PORTAL: Implementing a New Generation Archived Data User Service in Portland, Oregon

Robert L. Bertini, Steve Hansen, Spicer Matthews and Andrew Delcambre
Department of Civil & Environmental Engineering
Portland State University P.O. Box 751 Portland, OR 97207-0751
Phone: 503-725-4249 Fax: 503-725-5950 Email: bertini@pdx.edu

SUMMARY

The objective of this paper is to describe the implementation and status of the recently developed archived data user service for the Portland, Oregon metropolitan area. Included in this paper will be a discussion of the hardware and software implementation used in building this system. Also, an update to our future design plans will be included. Along with the implementation details the reader will be given a short tour of the user interface. Several of the above points are complemented with comparisons of other systems that are being implemented elsewhere.

INTRODUCTION

The Portland State University (PSU) ITS laboratory features roadway loop detector data that is streamed to servers located at PSU and than archived in a relational database. This system is know as PORTAL, the Portland Archived Listing. Once the data is retained, many different agencies can and will benefit from the archived data. Such agencies include transportation planners, administration research by key stakeholders including metropolitan planning organizations (MPOs), state transportation planners, traffic management operators, transit operators, and transportation researchers.

PORTAL'S ARCHIVED DATA USER SERVICE DESIGN

In July 2004 PSU and the Oregon Department of Transportation (ODOT) "went live" with their partnership to share loop detector data. Currently, the data come from 436 inductive loop detectors installed in the roadways around the Portland metropolitan area. These data are transferred over a secure direct fiber optic connection between ODOT and PSU. Every 20 seconds XML data is pulled down from the ODOT servers to the PSU archiving servers. Once retained at PSU, the 20 second data is inserted into a PostrgreSQL relational database management system (RDBMS) running on a Fedora Core Linux system. The 20 second data are not altered in any way. Each day at 3 a.m., the 20 second data is aggregated to 5 minute data. This data is appended to a table for the relevant month. The data is also aggregated to 1 hour and appended to a table that contains data for an entire year.

The tables in the RDBMS are dumped nightly to text files where the data is compressed and shipped to a backup server that is located several bocks from the main archiving server. It was felt that having the backup server off-site from the main archiving server was best to ensure total safe keeping of the archived data. In addiction to the backup server a yearly backup is made to a DVD-ROM.

PORTAL'S USER INTERFACE

As with other ADUS systems, such as California PeMS (http://pems.eecs.berkeley.edu/Public/), the user interface of choice is a web interface. It is assumed that anyone interested in the archived data has access to the Internet and has the basic understanding of using a web browser.

Since agencies may want to download the data into their own data processing system PORTAL has the feature of downloading all queried data into a comma separated text file. The Portland ADUS web interface (http://portal.its.pdx.edu) has a username and password login. An interested user registers for an account and then a member of the PORTAL team approves the account. With approximately 60 users at present ,the system is growing in popularity.

FIGURE 1 User Interface

Users login into Portland's ADUS for many different reasons. Some are looking for raw data that they can import and process with their own processing algorithms. Others are looking to the ADUS to organize the data in a fashion that meets their interest. From the web interface different specialized queries of the RDBMS can be done from just a few clicks of the mouse. For example, if a user only wanted to see the loop data from a certain loop detector on a certain day this is possible with very little effort.

Many users are not interested in seeing large amounts of data they want to see graphs of the collected data. The graphs are produced using the open source tool Gnuplot. For example, Figure 1 shows a sample speed plot for one freeway segment. On the data fidelity page a user can see a line graph showing the percentage of the data that is felt to be faulty. Such faulty data can be caused by many different reasons with malfunctioning loop detectors as the primary cause. This is useful to an organization such as ODOT. An employee of ODOT can log in and see that a loop detector has gone bad and send someone out to examine why. In a real-time system where data is discarded it is not as easy to spot broken loop detectors. Portland currently uses the SWARM and Portal algorithms for detecting faulty data. As another example, Figure 2 shows the archived weather data for one day, for both hourly temperature and hourly rainfall.

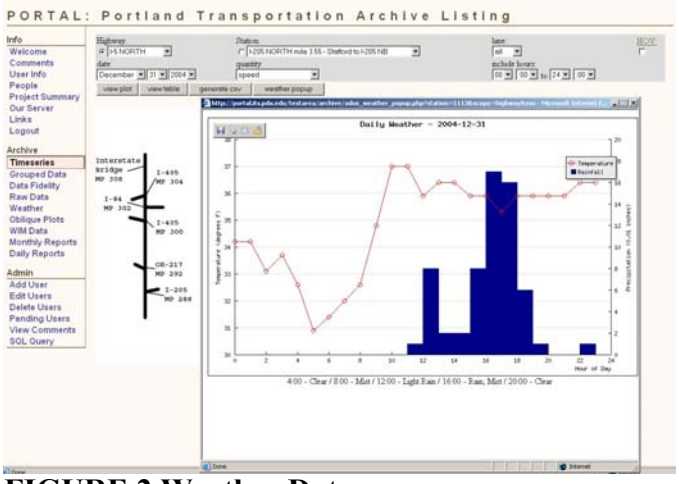

FIGURE 2 Weather Data

PERFORMANCE OF PORTAL

The RDBMS stores data on a redundant array of independent disks (RAID) providing both high-speed access and increased reliability through redundancy in the event of hardware failure. The technology chosen to implement this system is completely open source and has been proved to be the top quality technology in industry currently. Portland's ADUS uses an Apache web server with processing powered by the PHP hypertext preprocessor module.

Along with having the fastest software and top of the line hardware Portland's ADUS was written with optimized code. Much of the code written was written with performance in mind. Having such tremendous data to process every bit of optimized code helps in a noticeable fashion.

LOOKING TO THE FUTURE

There are many other transportation resources that we are interested in archiving. Some planned data sources are: TriMet bus data from the onboard GPS data loggers, Automated Traffic Recorders along different sections of the freeways in Oregon, and Weigh in Motion weight stations around Oregon that record information about trucks as they are weighed around Oregon.

The PORTAL team is beginning to implement some GIS features into the system. PORTAL currently has graphing capability to produce color maps from the timeseries and grouped data pages, and the next step is to display these data on a digital map. Some applications of this technology might include having the average volume per segment displayed as varying line thickness on a map of the system. Other ideas include having a click-able map that would link the user to the page for that station. One also might be able to click on the map at two locations and get the travel time between the two points.

The current loop controller firmware does not support vehicle length calculation. There is a firmware upgrade in progress by ODOT that would implement this feature. When this happens, PORTAL will be expanded to allow queries involving minimum or maximum vehicle length. This would enable planners to limit calculations only to long vehicles such as tractor-trailers. Some researchers are very interested in these statistics, another reason why the Weigh in Motion data would be so useful. Further travel time calculation improvements are also underway that will take into account measurements taken from adjacent freeway detectors.

Finally, the PORTAL system has implemented the calculations required to prepare the Federal Highway Administration's Urban Congestion Report (UCR). Figure 3 shows an example of the UCR for October 2004 with a row added for Portland. We expect in the future that the FHWA can simply access the PORTAL site and request a summary report for each month for the Portland metropolitan region. In the future we also intent to validate the congestion measures using probe vehicles as well as travel time data reported by the TriMet GPS-equipped buses that travel along arterials throughout the region.

URBAN CONGESTION REPORT October 2004

CITY	PERCENT CONGESTED TRAVEL (New 12-mo. 0 LOs / New 12-mo. His 0; 0% 50% 100%)	TRAVEL TIME INDEX (New 12-mo. 0 LOs / New 12-mo. His 2; 1.45 1.70 1.95)	BUFFER INDEX (New 12-mo. 1 LOs / New 12-mo. His 3; 1.65 2.30 2.95)	Max Usable Days:21	CONTRIBUTING FACTORS: Monthly Precip. vs. Norm	Rain > 0.5" (NUMBER OF DAYS)	Rain > 1.0" (NUMBER OF DAYS)	Frozen Precip. (NUMBER OF DAYS)	Fog (NUMBER OF DAYS)	Work Zone Count	Monthly Incidents Hours vs. Norm
Atlanta	24 44	1.70 1.85	2.20 2.52	16	70%	0	0	0	0	N/A	48%
Boston	39 46	1.56 1.65	1.95 2.32	20	50%	1	0	0	2	N/A	78%
Chicago	52 84	1.64 1.77	2.06 2.45	17	105%	2	0	0	1	N/A	91%
Houston	26 41	1.60 1.71	2.11 2.29	18	52%	1	0	0	3	N/A	63%
Los Angeles	55 65	1.61 1.70	2.04 2.21	20	1232%	4	1	0	3	N/A	105%
Louisville	12 21	1.56 1.80	1.87 2.83	17	313%	1	1	0	0	N/A	168%
Miami	57 74	1.51 1.63	1.84 2.10	20	89%	3	1	0	0	N/A	58%
Philadelphia	24 41	1.51 1.63	1.93 2.14	19	81%	1	1	0	0	N/A	23%
Pittsburgh	14 34	1.51 1.68	1.84 2.16	18	149%	1	0	0	3	N/A	38%
San Antonio	9 27	1.52 1.64	1.87 2.44	21	245%	0	0	0	0	N/A	108%
Seattle	36 53	1.50 1.59	1.83 2.04	20	89%	0	0	0	6	N/A	80%
Portland	17 36 1.21	1.31	1.94 2.37	21							

LEGEND: 12-mo. LO; 12-mo. HI current month; New 12-mo. HI; New 12-mo. LO

FIGURE 3 FHWA Urban Congestion Report