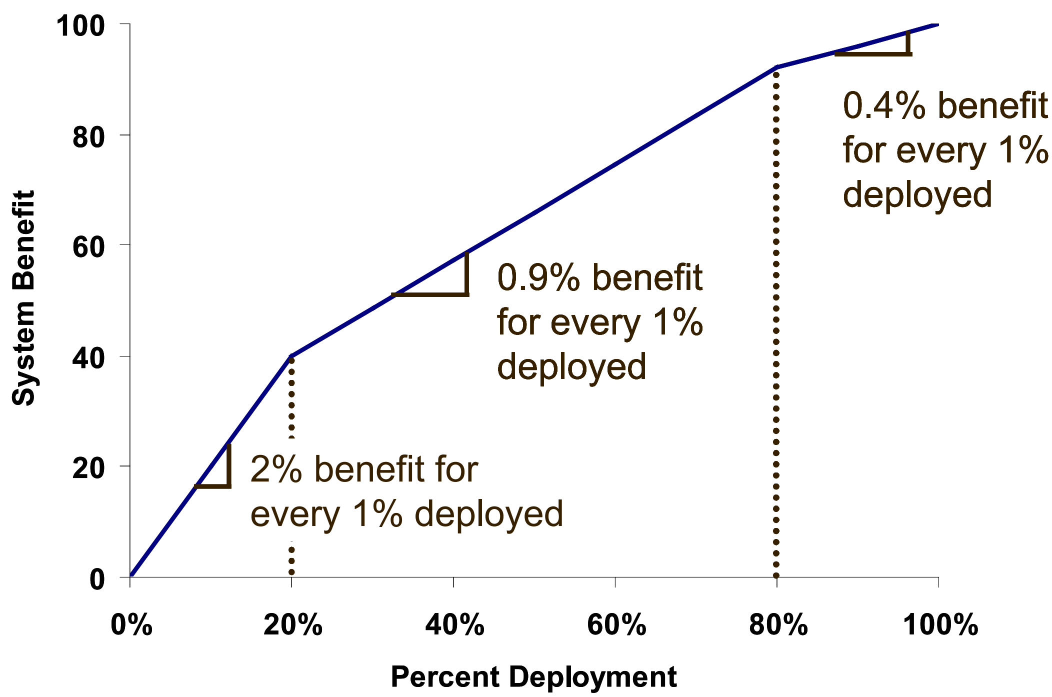

Figure 6. Percent Coverage vs. System Benefit

3.2. Percent Deployment and Costs to Deploy and Maintain

The relationship between percent deployment and cost (Figure 7) has a simpler linear form than the percent deployment and ATIS benefit relationship. Initial costs involve establishing the required underlying infrastructure for travel time data collection, processing and dissemination. This may involve the creation or augmentation of a Traffic Management Center (TMC). After the initial enabling infrastructure is in place, we assume a constant cost of instrumentation for each mile of roadway from the first mile to the last. This cost includes both deployment and operating costs.

The relationship presented here is independent of surveillance technology though the cost per mile will certainly depend on the type of detection system chosen. It assumes that some roadside infrastructure must be deployed to estimate travel times on new sections of the roadway network. In addition, we do not assume any economies of scale (declining unit costs as percent deployment rises) because the incremental deployment of surveillance technologies is likely phased in over several years. The relationship also assumes there is no mid-stream change of surveillance technology or inflation in deployment and operating costs.

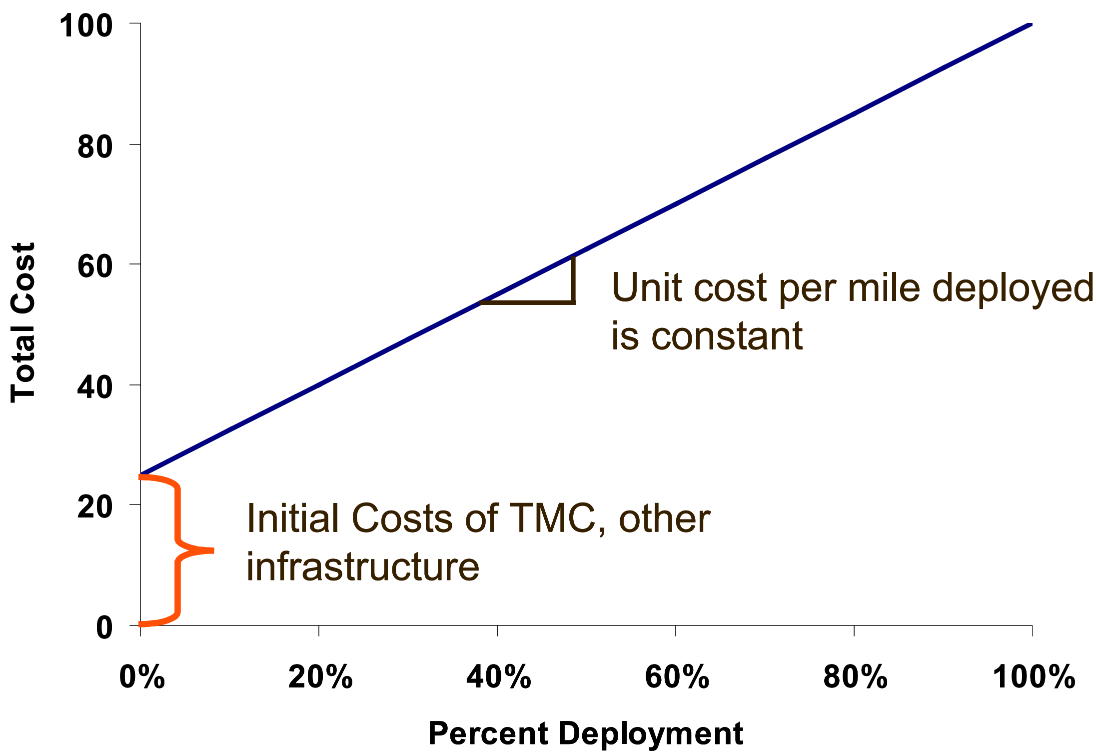

Figure 7. Percent Coverage vs. Total Costs

Once the enabling infrastructure is in place, the cost of providing detection has two components. First, there is the capital cost of cutting the initial loops, purchasing and installing the necessary hardware such as controllers and cabinets, and running telecommunications lines. Second, there are ongoing operations and maintenance costs. Operations costs include leasing of telecommunications lines and maintenance costs include the cost of loop replacement. The cost of a loop detector station according to one source

(11) is given below. The costs of the loops themselves were scaled down to be more representative of nationwide rather than region-specific prices and a five-year expected lifespan.

Description	Cost
Capital Cost	
Hardware Costs	$4,100
Installation Costs	$31,960
Total Capital Cost	$36,060
Maintenance Costs Per Year	$2,940
Operations Costs Per Year	$2,040
Total Annual Cost	$4,980

Table 1. Costs of Loop Detection – One Site, Six-Lane Highway

These costs are for one detector station, which includes six loops over both directions of traffic. If we assume, as is typical, that detector stations are placed ½-mile apart, these costs need to be doubled to reflect centerline miles of surveillance. If we assume a deployment life cycle of 20 years, the capital cost is distributed across that time and becomes $3,606 per centerline mile per year. The operations and maintenance is $9,960 per centerline mile per year, bringing the total cost of a centerline mile of surveillance to $13,566 per year.

As shown in generalized form in Figure 6 as surveillance expands to cover more miles of the network, the marginal benefit decreases. That is, the impact of adding an additional mile of surveillance is greater for new deployments than for mature deployments. And, because marginal cost is constant (i.e., the cost of adding an additional mile of surveillance is the same for new and mature deployments), we can identify the point of diminishing returns as the point where the marginal benefit (which varies) dips below the marginal cost (which is constant). Based on the linear regression strategy in Los Angeles, Figure 8 shows the marginal benefit as a function of network coverage. This figure shows that the marginal benefit to ATIS users of adding an additional mile of surveillance coverage dips below the marginal cost of $13,566 per mile when approximately 35% of the network is covered. This is based on $8 million per year of ATIS user benefit and 367 miles of freeway surveillance under full coverage in the Los Angeles area.

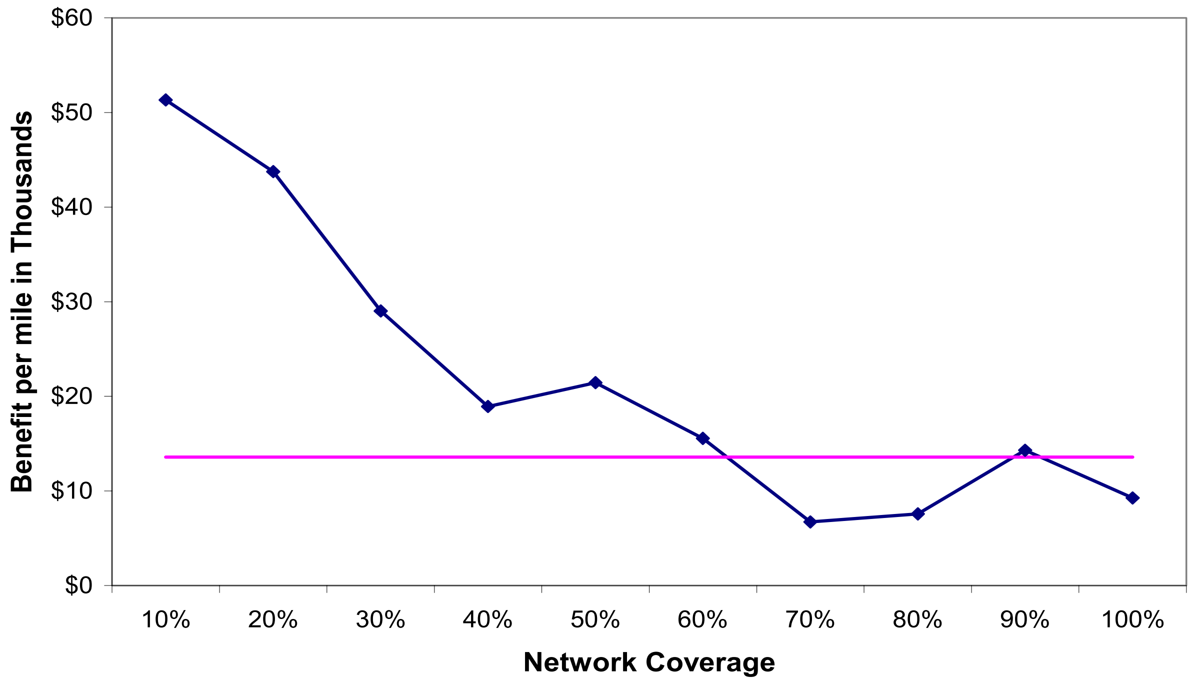

Figure 8. Marginal Benefit Per Mile vs. Network Coverage Level

4. Cost-Effective Next Steps

4.1. Existing ATIS Deployments

The four relationships between cost, benefit, ATIS accuracy and ATIS coverage discussed above examine each dimension of the ATIS investment issue independently and do not consider the effect of interactions. Of course, however, making investment decisions in the real world involves making sense of all four components and their interactions simultaneously. In Figure 9, we bring together what we understand about these interactions to provide a nomograph dealing with cost-effective next steps for decision-makers considering enhancements to an existing ATIS deployment.

On the x-axis we consider the accuracy of travel time estimation; on the y-axis we consider the percent of the regional network currently instrumented. A decision-maker can reference the most cost-effective next step by identifying their current position by accuracy level and percent deployment. For example, an ATIS service covering 20% of the regional network with average travel time estimation error near 15% references the “Expand Coverage” region of the nomograph. This region describes the conditions where the most cost-effective next step is the expansion of ATIS coverage rather than increased investment in improved travel time estimation. The “Improve Accuracy” region comprises the portion of the decision-space where investments in improved accuracy are likely the most cost-effective next step. The third region, “Stand Pat” is the area where additional investment (either in increased network coverage or in improved accuracy) is not likely to be cost-effective.

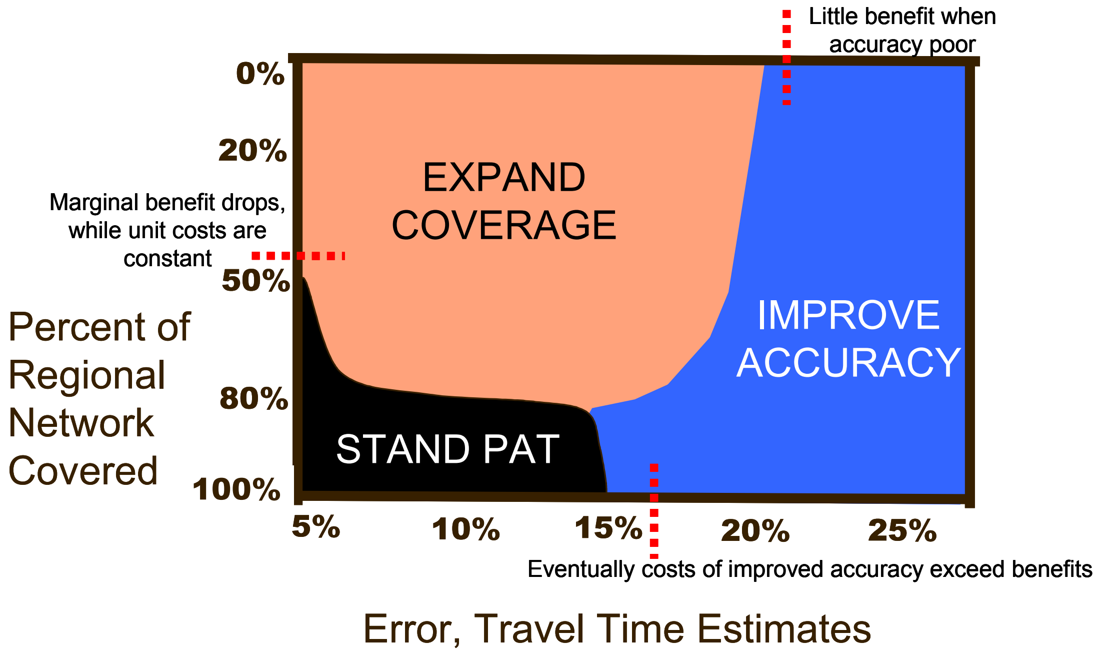

Figure 9. Existing Deployments: Most Cost-Effective Next Step

The key breakpoints on the nomograph mirror the critical breakpoints identified in Figures 3, 4, 6, and 7. For example, along the top of the nomograph in the decision space related to small-scale initial deployments (5-10% of the full network), we see a transition from improved accuracy to expanded coverage in the 18-20% error range. This reflects our concept from Figure 3 that a sizeable number of commuter trips will only find value from ATIS services when error is reduced below 20%. Along the bottom of the nomograph (ATIS near full deployment levels), the benefit to ATIS users of expanded coverage is too low to offset the cost of additional instrumentation. At these levels, however, improving accuracy to the 12-15% range through improved sensor maintenance may be cost-effective. Along the left side of the nomograph, corresponding to a maximally accurate ATIS system, there is little to no benefit to ATIS users in investing to reduce travel time error below 5%. However, building out such a system to a 50-80% level of deployment may be cost-effective.

The nomographs are provided here as a notional guide rather than as a tool to provide precise guidance on optimal deployments. We did not attempt to find precise breakpoints for costs and benefits given that there are many factors that we cannot detail when considering general deployment guidance. For example, we do not consider costs of switching sensor technologies or integrating multiple technologies for surveillance. Likewise, the level of travel time variability and the population using the roadways may cause the "Stand Pat" region to expand or shrink.

4.2. New ATIS Deployments and Projected End-States

With the same concepts of tradeoffs and interactions between our critical factors in the ATIS investment process, we present a second nomograph for decision-makers considering new ATIS deployments in Figure 10. Here the form of the nomograph differs from **Error! Reference source not found.** in that we attempt to define conditions under which ATIS systems should not be deployed from a cost-benefit standpoint. The x-axis and y-axis are defined in the same way as in Figure 9**Error! Reference source not**

found.. Here the decision-maker can consult a proposed ATIS deployment or proposed ATIS end-state and reference its expected accuracy and coverage level to determine its relative return on investment. The mapping provided divides the nomograph up into a "Do Not Deploy" section and a series of contours, each one rising to an eventual optimal end-state. The optimal end-state reflects the ATIS deployment with the highest benefit to cost ratio.

The corners of the nomograph characterize four different kinds of "Do Not Deploy" conditions:

- **Hi-Tech Sandbox**. This deployment has very low travel time estimation error but is too small to provide enough benefit to offset the costs of the start-up infrastructure required to initiate an ATIS service.
- **Inexpensive and Not Beneficial.** This deployment is likely very cheap as it combines both inaccurate travel time measurement and very low levels of roadway coverage. Such a deployment provides so little benefit to anyone that despite the low costs, it is still not cost-effective.
- **Expensive and Not Beneficial.** This is a full deployment of the same inaccurate sensor system that drives the inexpensive and not beneficial case above. Costs are significantly higher because of increased coverage but benefits are still low because of poor travel time accuracy.
- **Overkill.** This deployment covers every facility in the region with exceptionally high accuracy. While such a deployment may be a marvel of technology, the high costs of deploying and maintaining such a large complex sensor network exceeds the aggregate benefit likely realized by ATIS users.

Where we do foresee cost-effective ATIS deployments is in the sweet-spot where we expect deployment and operating costs of surveillance to be relatively low, covering the enough of the major facilities in the network with high enough accuracy so that ATIS usage is widespread and clearly beneficial to users. In general, we define this "sweet-spot" somewhere in the vicinity of 50-60% network coverage with travel time estimation error around 15%. Again, as with the nomograph presented in Figure 9, this tool is notional in nature and true cost-benefit for a particular locality will depend on congestion levels and other key local factors. Clearly, for some localities with low congestion, there may not be a cost-effective ATIS deployment at any level given current surveillance costs. In other highly congested metropolitan areas, the shapes of these contours may be different because of the high-return on more expensive surveillance technologies or expanded ATIS coverage.

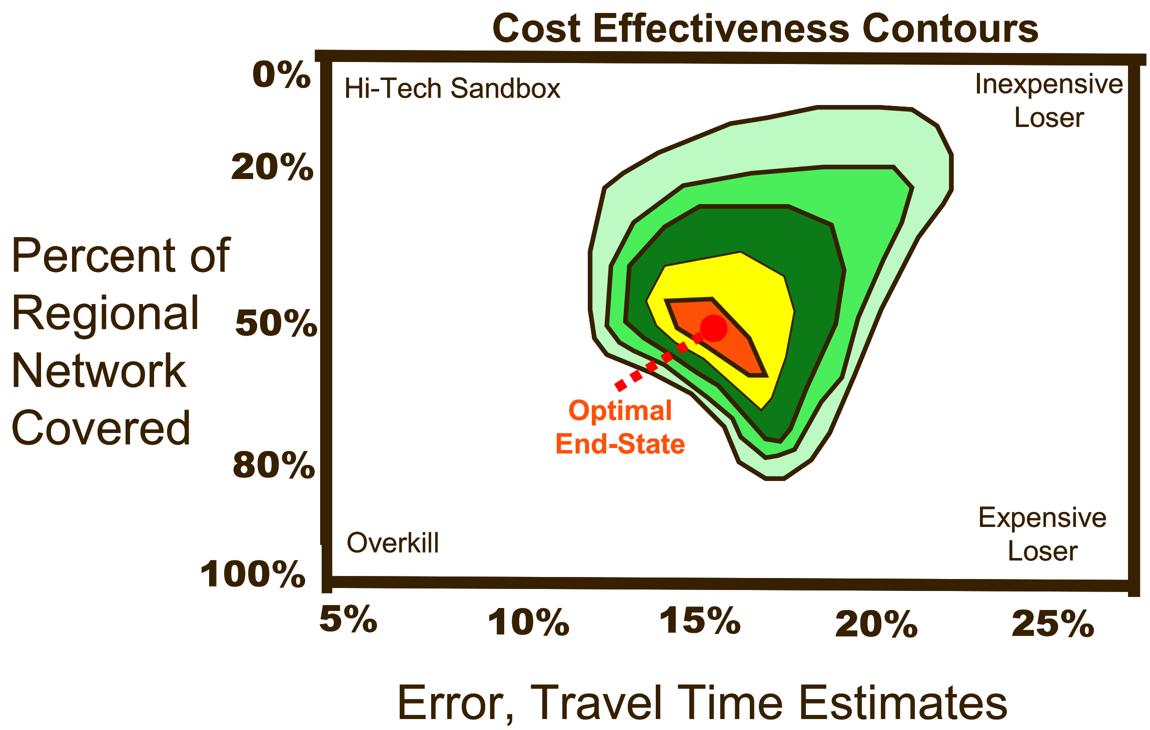

Figure 10. Evaluating New Deployments ("Greenfields") or Projected End-States

Where do current ATIS deployments stand with respect to accuracy and coverage? This question is difficult to answer because of the very few examples we have of ATIS providers independently measuring travel time estimation error. Further, it may not be clear what full-deployment conditions might look like in various cities. A uniform definition of full deployment conditions like those adopted by the Oak Ridge National Laboratory ITS Deployment Tracking program (1) could be utilized to address this issue; however, we have not done so at this point.

5. Future Work

Our high-level treatment of costs and benefits in this section helps to shed light on the complex questions facing decision-makers investing in ATIS deployments. Such a general treatment does not offer specific advice for a locality facing these decisions, however. Some of our key cost and benefit relationships (particularly with respect to the costs of surveillance technologies and expected accuracy) are still primarily notional in nature. In an upcoming series of projects sponsored by the ITS Joint Program office, we will be addressing critical gaps in our understanding:

- **Accuracy and Sensor Spacing**. Mitretek is currently working with the Minnesota Department of Transportation in a joint effort to relate accuracy of travel time estimation as it relates to the spacing intervals of loop detectors on a freeway system. Results from this study will help refine both the accuracy/cost relationship as well as the accuracy/benefits relationship.
- **ATIS Impact on Truck Movements**. Our benefit calculations to date have been restricted to consider only commuter benefits. We plan to estimate aggregate benefits from ATIS provision to urban truck movements as well.

- **Impact of Travel Advisories**. The accuracy spectrum we have defined relates only to the notion that travel time estimates are delivered to the ATIS user. These services differ from typical advisory services where congestion is described using relative terms such as “severe” or “stop and go.” We are currently examining the ability of such services to provide benefit to motorists using a similar technique to the one we use to evaluate more quantitative ATIS services.

One potential extension of this work that we have not yet pursued but would likely be highly revealing would be to work with a metropolitan area and plan out a cost-effective expansion program. This “case study” approach could be conducted either with a locality considering a new ATIS deployment or one considering additional investment in an existing ATIS deployment.

References

1. Gordon, S., Dixson, A., and J. Noltenius. ITS Deployment Tracking: 2002 Survey Results. http://itsdeployment2.ed.ornl.gov/its2002/default.asp. Accessed June 17, 2003.
2. http://www.its.dot.gov/511/PDF/511_overview.pdf. Accessed June 17, 2003.Jung, S., J. Larkin, V. Shah, A. Toppen, M. Vasudevan, and K. Wunderlich. On-Time Reliability Impacts of Advanced Traveler Information Services (ATIS) Volume II: Extensions and Applications of the Simulated Yoked Study Concept. U.S. Department of Transportation, ITS Joint Program Office, March 2002. EDL#13630.
3. Shah, V., Jung, S., J. Larkin, A. Toppen, M. Vasudevan, and K. Wunderlich. Accuracy and Coverage Implications of Advanced Traveler Information Services (ATIS) Benefits, Volume III: Extensions and Applications of the Simulated Yoked Study Concept. U.S. Department of Transportation, ITS Joint Program Office, February 2003.
4. Jung, S., J. Larkin, V. Shah, A. Toppen, M. Vasudevan, and K. Wunderlich. On-Time Reliability Impacts of Advanced Traveler Information Services (ATIS) Volume II: Extensions and Applications of the Simulated Yoked Study Concept. U.S. Department of Transportation, ITS Joint Program Office, March 2002. EDL#13630.
5. Fariello, Brian. Email correspondence. Texas Department of Transportation. May 2003.
6. Quiroga, C. Assessment Of Dynamic Message Sign Travel Time Information Accuracy. North American Travel Monitoring Exhibition and Conference (NATMEC). Middleton, Wisconsin, August 2000.
7. Park, D. and L. Rilett. Forecasting Multiple-Period Freeway Link Travel Times Using Modular Neural Networks. Presented at the 77th Annual Meeting of the Transportation Research Board. Washington, D.C. January 1998.
8. Sun, H., Liu, H., Xiao, H., He, R., and B. Ran. Short-Term Traffic Forecasting Using the Local Linear Regression Model. Submitted to the 82nd Annual Meeting of the Transportation Research Board. Washington, D.C. January 2003.
9. Zhang, X., and J. Rice. Short-Term Travel Time Prediction Using a Time-Varying Coefficient Linear Model. University of California at Berkeley. http://stat-www.berkeley.edu/users/rice/xTTfinal.pdf. Accessed May 2003.
10. Lin, Peter. Email correspondence. California Department of Transportation. November 2003.
11. Mouskos, K., Niver, E., Pignataro, L., Lee, S., Antonious, N., and L. Papadopoulos. TRANSMIT System Evaluation Final Report. Institute for Transportation. New Jersey Institute of Technology. June 1998.
12. ITS Benefits and Costs Database. Federal Highway Administration. Washington, D.C. http://www.benefitcost.its.dot.gov/its/benecost.nsf/SubsystemCosts?OpenForm&Subsystem=Roadside+Detection+(RS-D). Accessed December 2003.
13. Jung, S. A Study on Effective ATIS Deployment Plans: Implication for ATIS Investment Strategy. U.S. Department of Transportation, ITS Joint Program Office, Forthcoming, 2004.
14. Peirce, S. and J. Lappin. Acquisition of traveler information and its effects on travel choices: evidence from a Seattle-area travel diary survey. U.S. Department of Transportation, ITS Joint Program Office. April 2003.

Appendix: Extrapolating Per Trip Benefits to Regional Benefits

The HOWLATE methodology only allows us to estimate ATIS benefits on a trip by trip basis. However, in order for us to perform cost-benefit comparisons for ATIS in a metropolitan region, we need to compare system-wide costs and system-wide benefits. To estimate system-wide benefits of ATIS we need an estimate of the number of ATIS users, which we can multiply by the average user benefit to give us a measure of system-wide benefit. Here we present calculations of this value for Los Angeles in two different ways: one using PeMS data and the other using data on Los Angeles area Internet ATIS usage from the Traffic and News Network (TANN).

Method 1: PeMS Data

According to the PeMS database, the average number of vehicle-miles traveled daily in Los Angeles is approximately 100 million. From the HOWLATE analysis, the average trip length in Los Angeles is 23.71 miles. This gives an average of 4.2 million daily trips, according to the calculations below.

No. daily trips = 100,000,000 vehicle-miles per day / 23.71 miles per trip
No. daily trips = 4,217,630 trips per day

Based on a study of ATIS usage in Seattle, another large market for traffic information and a region with severe congestion delays, approximately 1% of the traveling public uses ATIS (14). Therefore, if we assume an ATIS market penetration of 1% in Los Angeles, that gives us 42,177 daily users of ATIS.

No. daily trips using ATIS = 4,217,630 * 1% = 42,177 ATIS trips per day

If we assume that only those trips that benefit will use ATIS, we can apply the average trip benefit for trips that benefit. At an error of 20%, this is $0.96 per trip. Therefore, the daily regional ATIS benefit is $40,489 per day.

Daily regional ATIS benefit = 42,177 * $0.96=$40,489 per day

If we assume 200 commuting days per year, the annual regional benefit is approximately $8.1 million per year.

Annual regional ATIS benefit = $40,489 * 200 = $8,098,000 per year

Method 2: TANN Data

Based on data from The Partnership, the Travel Advisory News Network (TANN) receives 3,600 hits per day. If we assume they have 10% of the ATIS market in Los Angeles and that each user hits a web page only once, there are 36,000 ATIS users per day. This is not far off from our earlier estimate of 40,500. The Partnership also estimates there are 6 million freeway trips made daily in Los Angeles. That is slightly more than our earlier estimate of 4.2 million. According to these numbers the market penetration for online ATIS services in Los Angeles is 0.6%. By these numbers:

Annual regional ATIS benefit = 36,000 * 200 * $0.96 = $6,912,000 per year

These are rough figures based on a number of assumptions and are simply meant to be used to get us within the ballpark. The fact that the benefit numbers are indeed in the same ballpark by the two separate approaches lends them some credibility.

The benefits from HOWLATE are based on the premise that commuters need to be on time. Of course, this is not always the case. There are also commuters who have flexibility in their departure and arrival times, whose objective is to avoid congestion and reduce trip time. Whether these and other commuter types benefit more or less than $0.96 per trip would increase or decrease the regional benefit figure accordingly.

One missing element in our analysis of system-wide benefit as a function of accuracy is how the number of users is affected by varying accuracy. Assuming traveler information is like any other product of service, it is logical that as the quality of the product improves it will attract more users. Given that the market for ATIS travel time provision is still in its infancy, it is not known how sensitive usership is changes in accuracy. For our purposes, we will assume constant usership within a reasonable range of accuracy.