Database Management Issues in Sensor-Enabled Distributed Environments

Sharad Mehrotra**, **Department of Information and Computer Science, University of California, Irvine, Irvine, CA 92612

Introduction

With the advances in embedded processors, low cost sensor technologies, and wireless communication (e.g., cellular and bluetooth technologies), smart devices that generate unprecedented amounts of diverse types of information about the real world and its activities will become an integral part of the ubiquitous computing fabric. Such devices will be capable of real-time capture, storage, processing, and transmission of information about objects dispersed in space and time that interact and communicate with each other and their surroundings. An integral component of a sensor enriched communication and information infrastructure is the database management technology that allows seamless access to information dispersed across a hierarchy of storage, communication and processing units -- from sensor devices, where data originates, to large data banks where the information generated is stored for analysis and mining. Such a database infrastructure must provide effective means for storage, representation, search and processing of massive volumes of highly dynamic spatio-temporal events. The database infrastructure will facilitate applications such as real-time monitoring, tracking and analysis of objects/events/phenomena dispersed in space and time. In this paper, we discuss some of the key challenges we are exploring in developing adaptive and scalable database infrastructure for sensor-enriched distributed environments as part of the Quality Aware Sensor Architecture (**QUASAR** for short) project [1-6] at UC, Irvine. Before we discuss the challenges, we first elaborate on a specific application scenario that motivates our research.

A Motivating Application[1]

While computing environments with embedded smart devices have the potential to revolutionize a wide variety of application domains including avionics, emergency response and disaster relief operations, science we focus on the domain of emerging high-speed motorways and vehicles that drive upon them. In the not too distant a future, we expect the following core technologies to become ubiquitous with the potential to revolutionize the driving experience on modern highways:

- **Location awareness**: with declining costs and economy of scale, we can expect that most vehicles will be equipped with GPS receivers providing fairly accurate geographical position coordinates. (Note: many luxury vehicles are already being outfitted with GPS receivers.).
- **Inter-Vehicle Communication:** many high-end vehicles are already being equipped with sophisticated computing components interconnected via a LAN that control many mechanical components. We believe that this trend will very likely continue and devices will be installed in vehicles that will allow for short-range inter-vehicle communication while forming ephemeral rapidly evolving ad hoc networks.
- **Vehicle-Infrastructure Communication:** Roadside infrastructure comprised of possibly a number of base stations strategically positioned in close proximity to the highways will be installed enabling travelling vehicles to communicate with the fixed infrastructure. Such an infrastructure will enable information to flow both ways – from vehicles to the infrastructure and vice versa. The infrastructure can be used for monitoring current traffic conditions, as well as managing traffic. (Vehicles elsewhere can also query the fixed infrastructure for trip planning purposes.)

The above listed technologies will provide enhanced situational awareness to drivers that will besides facilitating their decision making tasks (e.g., trip planning based on traffic congestion on the road), will improve highway safety (by bringing information about catastrophic events and road conditions to the driver’s attention). For example, data captured by the traffic network, when properly aggregated, can be fed into the traffic monitoring and flow control system for real-time traffic management. Alternatively, such information can be archived for off-line analysis to understand traffic bottlenecks and devise techniques to alleviate traffic congestion. While there are numerous possibilities, some of which may spark debates over privacy rights and apprehensions over undue monitoring, we describe below two applications that we believe are non-intrusive and of an assistive nature (hence less controversial).

---

[1] The principal ideas of this application domain listed below have emerged from joint work [6] with other colleagues whose research explores networking and communication, security, and middleware issues related to sensor enriched distributed environments.

One of the most compelling application examples is what we call "**V**ehicle **I**mmediate **V**icinity **A**wareness" or **VIVA** for short. **VIVA** aims to communicate to each vehicle *vital signs* of other vehicles that are traveling in close proximity. Proximity in this context means the area that falls within direct range of transmission of the wireless networking device found in each vehicle. Such vital signs may include the status of: turn signal indicators, brake application, relative/absolute speed, headlights, etc. It is important to note that all of these signs are, in any case, intended for external display, e.g., the status of the turn signal is always in plain view. In other words, no new information (that might be construed as private) is intended for communication outside a vehicle. We anticipate that the information collected via **VIVA** will assist the drivers by offering them better awareness of their immediate surroundings as well as of the current and intended behavior of the nearby vehicles. Drivers will be able to better concentrate on the road ahead if they no longer have to look sideways to observe flanking and tailgating vehicles. Also, notorious blind spots can be effectively eliminated if drivers are continuously made aware of the surrounding space. Furthermore, traffic conditions ahead can be observed faster if drivers are warned of braking activity one or two vehicles ahead. Consequently, we expect **VIVA** to increase highway safety by preventing some accidents.

Another, very different, application has to do with trip planning. We refer to it as "**Hi**ghway **T**raffic **C**ondition **H**elper" or **HITCH** for short. In it, vehicles communicate directly to the fixed infrastructure (base stations) and report their vital signs (as above) as well as the speeds of surrounding vehicles. This information is efficiently gathered and coalesced into regional traffic snapshot databases. Vehicles both on and off the highway can query these databases and obtain immediate information on traffic conditions towards intended destinations. (Of course, wired clients, i.e., Internet users at home can also use **HITCH**.)

Data Management Challenges

Data management infrastructure for sensor-enriched distributed environments poses a large number of challenges. Let us first look closely at the nature of data generated, the nature of applications, and the resource constraints imposed on the solutions by the underlying communication and computation infrastructure.

Data generated in sensor environments is very dynamic and spatio-temporal in nature—e.g., a vehicles location may change continuously. Given the impossibility to constantly monitor a spatio-temporal event as well as the inherent limitations of device precision, data representation will have inherent imprecision and/or uncertainty. Mechanisms to represent data imprecision and furthermore account for it in applications will need to be devised. Sensors allow us to capture and create very large data repositories. For example, vehicular information captured over a period of an year may lead to several 100s of terabytes of information that can be utilized for a variety of purposes. Current data may be used for real-time traffic control (as in **VIVA** and **HITCH** applications discussed above). Suitably aggregated over a period of time, it could also be used for traffic analysis and demographic studies. Resource limitations that must be considered include, at least in the short term, the limited computation, storage, and communication capabilities of sensors, power supply at the sensor, and limited communication bandwidth among sensors and between sensors and the fixed infrastructure.

The first and foremost challenge is that of database architecture that establishes where data generated at sensors is stored and where queries execute. It should be obvious that sensor-enriched distributed environments, such as the vehicular infrastructure discussed in this paper, can be viewed as a highly distributed and dynamic database in which sensors/data producers (e.g., vehicles) may be considered as smaller (light-weight) databases. Data may reside at the data producers and/or may, alternatively, migrate to the backbone storage and computation units. Furthermore, information may be cached, at different levels of aggregation depending upon the application, from the infrastructure and/or other data sources to the vehicles to expedite query processing. The data management architecture will need to accommodate various resource constraints and itself depends upon the strategy used to represent dynamic data as well as the nature of queries.

Challenge in representing dynamic data (e.g., vehicle location) at the fixed infrastructure (and/or other vehicles) arises due to the impossibility in reflecting each change via an explicit update given the database, network, and power consumption costs of updates. Instead, a dynamic attribute may be represented in the database using a prediction model and a set of parameters to that model that can be used to estimate the value at any given time. The prediction model may also include an error model that establishes (possibly tight) bound on the error in prediction. Based on this representation, the system will attempt to answer the query locally if the answers can be produced at the desired level of error tolerance. Else, the query may execute at the infrastructure that may in turn have to probe other data producers if the quality tolerance of the query cannot be met. Many research issues arise. Given the prediction models used for dynamic data at different levels of the system, the resource constraints, as well as some information about the nature of queries and their error tolerances, information collection/caching mechanisms that control the information flow through various components of the system need to be devised. These mechanisms should attempt to minimize the resource consumption. Collecting data more frequently than required results in resources wastage while collecting data at a slower rate will result in probes during query execution causing increased latency of applications. A related issue is that of model selection for dynamic attributes. Different models might be suitable at different times for a given data source and techniques to choose the optimal representation will need to be devised. For example, a vehicle may initially be stationary, may then accelerate to a constant speed, and then maintain its speed. Different prediction models may capture the vehicle behavior at different times optimally.

Another fundamental issue is that of quality-based query processing. Given the underlying imprecision in representation and tolerance in queries, techniques to correlate the quality of results based on quality/age of data will need to be developed. Innovations are needed at the language level (e.g., SQL level) to enable queries to specify their error tolerances and quality requirements. Query optimization and processing techniques that trade quality for resource consumption need to developed. Similar issues of tradeoff between quality and resources have previously been explored in the area of multimedia systems where resource provisioning approaches that meet the QoS requirements (frame rate, jitter, frame quality) of applications have been extensively studied. However, in our case, we need to define a notion of quality for information. While quality measures for values (e.g., aggregates) such as absolute error are easy to define and exploit during query processing, quality measures for set valued attributes such as earth movers distance or match and compare are more complex. Instead, in [2] we explore a different mechanism that generalizes the well known precision and recall measures for this purpose.

Above we have discussed only two of the challenges in sensor enriched distributed environments – adaptive information collection and quality-based query processing. Other challenges include indexing, and query optimization in presence of dynamic data. Furthermore, monitoring and situational awareness applications in sensor environments result in numerous interesting queries that have erstwhile been not studied in database literature. A taxonomy of such queries along with some solutions is described in [5]. Possibly the most interesting is the notion of continuous query studied in [3]. Traditionally, in databases queries are issued explicitly, the database evaluates it and returns the results one time. In contrast, in a dynamic query a user (e.g., a driver) who himself is mobile may register a query with the system which is continuously evaluated. An example of such a query is to continuously monitor traffic en-route. Implemented naively, the vehicle will generate a fresh query every time its location changes or an event occurs that changes the result set of the previous query. Given the continuity of motion in space and time, it is possible to optimize such continuous queries extensively. The focus of the discussion has been on data management challenges to support real-time applications such as situation monitoring. Many research challenges also arise in collecting and storing data generated from the devices for offline analysis and mining. Given the gargantuan amounts of data mechanisms to appropriate compress the data (without losing much content), as well as approaches to support multiresolution representation of spatio-temporal data sets that allow for efficient approximate and progressive query evaluation need to be explored [4].

Below we list a set of references from which much of the content of this paper has been extracted. We note that besides our effort, many other researchers are working on related problems. The references to their work can be found in the citations included below.

Selected Quasar References

[1] Iosif Lazaridis, S. Mehrotra, K. Porkaew, R. Winkler, "Database Support for Situational Awareness", *Computer Science Handbook, M.S. Vassilou and T.S. Huang (eds.), Army Research Laboratory, 2001*

[2] I. Lazaridis and S. Mehrotra, “Quality Aware Query Processing”, *Submitted for publication. 2002*.

[3] Iosif Lazaridis, K. Porkaew, S. Mehrotra, "Dynamic Queries over Mobile Objects", *To appear in to the 8th International Conference on Extending Database Technology, 2002*.

[4] Iosif Lazaridis, S. Mehrotra, "Progressive Approximate Aggregate Queries with a Multiresolution Tree Structure", *2001 SIGMOD Conference on Management of Data, 2001*.

[5] K. Porkaew, I.Lazaridis, S. Mehrotra, "Querying Mobile Objects in Spatio-Temporal Databases", *7th International Symposium on Spatial and Temporal Databases, 2001*.

[6] M. El Zarki, S. Mehrotra, G. Tsudik and N. Venkatasubramanian, “Security Issues in a Future Vehicular Network,” to appear in EuroWireless 2002, Florence, February 2002, to appear.