SenseWeb: Browsing the Physical World in Real Time

Andre Santanche, Suman Nath, Jie Liu, Bodhi Priyantha, Feng Zhao
Microsoft Research
One Microsoft Way
Redmond, WA 98052, USA

1 Introduction

Geo-centric web interface such as MSN Virtual Earth (`http://local.live.com`) and Google Maps (`http://maps.google.com`) are useful to visualize spatially and geographically related data such as addresses, neighborhoods, weather, traffic, and so on. Desires to augment additional useful information to these interfaces have led people to create custom applications that overlay their own data on top of browsable maps. Examples of such applications overlay housing information (`http://www.housingmaps.com`) and crime-rate data (`http://www.chicagocrime.org/map/`) on Google maps, locations of vehicles (`http://jprestonsystems.2mydns.com/vemap.aspx?name=demoacct2`) and podcasters (`http://www.podlook.com/map.aspx`) on MSN Virtual Earth, weather data on custom maps (`http://www.wunderground.com`), etc. Such applications have been possible after Google Maps and MSN Virtual Earth have published useful APIs to overlay location data on their maps.

We envision publishing and querying real-time data (e.g., from sensors) over such geo-centric web interfaces. Existing solutions, although useful to write simple applications as above, have several drawbacks in achieving this vision. First, publishing even a single stream of data as a useful service is a nontrivial task. Many useful data is not being published yet because the data owners do not have enough programming expertise, or publishing it requires too much effort. Second, all the existing applications are mutually incompatible. One can't bring up a single map that shows both the housing information and crime-rates in an area. Third, existing solutions do not provide useful primitives such as querying live sensors based on keywords or location and aggregating the results in useful ways.

The SenseWeb project at Microsoft Research aims to address these challenges by providing a common platform and set of tools for data owners to *easily* publish their data and users to make useful queries over the live data sources. The SenseWeb platform transparently provides mechanisms to archive and index data, to process queries, to aggregate and present results on geo-centric web interfaces such as MSN Virtual Earth, etc. We believe that such a platform will encourage the community to publish more live data on web and users to build useful services on top of it.

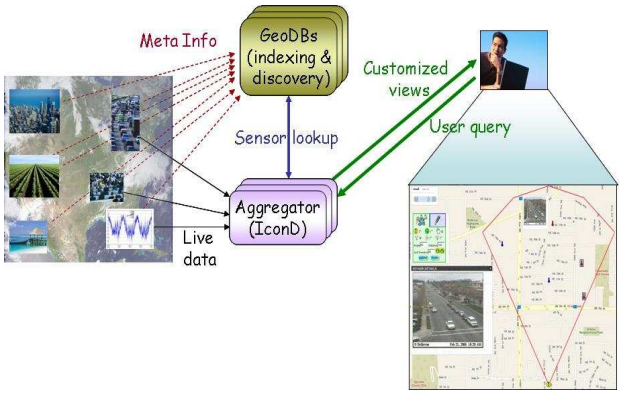

Figure 1: *SenseWeb Architecture.*

2 SenseWeb Architecture

The SenseWeb platform consists of the following four components (Figure 1): the *data publishing toolkit* publishes sensor data and metadata, a *GeoDB* indexes the data so that it can be queried efficiently, an *aggregator* aggregates clusters and summarizes data in useful ways, and the *client-side GUI* lets users query data sources and view results.

Data Publishing Toolkit. This toolkit runs on nodes connected to the sensors and the Internet. It provides sensor metadata to GeoDB and sensor data

in response to user queries. The metadata includes sensor location and type. SenseWeb uses a sensor ontology (an extension of the namespace defined by Open Geospatial Consortium [1]) to define inheritance, associates, and compositional relationship among different sensor types. This helps the client software to filter, fuse, and automatically visualize sensor data.

GeoDB. GeoDB is the portal for registering sensor metadata. We envision that typical user queries will be based on sensor types, descriptive keywords, and geographic locations, such as list of all cameras along a route or average temperature reported by all the thermometers inside a geographic region, etc. To efficiently support this type of queries, GeoDB indexes data by using hierarchical triangular mesh (HTM) indexing scheme [2] which is particularly suitable for geographic queries. The indexing is implemented as table valued functions in SQL server.

Aggregator. The aggregator mashes up sensor data with maps. It accepts queries from the client and redirect the geographic components of the queries to the GeoDB. After obtaining the metadata of a set of sensors that satisfy a client query, it contacts the sensors (data publishing toolkits) for their real-time data. It then aggregates the data accordingly (e.g., depending on the zoom level of the underlying map shown to the client). By doing so, SenseWeb provides useful summarization of data to the client. The particular aggregation performed by the aggregator depends on sensor types. For example, an average of the temperatures in a neighborhood can be displayed for data collected from thermometers.

Client-side GUI. It lets end-users to pose queries by drawing geometric shapes (e.g., a region, a route) and by specifying keywords and sensor types. It also overlays the results on MSN Virtual Earth.

Note that the data owner only needs to install the Data Publishing Toolkit and the end-users only need to browse a web page to query the data. The GeoDB and the Aggregator are transparent to both the data owners and users.

3 Demonstration Scenarios

We will demonstrate the following features of SenseWeb. First, we will show how the data owner registers his data source in SenseWeb. Second, we will show how an end user queries the data sources based on keywords, sensor types, and geographic regions or routes and how the data get displayed over MSN Virtual Earth. Finally, we will show how the live data from different sources get clustered and aggregated as the end user changes the zoom level of the map in MSN Virtual Earth.

References

[1] Reichardt, M. Sensor web enablement. Open Geospatial Consortium (OGC) White Paper 05-063, 2005.

[2] Szalay, A., Gray, J., Fekete, G., Kunszt, P., Kukol, P., and Thakar, A. Indexing the sphere with the hierarchical triangular mesh. Tech. Rep. MSR-TR-2005-123, Microsoft Research, September 2005.