Database Architecture for Autonomous Transportation Agents for On-scene Networked Incident Management (ATON)

Mohan Trivedi, Shailendra Bhonsle, and Amarnath Gupta* Department of Electrical and Computer Engineering, *San Diego Supercomputing Center, University of California, San Diego, USA. Email: trivedi@ece.ucsd.edu

Abstract

A collection of distributed databases forms an important architectural component of the ATON project for networked incidence management of highway traffic. The database sub-architecture supports the architectural integration of many thematic areas of the ATON, and provides many high level abstractions that semantically correspond to traffic incidents. These databases are queried for the detection of local or distributed traffic incidents by many distributed control and analysis algorithms. The abstract database sub-architecture leads to many concrete databases of different types having different functions. The semantic event/activity database is one of them. The distributed multi-sensory subarchitecture detects atomic semantic events that occur in the environment. The event/activity database stores them and their temporal structures. An activity is a temporal composition of atomic events. A specialized query language is used to flexibly define and detect activities of interest. The query language embodies high level semantic pattern matching abstractions. Preliminary results of using the event/activity database are also presented in this paper.

1. Introduction

The eventual goal of the autonomous transportation agents for on-scene networked incident management (ATON) system is to automatically detect, verify, and manage traffic incidents. A traffic incident is defined as "an event that causes blockage of traffic lanes or any kind of restriction of the free movement of traffic" [5]. To achieve this goal, ATON uses clusters of video and acoustic sensors, mobile robotic agents and interactive multimedia workstations and interfaces, all connected using high-speed, high-bandwidth communication links. The project is divided into five thematic areas of robotics agents with coordinated behavior, distributed multisensory networks, ubiquitous tele-existence, distributed virtual environments, and distributed databases. Figure 1 shows a scenario for traffic incident detection, verification, and management. Many steps taken for incident management in this scenario are the automatic detection and reporting of incidents by sensor clusters, dispatch of a robotic mothership [6] or remote agent to the scene of the incident, dismounting of a team of little robotic agents [6] from the mothership to form a safe zone in a coordinated way, and assistance offered by the mother ship through tele-existence terminals.

This paper addresses the issues in the design of an integrated software architecture for coordination of many distributed control and analysis algorithms present in various thematic areas. The novel concept of using a set of distributed databases, having many different functions and types, is proposed for distributed coordination of these algorithms. This coordination paradigm using databases makes the whole architecture robust by providing means to efficiently manage current and past *states* of the *system* and the *monitored environment*. Such *states* are typically large in multi-sensory environments where some of the sensors could be as complex as omnidirectional video sensors.



Figure 1. Scenario depicting steps associated with traffic incident management

Specialized databases are also used in ATON to provide high level semantic abstractions for flexible definition and detection of traffic incidents. These abstractions are those of semantic events [3, 4], activities [1, 2], and behaviors [2]. **Events** are atomic semantic units that are detected over a bounded temporal extent by fusion of multi-sensory information. **Activities**, which have much higher semantic content than events, are spatio-temporal compositions of events. A **behavior** is defined to be a set of activities. A traffic **incident** has an even higher level of semantic content, and is modeled using the abstractions of *events*, *activities*, and *behaviors*.

Apart from providing a direct semantic link to the notion of a traffic incident, high level semantic abstractions are also required by many distributed control algorithms to control robotic actuators in the real world. Computations of these abstractions from raw sensory data have prohibitive computational costs, and when the amount of sensory information produced is very large, it is best to delegate the role of these computations to a database.

In section 2, databases providing different functions are briefly described. The event/activity database is discussed in section 3, and preliminary results of using a semantic event/activity database are described in section 4.

2. Databases in the ATON architecture

Many functional classes of databases are required for the distributed decision making to detect and manage traffic incidents. In this subsection, we only concentrate on *world databases* or databases that deal with the *states* of the monitored environments. Figure 2 shows a few of these databases and functions that they provide. The distributed nature of these databases is not shown in this figure.



Figure 2. Many types of databases used in incidence detection and management

At the core representation level, the information gathered by sensors in a multi-sensory environment can be treated as a huge multi-sensory database. A part of this database may consist of the video of the segment of highway being monitored. Another part of this database may consist of just the average intensity of light at the monitored site sensed by an appropriate sensor. Using appropriate sensor fusion algorithms, many functional classes of information, stored in relevant database, are derived from this raw multi-sensory database.

Spatial databases provide information about the spatial structure of the monitored environment like its

division into regions corresponding to traffic lanes, curb, exit lane etc. The **environment** databases provide static or dynamic environmental conditions like fogginess, rain intensity, luminance etc. Spatial and environmental databases depend only on the environmental conditions at the monitored site, and these are independent of the actual traffic related facts. **Traffic parameter databases** provide a number of traffic related parameters like counters that indicate the rate of flow of traffic, parameters to indicate the *burstiness* of the traffic etc. These databases make use of spatial and environment databases for computations of these parameters.

There are two classes of databases that deal with high level, domain-dependent semantics of observations: the feature and the event/activity databases. The feature databases deal with features of either mobile objects or spatial regions of interest. For example, these features could relate to the color, size, centroid etc. of certain monitored mobile objects or to statistical properties of the backgroud region. Features are the basic domaindependent semantic information associated with monitored mobile objects or spatial regions. Events usually represent the state-transition of observed mobile objects or spatial regions. They represent a form of atomic semantic units, and the monitored environment can be described as a collection of these events. In the next section, we describe a specific event database that stores events and their temporal structures.

A set of databases was listed above that provide basic information about the monitored objects and spatial regions. These databases are queried, and retrieved data is processed to provide information required by the ATON distributed decision making algorithms for incidence detection and management. For example, traffic flow analysis algorithms use these databases to produce flow networks. Many time-stamped versions of different types of such flow networks need to be kept in a separate database. Such specialized databases are our next level of databases providing a very high level semantic view of the physically distributed monitored environment.

The *type* of a database refers to its underlying data model. The range of database types that provide the required functions vary from the simplest and widely used *relational model* to data models that represent *complex spatial and temporal structures*. In the following section, we describe one such complex database used in the ATON architecture.

3. Semantic event/activity database

Semantic activities are complex spatio-temporal compositions of semantic events. The event/activity database stores events and their spatio-temporal interrelations. A transducer is used to detect events. The transducer uses raw sensor data, spatial information from a spatial database, and feature data of detected objects or regions of interest. The event/activity database query language provides high level semantic abstractions for specification of activities through its set of operations. These queries are executed against a database instance to detect activities. Also, there is provision for *active queries*. These use query abstractions similar to those provided by the activity query language, but they execute in real-time and can notify different entities of the ATON architecture of occurrences of activities. Many components of such event/activity databases together with their interactions with the sensor processing elements and the event detection transducer are depicted in figure 3.



Figure 3. Detection and storage of semantic events and retrieval of semantic activities

Design of such semantic event/activity databases depends on many factors. Some of these factors are whether events represent states or state-transitions of objects, whether spatial, temporal, or spatio-temporal aspect is emphasized, how spatial and temporal *uncertainties* in event occurrences are represented, and whether the model of composition of activities from events is statistical or combinatorial. For our modeling purposes, we defined events to represent state-transitions of observed objects. These objects may be the mobile objects like moving cars in a highway segment, or spatial regions like a specified lane of the monitored highway.

We focused on the temporal composition of activities from events, and an important consideration was to take into account the *temporal uncertainties* in event occurrences. Assuming that such uncertainties for a given set of events are bounded by a constant, a model of composition of semantic activities from events is obtained that suitably represents the *concurrent occurrences of events* in the presence of such temporal uncertainties. The concurrency may be between events associated with *multiple objects*, or between many events associated with a *single agent*. Such concurrency in the presence of temporal uncertainties is *not modeled* by ubiquitous sequential temporal composition rules, or by its simple extensions [2].

The combinatorial structure obtained above corresponds to semiorders (a proper subclass of partial orders) under certain binary relation. The binary relation is < precede, **D**> such that event x < precede, **D**> event y if and only if the occurrence time of y is greater than the occurrence time of x plus D, where D is the fixed temporal uncertainty interval. Semiorders based temporal compositions represent a natural evolution of sequential composition rules. We designed a semiorder based data model and a corresponding query language that also embeds a semiorder pattern definition language. Many algorithmic and architectural issues associated with the design and implementation of this database is discussed in [2]. In the following section, we briefly describe a set of simple queries to illustrate our use of this database in the ATON architecture.

4. Preliminary results

A prototype of the semantic event/activity database was designed, developed, and used for detection of complex traffic related activities. A set of traffic related atomic events like startLeftTurn, endLeftTurn, enterRegion, exitRegion etc. were defined and the corresponding modules were inserted in the event detection transducer. Many atomic events are defined only with respect to the specific spatial structure of the monitored highway segment as described below. The spatial structure, depicted in figure 4, is stored in a surrogate spatial database. The highway segment is divided into many layers, where each layer consists of non-overlapping regions like ExitLane, Lane1, Lane2, Curb, etc. The event detection system queries this spatial database to appropriately detect atomic events like LaneChange by a car. The event/activity database stores the detected events and their spatial and temporal attributes.



Figure 4. Spatial structure of the highway segment being monitored

The dynamically changing monitored environment was represented in the event/activity database as a collection of occurred events, their temporal structure, and their parameters. Complex traffic activities were detected using queries defined using the activity query language. Many complex temporal structures over a set of events were flexibly defined using *semiorder* patterns, which, in turn, were used by these queries. The execution of the query extracted the set of matching activities from the semiorder representation of the environment in the database. Although the prototype supports complex queries, for brevity, we only discuss a few simple representative queries. The Δ uncertainty bound was set to 100 ms in our experiments.

Query 1: (*Exit*) Find cars that entered *lane 1*, then entered the *exit lane*, and finally exited from the *exit lane*.

Query 2: (*Tailgate*) Find a set of cars that entered *lane 1* almost simultaneously.

Query 3: (*Simultaneous*) Find an object that *while* entering *lane 3* was also turning left.

The first query illustrates the sequential composition of an activity from atomic semantic events. An important fact is that *exponentially* many such *activities* can be defined on the same base set of *events*. Databases provide a tool to the user for *flexible* definition and detection of what is needed. This compares well to the currently used alternative of fixing a set of high level activities and then detecting them in an ad-hoc but efficient manner [7]. One of the results produced by this query is shown in figure 5.



Figure 5. Result of Exit query

The second query shows a semantic activity that is *multi-agent* and involves the specification of *concurrency*. The third query involves *multiple concurrent events* associated with a *single agent*. Some of the results produced by these queries are depicted in figure 6. Being able to handle the requirement of modeling concurrency of events in the presence of temporal uncertainties is an important strength of the semiorder database model.

These simple queries demonstrate the usefulness of the event/activity database model in the ATON architecture. Combinations of the essential elements of these queries specify semantic activities with very complex structures that the query language is able to deal with.

5. Conclusions

The role of databases in the integrated ATON architecture was described. It was pointed out that for detection and management of traffic incidents a number of databases having different *functions* and *types* are required. One of these many types of databases is the semantic event/activity database. These databases provide high level of abstraction for *flexible* definition and *detection* of semantic activities. These high level abstractions in turn help in the flexible definition and detection of *incidents* or *conditions* that lead to traffic incidents.

On the other hand, use of databases in such systems is a new paradigm for the integration of many distributed control and analysis algorithms that are present in many thematic areas. This is especially true in multi-sensory monitoring environments where the *global and local states* have large spatio-temporal extents and these states need to be made persistent for robust interactions amongst many distributed algorithms.



a. Tailgating



b. Simultaneous Figure 6. Results of (a)*Tailgate* and (b) *Simultaneous* queries

6. References

[1] S. K. Bhonsle, et.al., "Complex visual activity recognition using a temporally ordered database", *3rd Int. Conf. Visual Information Management*, Amsterdam, June 1999.

[2] S. K. Bhonsle, "Semiorder Model for Temporal Composition of Activities from Events in Multi-Sensory Environments", *Ph.D. dissertation*, Dept. of Computer Science and Engineering, Univ. of Cal. San Diego, March 2000.

[3] W. I. Grosky, "Managing multimedia information in database systems", *Comm. ACM*, Vol. 40, No. 12, Dec. 1997.

[4] A. Gupta, S. Santini, and R. Jain, "In search of information in visual media", *Comm. ACM*, Vol. 40, No. 12, Dec. 1997.

[5] K. Ozbay and P. Kachroo, "Incident Management in Intelligent Transportation Systems", Artech House, 1999.

[6] M. M. Trivedi, K. C. Ng, N. Lassiter, and R. Capella, "New generation of multirobot systems", *IEEE Int. Conf. on Systems, Man, and Cybernatics*, San Diego, USA, Oct. 1998.

[7] R. Weil, J. Wootton and A Garcia-Ortiz, "Traffic Incident Detection: Sensors and Algorithms", *Math. Comput. Modelling* Vol 27, No. 9-11, pp. 257-291, 1998.