Real-Time Rider Information

Point of contact: Aaron Steinfeld, steinfeld@cmu.edu, 412-268-6346

Project summary

This project is focused on supporting transit riders and the Port Authority. Riders want to reliably get from A to B at time T. They want low cost and to enjoy the experience. Port Authority wants to increase ridership, lower costs, lower fares, and remain a premier transportation system. A key functionality to make progress on these goals is to provide riders and agency staff with real-time rider information about the arrival time and fullness of nearby buses. This work seeks to integrate several existing pilot efforts to develop this functionality at a low cost with existing work in the delivery of information to riders.

Problem definition and the desired results

The team seeks to collect and disseminate real-time information about bus arrival times (based on current bus location) and fullness to Port Authority riders. The target outcomes are accurate and easy to use interfaces that are valuable to both riders and Port Authority personnel. For example, Ben is leaving his house on Beeler to see his doctor at UPMC Montefiore. It is 10 AM and cold outside. He's not sure when the next bus is coming or if there is room for him to get on. He makes a quick call to Port Authority, or looks at his smartphone, and learns there is a fairly empty 61A coming in ten minutes. He waits five minutes before braving the cold.

There are two parts to this project. One is to combine several sources of information in order to make a precise as possible prediction of arrival time for the next bus. The other is to deliver this information to riders in an easy to use and accessible manner.

Sources of Information

Members of the team have been working towards collection of arrival information through three methods. First, one group has been tackling the problem directly by trying to develop automatic vehicle location (AVL) technologies at low cost using a mobile phone platform. Two other groups have been working on complimentary crowdsource methods for gathering AVL data from the riders themselves, also using mobile phone data. While basically free to the transit agency, this is more difficult in that it is necessary to determine whether a rider is on a bus or not. The two methods under development are sensing rider transportation mode from existing sensors in the phone (e.g., GPS, accelerometers) and obtaining self-report data from riders through a user interface on their mobile phone. It is important to note that the crowdsource information provides extremely valuable data to the transit agency since it documents the location of both rider entry and exit from the system. Existing systems, such as smart fare boxes only document the location of people entering the system.

Delivery of Information

Using a rider's mobile phone as a communication platform allows the presentation of tailored information about real-time bus arrivals (e.g., Figure 1). However, not all riders are likely to participate in supplying crowdsource data to the system. Providing information to these riders through their mobile devices is important since dynamic signs at bus stops are expensive and most riders need this information before leaving their home or workplace. Calls to Port Authority are a popular method for accessing such information. For this population group, the team will modify the backend of the existing Let's Go! system, a speech recognition based call-in system developed by members of the team which has been deployed since March 2005 on the Port Authority phone line in the evening and on weekends. The information given to one caller might result in, "The next 61C is scheduled to arrive at 1:14. Carnegie Mellon Ride estimates this bus will arrive at 1:16 and be 75% full."

Need and potential for involvement of public or private partners

Figure 1. Example smartphone concept

Accurate predictions of real-time bus arrival times, in place of scheduled arrival times, has been shown to be very important to riders. The team has collected a variety of evidence from Port Authority riders on this topic and the vast majority of calls to Port Authority's customer service line are queries about bus arrival times. This evidence is supported by a variety of results from other agencies suggesting anywhere from an 8% system-wide to a 40% increase on specific routes in ridership when routes have real-time predictions. Knowledge about arrival times permits riders to avoid waiting – a feature particularly important to riders with disabilities and riders who have the option of driving their own car. According to Port Authority 2008 budget figures, a 1% increase in ridership corresponds to $500,000 increase in rider fee revenue. Automated provision of such information to riders also reduces load on human operators at the call center as well as a shorter wait time for callers. Collection of paired entry-exit data supports better system-wide route and load planning and opportunities for transfer predictions (e.g., eight riders on bus X transferred to bus Y yesterday at this time, hold bus Y until X arrives). A mobile communication platform with riders also allows transit agencies to disseminate information about dynamic rerouting (e.g., real-time detour info during parades, road-work, etc.). Currently this dissemination is done in person. This mobile platform can also be utilized to disseminate emergency information for citywide emergency events.

Current AVL systems are expensive, both in capital and operating costs. The installation of a commercial AVL system on the Port Authority's approximately 1,000 buses would cost millions of dollars. Commercial systems are also narrowly focused. They do not usually support paired entry-exit data on riders, nor do they offer a method to communicate directly with riders. We propose a low-cost approach that leverages a free and already installed infrastructure – the mobile phones that nearly all transit riders currently carry. By leveraging this resource we gain access to the GPS units within many of these phones, and we can leverage their displays to feed riders personalized information without the cost of installing and maintaining expensive, electronic screens at every bus stop. We are particularly interested in interoperable, modular, technologies that support combinations of multiple information sources and provide a low-cost avenue for third party innovation. These mechanisms can be in the form of schedule data (e.g., http://www.bart.gov/schedules/developers) and opportunities for new functionality (e.g., Let's Go! spoken dialogue testbed). While the specific effort here is aimed at Port Authority, we also envision other transit agencies re-using and composing system components as appropriate to their community with the same delivery mechanisms and low-cost model.

In terms of partnership with public and/or private partners, the team has already briefed Port Authority on the goals of this project and solicited valuable feedback from the agency. This is only the latest contact between the agency and team members – the Let's Go!/DialRC project, RERC on Accessible Public Transportation, and Heinz Systems course have all been in steady contact with the agency for their respective work.

Members of the team have also been in contact with potential industry partners and will seek their involvement, through monetary contributions, access to data and/or technology, and/or donated engineering time. For example, the RERC-APT and Let's Go! have recently obtained access to the Google Transit feed for Port Authority and are in contact with the Google Transit development team. The wireless industry is another potential partner (see Tim McNulty).

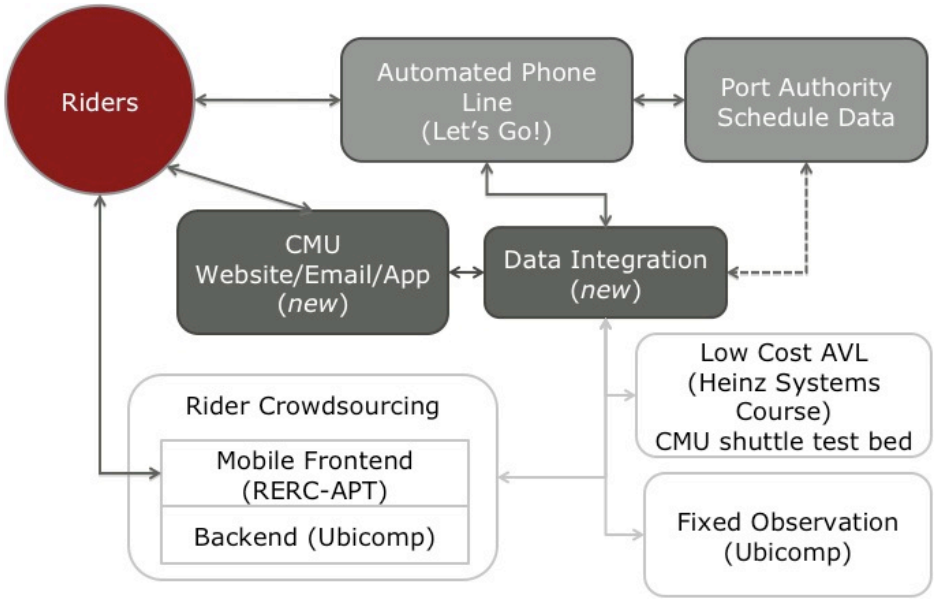

Figure 2. Integration of existing and new components

Outline of major tasks and rough timetable

The project is designed in a modular form. The minimum level of work consists of the *Base* project. Additional options can be incorporated as desired or funding becomes available. The rough timeline for the Base project and options is provided in Table 1 at the end of this section.

Base project

The team will integrate a proof of concept system (Figure 2) and deploy it for beta testing within the Oakland Forbes-Fifth corridor for the 61 series buses. Coverage will be from the Craft Ave to Margaret Morrison. The base system will be delivered via a website, mobile client, and the Let's Go! system to callers to the Port Authority. The mobile client will be implemented on the Android platform.

The team will utilize existing work in the associated projects to ensure that the system is easy to use and continue to find methods for student involvement beyond research funding. Students have already made significant contributions through coursework and independent studies. For example, two HCII course teams and a Design student have already conducted background research and interaction design on the Mobile Frontend and Backend (Figure 1). Similarly, a team from the Heinz Systems course has conducted the initial work towards probe data from Low Cost AVL technologies.

Scope Option

Under this option, the team will either increase the geographic territory for the 61 series (e.g., Market Street to Murray Avenue) or add in additional lines in the Base corridor (e.g., 59U, 67 series, 44U, etc). This is intentionally a small increase in scope since increasing too much could lead to inaccurate expectations from the general public.

Robustness Option

The Base level is intended to demonstrate basic functionality on phones owned by the research team. The robustness option focuses on software reliability and support for downloading and installation into phones owned by the general public. This option will involve more aggressive software engineering and support for non-research phones. We will also allocate effort to port a reduced functionality mobile client to the iPhone platform. Unlike the Android platform, technical restrictions of the iPhone requires the user to

keep the application running when contributing data since the iPhone does not support background logging of location or internal sensors.

For the Let's Go system, the backend schedule lookup will be modified to accept the second source of scheduling information, the necessary speech synthesis prompts will be added and the speech generator and dialogue manager will also be modified in order to accept these changes. We will ensure that the quality of service will remain the same for the callers.

Evaluation Option

Limited data collection will occur under the Base project in order to document functional success of the proof of concept (especially in the precision of the predictions of time and fullness) to help justify pursuit of additional funds. This option will permit more in-depth data collection and analysis of rider and agency cost and benefit. Team members will (a) conduct human participant experiments including measurements of rider satisfaction before and after deployment, interviews with selected users to better understand when and how they use the system, and an analysis of the logged data detailing how often and at which locations users of the system provide data and request data; and (b) investigate the policy implications and potential impact of merging the official schedule information provided by Port Authority with the unofficial arrival time on agency operations.

Table 1. Rough timeline for the project (2010)

Task	Jan-March	April-June	July-Sept	Oct-Dec
Base Project				
Software architecture	X			
Implementation	X	X	X	
Testing & refinement		X	X	
Demonstration & documentation			X	X
Scope Option	X	X	X	X
Robustness Option				
Software engineering	X	X	X	X
End user installation		X		
Port mobile client			X	
Evaluation Option				
Develop Protocol	X			
Collect Data		X	X	X
Analyze				X

Planned Project Team

As implied in the material above, the team consists of a group of researchers already working on related components under other funding. These are:

RERC-APT	Aaron Steinfeld, Anthony Tomasic, John Zimmerman
Ubicomp	Anind Dey
Let's Go!	Maxine Eskenazi, Alan Black
Heinz	Robert Hampshire

This will be a combined effort but two members of this group will take on coordination roles. Anthony Tomasic, the Director of the MSIT-VLIS profession master degree program in Very Large Information Systems, will serve as the lead integrator due to his direct experience with software engineering and issues of scale. Aaron Steinfeld will handle internal coordination and act as a centralized point of contact when needed. This is an extension of his exiting coordination duties with Dey, Tomasic, and Zimmerman through the RERC-APT and Quality of Life Technology ERC.

For the *Base* project, effort from existing student and staff team members within Let's Go! will work on integrating output from the Data Integration component. The Data Integration and non-voice user

interfaces will be developed by programmers within the RERC-APT. For each of the options, additional effort will consist of:

Scope: greater effort for Base developers, additional Ubicomp effort
Robustness: additional members from component teams
Evaluation: existing HCII, LTI and Heinz students

Note that the team will be contributing effort through other projects as a result of their current research programs. For example, interface and interaction design will be developed under the RERC-APT. Sensing technology will be developed by Dey's Ubicomp team and the RERC-APT. Probe data will be provided by Hampshire's Heinz Systems Course.

Cost estimate

The estimated Total Direct Cost for the Base project is about $75,000. This includes a month of time for Tomasic, another month of faculty time for the Let's Go! team, 3 months of programmer time from RERC-APT/Ubicomp, 5 months of Let's Go! graduate student/post doc time, and approximately 1,000 hours of hourly student programmer time for RERC-APT/Ubicomp.

The Scope and Robustness options would be an additional $35,000-40,000 TDC each, each reflecting an increase in 50% effort for the team. The specific distribution within the team will alter for each of these due to shifts in technical emphasis. Note that further increases in Scope would lead to higher costs. The Evaluation option would be approximately $45,000 TDC, reflecting a month of time from Zimmerman and Hampshire (each) and approximately 1,000 hours of hourly student time to assist with data collection and analysis. LTI graduate students working on Let's Go! will gather caller satisfaction information. We would also recruit student teams from existing HCII and Heinz courses to assist in this effort.

Laying the groundwork for future R&D funding

We are quoting shoestring budgets and additional funding from external sources would improve impact and deployment due to ability to instill good software engineering practices, greater emphasis on usability and ease of use, and more community outreach. This proposal embodies the kernel of a larger effort focused on implementation and evaluation of a deployment. Ideally, we envision $500k - $1M for a large rollout and early testing, with an increase to $2M for a full-blown field operational test study.

Our vision of the future involves deploying three variants: Foreground AVL, Background AVL, and Combination AVL (both foreground and background). The former involves systems that require user action, Background are systems that do not require user action, while the latter is a combination of the two. This permits direct evaluation of each approach and their interaction.

In the long term, this project is contributing to the solving the general problem of a *rendezvous* between people. That is, transportation from a rider's point of view is simply a means to an end: to arrive at the same location at the same time as someone else. In fact, in interviews of riders discussing their needs, some riders have mentioned that they coordinate meeting *on a bus* as part of a trip.

Potential sources of financial support other than Traffic21

We see numerous potential funding targets. For example, opportunities at the National Science Foundation include CISE IIS Medium & Large (all three variants), Social Computing (Foreground & Combination), and Cyber-infrastructure (Background & Combination) grants. This work is directly relevant to the UTC effort underway and efforts in US DOT RITA and FTA. PennDOT has shown interest in advancing support for public transit and is also an option. Finally, we think there are numerous companies with active interest in rider information systems (e.g., Google, ITS technology suppliers, etc).