Demo Abstract: PCP: The Personal Commute Portal

Hari Balakrishnan, Nikolaus Correll, Jakob Erikkson, Sejoon Lim, Samuel Madden, and Daniela Rus
MIT Computer Science and Artificial Intelligence Laboratory
The Stata Center, 32 Vassar St., Cambridge, MA 02139
{hari,nikolaus,jakob,sjlim,madden,rus}@csail.mit.edu

ABSTRACT

The Personal Commute Portal (PCP) is a Web-based traffic information system that provides a good driving direction and personalized route recommendation using historical and real-time traffic data obtained by a vehicular sensor network.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous

General Terms

Design, Algorithms

1. OVERVIEW

Traffic congestion is a global problem that results in tremendous personal and social costs. Drivers go along familiar routes every day without understanding their travel costs. Although several Internet services provide driving directions, their level of information about the roads is sparse and not up-to-date, and they do not include any kind of personalized recommendations based on users driving habits and patterns.

Our research objective is to provide an effective mobile-sensor network based personal commute portal that includes a set of services that analyze a user's commutes and recommend changes to driving habits that will reduce fuel consumption, time spent in traffic, and driver frustration. This builds on our previous work on CarTel [3] (`http://cartel.csail.mit.edu`) which focused more on system software for capturing, storing, and delivering data from a variety of mobile sensors rather than the personalized driving that relates to other users and to traffic patterns discovered by our system.

The Personal Commute Portal (PCP) has three key components: first, an embedded node that continuously monitors car's location and opportunistically transmits that data off of cars using open WiFi networks. Second, a route planning web site that uses historical and real-time traffic data to determine good driving directions that take traffic delays into consideration and that is personalized to a user's own driving habits. Third, a set of recommendation services that allow users to visualize and analyze their driving habits and compare their drives to friends and other users, including a "leaderboard" that shows which users have shown the greatest reduction in overall driving time and fuel consumption. We describe these three components of our demo in the remainder of this paper.

SenSys'08, November 5–7, 2008, Raleigh, North Carolina, USA.
ACM 978-1-59593-990-6/08/11.

2. DATA COLLECTION FOR TRAFFIC ESTIMATION

Over the past three years we have developed and deployed a data collection system on a fleet of taxis in the Boston area (this system currently continuously runs on about 30 cabs); each car carries a wireless embedded computer and a collection of sensors, including GPS. The result is a mobile sensor network capable of obtaining data about a large metropolitan area at relatively low cost, compared to current approaches that deploy traffic sensors in roads. In the current implementation, data is transmitted off of these cabs using an opportunistic WiFi system called Cabernet [2], that rapidly associates with open WiFi access points in the Boston area. We have gathered many gigabytes of traffic delay information over the past months, covering several thousand hours of driving.

Converting this data into actionable statistics that can be fed into a shortest-path planning algorithm is a challenging database problem: first, data volumes are quite high, so efficient algorithms and indices are needed. Second, the data is noisy, so algorithms that match GPS points to an underlying road database in a robust and error free way are essential, as discussed in prior work [1]. Third, determining how to segment the data into a collection of road segments is an interesting and challenging problem: simply computing statistics on a per-intersection granularity leads to a huge road graph, which substantially slows performance. Our demo includes implementations of all of these components.

3. TRAFFIC PORTAL

The goal of the traffic portal is to provide an effective navigation system for cars that uses historical and real-time traffic data to determine good driving directions that take traffic delays into consideration. By collecting traffic delay data and categorizing it according to location (road segment), time of day, day of week, and time of year, as well as correlating it with the occurrence of various events (e.g., concerts, sporting events, etc.), we can model the statistics of traffic delays on road segments. These statistical models will then be used for two purposes: first, to identify the locations and times of traffic "hot spots" that drivers should avoid, and second, to use in new algorithms for computing efficient paths according to various delay criteria. In particular, our interest is in developing algorithms that use delays rather than simply distance.

Our traffic portal runs on data that is continuously collected by the CarTel system and processes it for input into a delay-aware navigation system for vehicles. The data analysis engine compiles a historical traffic conditions and generates statistical models of this data. In contrast to existing methods for measuring traffic flow and performing route planning and prediction, the approach we will demonstrate is novel in two key ways. First, by using a mobile sensor network deployed on taxis that drive many hours each day, we are able to develop, measure and model delays directly, in contrast to road sensors that usually only obtain information about how many cars are driving past a location in a time interval (flow information). Since users care about delays and not flow rates, measuring it directly is beneficial.

Second, no current system allows users to find good paths according to various delay criteria, such as finding the route with the minimum expected travel time at a particular time of day based on historical travel data, or finding the route with the maximum probability of reaching a destination within a given deadline. Our system can explore both types of routes (in addition to traditional shortest-path routes recommended by current routing tools.)

4. COMMUTE ANALYSIS

In addition to using cabs and user's cars to measure traffic information, our system provides a personalized interface that allows users to browse their own historical driving patterns. Our system is currently in daily use by seven members of our lab, in addition to the 30 taxis (which are primarily used to capture traffic information.) Key features of the interface are:

1. A tool to compare the expected drive times of different routes that the user commonly drives at different times of day, including tools that show how total commute time would vary if the user were to leave for or from work a minutes earlier or later.
2. A tool to break down driving time, distance, and fuel consumption by time of day, day of week, and hour of day, and to compare how long different trips actually take depending on their start time.
3. A "competitive commuting" tool that compares two friends' driving habits and that provides a leaderboard showing drivers who show maximal improvement in their driving habits over time.

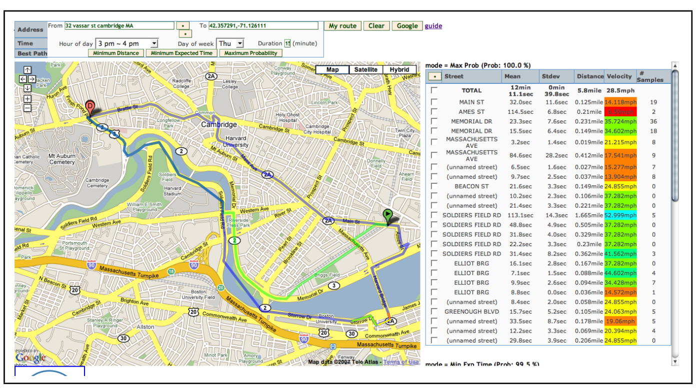

Figure 1: Demonstration traffic portal. The figure shows different optimal paths from a start location (the green arrow) to a destination (the red "D"), according to the three different criteria: the route recommended by Google's Maps (the topmost route), minimum expected time (the middle route), and maximum probability of arriving by the deadline (the bottom route). Here, our system estimates that Google's route – while 3.3 miles – will take almost 16 minutes on a Thursday afternoon, whereas our system's minimum expected time route takes only 11 minutes and 45 seconds, even though it is 4.7 miles in length.

5. THE DEMONSTRATION

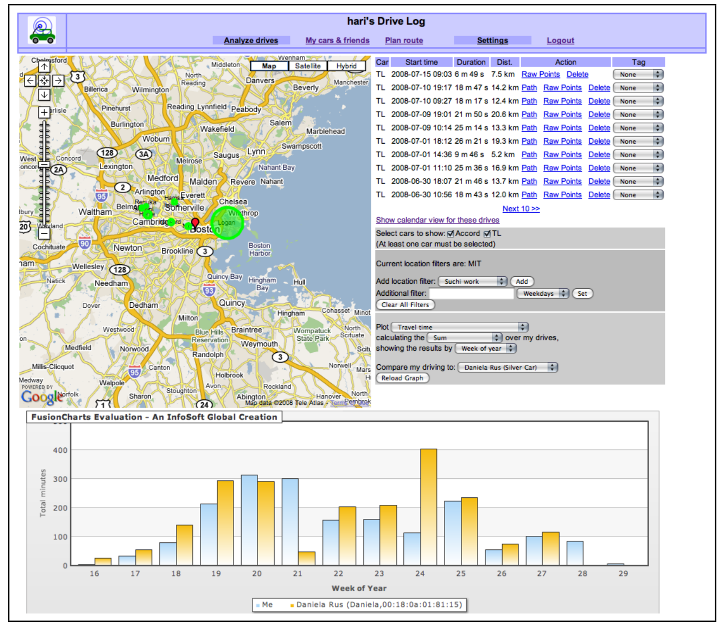

Figure 2: Demonstration commute portal interface. Drives are shown on the right, and map showing the user's favorite locations (green circles) as well as the start and end of routes as the user mouses over them is shown on the left. On the lower right, the user has controls to select drives of interest, such as those that go from work to home. At the very bottom is a plot comparing the user's drives to those of a co-worker on a week-by-week basis.

We will demonstrate PCP, which is a browser accessible web application, with data from our taxi-cab deployment in the Boston area as well as from a number of individual drivers. Our demonstration will include:

1. A live view of the current data from cabs (showing their position in Boston, as well as a composite "heat map" visualization of the aggregate data collected by the cabs.
2. A web service that allows users to select start and destination routes and observe paths subject to a variety of delay criteria. The interface includes the ability to compare predicted routes to those produced by the Google Maps software (see Figure 1), as well as see how total travel time for different routes would vary if the user were to shift his or her driving a few minutes earlier or later. PCP gives personalized route recommendations from the traffic portal using their own driving patterns as a guide.
3. A demonstration of our "commute portal" that allows users to browse their drives, compare their driving patterns to other users, and measure their own fuel consumption and carbon footprint (Figure 2).

6. REFERENCES

[1] S. Brakatsoulas, D. Pfoser, R. Salas, and C. Wenk. On map-matching vehicle tracking data. In *VLDB*, 2005.

[2] J. Eriksson, H. Balakrishnan, and S. Madden. Cabernet: A wifi-based vehicular content delivery network. In *Mobicom*, 2008. To appear.

[3] B. Hull, V. Bychkovsky, Y. Zhang, K. Chen, M. Goraczko, A. Miu, E. Shih, H. Balakrishnan, and S. Madden. Cartel: a distributed mobile sensor computing system. In *SenSys*, 2006.