MetroNet: Case Study for Collaborative Data Sharing on the World Wide Web

Robert F. Dickerson Jiakang Lu Jingyuan Li Billy Chantree Jian Lu
John A. Stankovic Kamin Whitehouse

Computer Science Department
University of Virginia

Abstract

We demonstrate MetroNet, which is an application that illustrates collaborative data sharing on the Wide Wide Web. MetroNet has two parts. First, our sensors gather data about pedestrian foot traffic in front of and into stores, which is made available online to shopkeepers. Second, this sensor data can be made public, at the discretion of the shopkeeper, for use by city planners, other shopkeepers, or residents of the city. MetroNet is an application through which we study fundamental problems of sharing data on the Web, such as search, data fusion, and privacy.

1 Data Sharing on the Web

The value of anywhere/anytime information about the physical world has long been clear, and as sensor networks are deployed across the globe by many different organizations such as shopkeepers, home owners, scientists, and governments, it is also becoming clear that *data sharing* will be an important part of achieving this goal. For example, 3000 motion and perimeter sensors are soon to be deployed in the New York City rails, bridges and tunnels for anti-terrorism purposes. If this data were shared, it's value would multiply: it could also be used by passengers to get real-time train information or to see when rush hour crowds have gone.

Mechanisms for sharing data are sometimes called the *World Wide Sensor Web*, and have been proposed in many forms. IrisNet stores sensor data in a distributed database which is dynamically partitioned for geographic indexing [4]. Other proposals use distributed hash tables [5], push data through an aggregation network [7], or move data to a centralized location and run continuous queries [2, 1]. The SensorMap project at Microsoft Research also uses a database-oriented approach but generates the queries from a graphical map interface [6].

We envision a different kind of Sensor Web – one that more closely resembles the World Wide Web – in which people publish sensor streams on the Internet through data hosting services and consume data with the help of search engines. The primary advantage of this approach is simplicity; there is no system, schema, or API to conform to. However, there are also a number of disadvantages: data will be difficult to find, and may be posted in non-standard data formats.

In this work, we explore mechanisms to allow data sharing over the World Wide Web. We proposed a web abstraction called a *StreamFeeds* [3] for publishing data on the Web, and are currently investigating mechanisms for search, data fusion, and privacy. To make the discussion concrete, we are developing a project we call *MetroNet*, in which sensors in the storefronts of Charlottesville, VA monitor pedestrian traffic on the downtown pedestrian zone. MetroNet has a key property that would be true of all networks on the Sensor Web: users are willing to build new sensor infrastructure for themselves, and have a vested interest in sharing data with others.

2 MetroNet Description

MetroNet is a system we are designing to measure foot traffic near store fronts in downtown Charlottesville, VA, where shops are densely packed and rely primarily on pedestrian traffic. This simple system will be used by shops to identify the effectiveness of window displays, signs, or other techniques designed to attract foot traffic into the shop. It will measure (i) how many people walk by the shop (ii) how many people enter the shop. Future versions of this project will also measure activity within to store to help the shopkeeper evaluate internal displays and their effect on *conversion rate*.

This type of information is extremely valuable and is already collected in many retail stores. However, there are also a number of compelling reasons to *share* the data. For example, shopkeepers could use data from neighboring shops to normalize their own data, allowing them to differ-

978-0-7695-3157-1/08 $25.00 © 2008 IEEE
DOI 10.1109/IPSN.2008.68

Figure 1. The Metronet Application allows shopkeepers to acquire new sensors, deploy their sensors, analyze data, and view a map of public sensor streams.

entiate the effect of a bad storefront display from the effect of bad weather, vehicular traffic, and downtown events such as fairs or concerts. Once in place, the sensors and their data would also be extremely valuable to the residents and planners in the city. For example, potential residents who would like to buy a new house could query for aggregate activity levels in each part of the city to find a location that is close to the city's night life. This type of query could be augmented in a subsequent phase by collecting supplemental data such as average levels of noise, sunlight, humidity, or air pollutants. Potential business owners would have empirical data upon which to base the value of commercial real estate, and city planners could gear zoning and traffic decisions to increase business and quality of life for residents. In fact, one motivation for the MetroNet testbed is to provide empirical data for an ongoing debate in Charlottesville about the effect on business of vehicular traffic crossing the downtown pedestrian zone.

MetroNet illustrates the principle of *multiplicative benefit*: a small effort on the part of a few can be multiplied into large benefits for many. Examples for this principle abound on the Internet, eg. YouTube and Wikipedia. Our goal is to use MetroNet as a platform to explore the fundamental computer science questions that would allow this principle to be applied to sensor data: search, data fusion, and privacy.

3 Demo Overview

Our demo consists of PIR sensors that monitor pedestrian traffic during the demo session. These sensors will wirelessly transmit the data to a base station, where it is displayed in MetroNet's Web interface shown in Figure 1. This interface is designed to be used by a shopkeeper who owns MetroNet sensors to enable viewing of the data. It also enables the creation, tagging, and general management of MetroNet data streams.

Once the data stream is created by the shopkeeper, simple data processing operations can be performed. These can be used, for example, to calculate the number of people walking in front of or into the store. From the interface, the user can also enable or disable sharing on the data stream, or on processed values of the stream. Shared streams are available through a URL on the World Wide Web, using the StreamFeeds abstraction.

4 Acknowledgments

This work is funded by Microsoft Research.

References

[1] S. Chandrasekaran, O. Cooper, A. Deshpande, M. J. Franklin, J. M. Hellerstein, W. Hong, S. Krishnamurthy, S. Madden, V. Raman, F. Reiss, and M. Shah. Telegraphcq: Continuous dataflow processing for an uncertain world, 2003.

[2] M. Cherniack, H. Balakrishnan, M. Balazinska, D. Carney, U. Cetintemel, Y. Xing, and S. Zdonik. Scalable Distributed Stream Processing. In *CIDR 2003 - First Biennial Conference on Innovative Data Systems Research*, Asilomar, CA, January 2003.

[3] R. Dickerson, J. Lu, J. Lu, and K. Whitehouse. Streamfeeds: An abstraction for the world wide sensor web. In *Internet of Things*, 2008.

[4] P. B. Gibbons, B. Karp, Y. Ke, S. Nath, and S. Seshan. Irisnet: An architecture for a world-wide sensor web. *IEEE Pervasive Computing*, 2(4):22–33, 2003.

[5] R. Huebsch, J. M. Hellerstein, N. L. Boon, T. Loo, S. Shenker, and I. Stoica. Querying the internet with pier, 2003.

[6] S. N. J. Liu and F. Zhao. Challenges in building a portal for sensors world-wide. First Workshop on World-Sensor-Web: Mobile Device Centric Sensory Networks and Applications (WSW'2006), 2006.

[7] J. Shneiderman, P. Pietzuch, J. Ledlie, M. Roussopolous, M. Seltzer, and M. Welsh. Hourglass: An infrastructure for connecting sensor networks and applications. Technical report, Harvard University, 2004.