Spatio-temporal Databases in Urban Transportation*

Ouri Wolfson and Bo Xu
University of Illinois, Chicago, IL 60607
{wolfson,boxu}@cs.uic.edu

Abstract

In this paper we describe applications, research issues, and approaches related to Intelligent Transportation Systems (ITS). More specifically, we focus on spatio-temporal databases in urban transportation. We address the issues of trip planning and navigation, abstraction of concepts from spatio-temporal sensor data, mobile peer-to-peer data management, and social networks and crowd-sourcing. These issues are strongly related to ITS efforts currently undertaken throughout the world, particularly the IntelliDrive initiative of the US Department of Transportation.

1 Introduction

The impact of Computer Science (CS) and Information Technology (IT) on transportation systems is not as dramatic as the one on finance or business in general. But in the last few years we have witnessed significant penetration of IT in surface transportation. Navigation systems with real-time traffic information, color coded traffic maps, and real-time information about public transportation vehicles (e.g. Nextbus and Chicago Transit Authority's bus-tracker) are some examples of the improvements in urban transportation brought about by IT. Rapid advances in mobile and ubiquitous computing, sensor networks, and Geographic Information Systems are creating opportunities to revolutionize urban transportation. Indeed, the purpose of the IntelliDrive initiative of the U.S. Department of Transportation is "advancing connectivity among vehicles and roadway infrastructure in order to significantly improve the safety and mobility of the U.S. transportation system" [18]. Additional compatible goals are reduction of environmental impact and energy consumption.

In the envisioned environment, billions of sensors embedded in the infrastructure, in portable devices, and in vehicles will generate vast amounts of data whose interpretation could be exploited to spur the creation of innovative transportation services and policies. Advances in social networking and data mining research are increasingly creating new sophisticated mechanisms which can foster seamless information integration among travelers, provide alternatives, and support sustainable economic and social policies.

In this paper we discuss data management issues in the envisioned environment. Specifically, we discuss the applications, research issues, and approaches in the following research areas.

- **Route Planning (Routing), Navigation, and Tracking.** This area studies methods of routing, navigation, and tracking in transportation networks (the spatio-temporal database) that may involve multiple modes

Bulletin of the IEEE Computer Society Technical Committee on Data Engineering

*Research supported by NSF DGE-0549489, IIS-0957394, and IIS-0847680.

such as train, bus, private car, and bicycle. Optimization criteria are traditional such as time, distance, and cost, but may also involve environmental and energy aspects. For example, the requested route should minimize exposure to, and/or generation of, pollution. Data models, query languages, and processing algorithms are key research issues. The maintenance of dynamic, up-to-date, and reliable travel time and incident information is another important research issue.

- **Abstraction of concepts from spatio-temporal sensor data.** This area studies mining techniques for analyzing vast amounts of data related to moving objects. For example, map matching (i.e., matching the GPS traces of a traveler to the road network) extracts route information. A higher level abstraction would be extraction of semantic location and activity knowledge from GPS traces, possibly augmented with other information such as visited-web-sites lists. This analysis may significantly improve the accuracy and efficiency of household activity surveying, which is an important data source for transportation planning.
- **Mobile peer-to-peer data management.** This area studies the problem of querying and dissemination of spatio-temporal traveler information via short-range wireless communication among vehicles, pedestrians, and road side facilities. The spatio-temporal traveler information includes, for example, accidents, transit vehicle on-time performance, parking availability, and ride share opportunities. The key research issue is query processing in a distributed and mobile ad hoc environment, without data-directory services. Bandwidth, energy, and memory constraints complicate the problem.
- **Social networks and crowd-sourcing.** This area studies techniques based on social networks to crowd-source information that is valuable to travelers, traffic administrators, and long-range transportation planners. For example, traffic congestion, parking slot availability, ride share opportunities, and cooperative driving [29] information will be obtained and disseminated via (possibly ad hoc) social networks. A major research issue is providing incentives for participation.

In sections 2 to 5, we discuss each of the above research areas separately.

2 Route Planning (Routing), Navigation, and Tracking

2.1 Applications

For trip planning, travelers often ask queries that pertain to routes involving various constraints, optimization criteria together with uncertainty clauses (see e.g., [21]). For example, the queries can be: ”Find a route that will get me home by my designated time with at least 90% certainty”; or ”Using public transportation, find a route that lets me stop at a grocery store for 30 minutes and reach home by 7:00pm”. A special application that has been studied extensively is evacuation route planning: given a transportation network, locations of a vulnerable population, and a set of destinations, identify exit routes to minimize the time to evacuate the vulnerable population [12]. While the planned trip is being executed, it is often useful to track the traveler, an activity that consumes resources such as bandwidth, energy, and processing power.

2.2 Research Issues and Approaches

Relevant research issues include: 1. data models and appropriate query languages to represent the routing problem; 2. performance of the routing and tracking algorithms (see e.g. [25]), and their incorporation into the query processing engine; 3. efficient and effective collection and distribution of travel-time information; 4. travelers are usually interested in travel time of a road link when they get to it, so it is important to *predict* travel-time based on past travel times, and based on travel-times on near-by road links (see e.g., [13]); 5. scalability and performance of transportation network micro-simulators such as CORSIM, VISSIM, and TRANSSIM with increase of network size, traffic intensity, and fidelity of vehicular behavior [16]. In the rest of this subsection we will elaborate on existing approaches to some of the above problems.

2.2.1 Transportation network data models, query languages, and processing

One technical difficulty with multi-modal public-transportation trip planning is the heterogeneous nature of the data and the lack of a coherent data model that can be used effectively. In [4] we proposed a method to integrate the key aspects of spatio-temporal, moving objects, and graph-based databases to facilitate trip planning in multi-modal urban transportation networks. Our solution is based on a graph model of the network. This graph model contains not only topology of the network but also various other information such as real-time traffic information including predicted travel times, dynamic schedule information, available facilities, etc. A trip corresponds to a path from an origin to a destination in this graph.

We defined the language for specifying queries in our graph model. The language focuses on querying trips subject to various constraints – including uncertainty of travel times on the links. The query structure extends the basic syntax of SQL with additional clauses. For example, the following query retrieves trips/paths from work to home where only the pedestrian and bus modes are allowed. The CERTAINTY and WHERE clauses specify that the trip must end by 5:00 pm with probability greater than 0.8. Among all such paths the shortest path (by length) is selected.

*SELECT * FROM ALL_TRIPS(work, home) AS t WITH MODES pedestrian, bus WITH CERTAINTY > .8 WHERE FINISHES(t) <= 5:00pm MINIMIZE LENGTH(t)*

The reader is referred to [4] for the detailed syntax and semantics of the query language. That work also presents algorithms for processing a restricted class of queries specified in this language, and proposes a plug-and-play approach in which existing and new algorithms developed in the Operations Research field can be incorporated. For a prototype implementation see [3].

2.2.2 Maintenance of link travel-times

Accurate information about travel-times on road-links of the transportation network is critical for trip planning. **Travel-times Collection.** Point detection by inductive loop detectors, ultrasonic detectors, remote traffic microwave sensors, video cameras, etc. is a common travel-time collection method. These devices are typically deployed on selected highways, but not on urban streets for both economical and technical causes. A promising alternative to point detection is floating car data (FCD), or probe car data, which uses probe vehicles to collect travel-time information. FCD can cover both highways and arterial roads without dedicated sensor deployment. A similar method is to use mobile phones to collect travel-time information as demonstrated by the Mobile Millennium project (traffic.berkeley.edu). For this method a technical challenge is to identify and eliminate the data generated by pedestrians. In the rest of this subsection we focus on FCD.

In terms of system architecture, there are two ways to collect floating car data. One way uses a client-server architecture, where vehicles (clients) send the travel-time that they experience to the server. This informs the server of the real time traffic condition of each road segment. In turn, the server broadcasts updated travel-times to the vehicles. In [2] we proposed methods that reduce communication by using randomization. The other way is based on a mobile peer-to-peer (P2P) architecture, where vehicles send travel-times of road segments to each other via vehicle-to-vehicle communication, without server facilities (see TrafficView [8] and SOTIS [24]).

Travel-times Aggregation. Collected travel-times may be aggregated in different ways for different purposes. In mobile P2P FCD, travel-times may be aggregated to reduce the communication volume [8]. In this case, the level of aggregation varies such that the travel-time information is less accurate about regions that are farther away from the vehicle at which the aggregation occurs. In client-server FCD, travel-times may be aggregated to generate a travel-time profile of a link, namely how the travel-time of the link varies depending on the time of a day, the day of a week, etc. Trip planning based on time-dependent travel-times enables more efficient navigation. The authors in [17] proposed a data warehouse for building travel-time profiles. The basic fact table contains the collected travel-time data. Travel-times are aggregated along the spatial and temporal dimensions to create a data warehouse enabling efficient processing of queries such as ”What is the average travel-time of the links in downtown Chicago on Mondays from 9:00-9:15?”.

3 Abstraction of Concepts from Spatio-temporal Sensor Data

3.1 Applications

The transportation environment consists of various mobile sensors such as on-board GPS receivers, sensors mounted on public transportation vehicles (to monitor, e.g., traffic, air quality, pollution, and toxic gases) or pedestrian cell phones. These sensors continuously generate spatio-temporal data and enable applications such as moving objects tracking and environmental monitoring [10]. The abstraction of concepts from spatio-temporal data can derive, for example, the context of a traveler which indicates where she is and what she is doing. The context information may be used to reduce user burden when dealing with the restricted user interfaces of mobile devices. For example, if the traveler's smartphone finds that its owner is at an electronic store, then the smartphone may automatically retrieve appropriate coupons. The context information may also be used to complement the traditional paper based household travel survey by automated activity detection. Another application of concepts abstraction is to identify energy efficient driving patterns, and optimal routing patterns for minimization of pollution and energy consumption.

3.2 Research Issues and Approaches

Relevant research issues include: 1. data mining to extract semantic location and activity knowledge from sensor traces; 2. efficient monitoring of spatio-temporal streams in terms of communication, storage, and query processing; 3. management of uncertainty which is inherent to continuously changing variables such as traveler locations; 4. fusion of information from multiple sensors generating heterogeneous data. In the rest of this subsection we will elaborate on existing approaches to some of the above problems.

3.2.1 Map matching

For the purpose of abstraction, usually GPS reported locations need to be mapped to the road network. This procedure is called *map matching*. A straightforward way to do map matching is to snap each GPS point to the closest road segment. However, this method often produces incorrect results [28, 11]. A better method is to snap subsequences of GPS points in 3D space-time. For example, in [28] we proposed an algorithm based on a weighted graph representation of the road network in which the weight of each road link represents the distance of the link to the trajectory. The matched route in the road network is then found by computing the shortest-path in the weighted graph.

3.2.2 Extracting semantic locations from GPS traces

A semantic location carries context information about the physical (x,y) location, e.g., ”office at 851 S. Morgan St”, or ”grocery store at 1340 S. Canal St”. In [14] we proposed a method to extract the semantic locations from traces of GPS points augmented with behavioral information. The principle is as follows. First, we extract some 2D positions where the user stays for a while. Then, we use reverse geocoding (i.e., translating a physical location to a street address) to obtain the street address candidates for each stay. Finally, correlating with other available sources (such as semantic locations history, map, yellow pages, address book, calendar, phone-calls trace, visited web pages trace), we determine the semantic location of a stay.

3.2.3 Automated survey data reduction

In traditional travel surveys, respondents have difficulties giving accurate data due to cumbersome data entry requirements and recall limitations. An effective approach to reducing respondents' burden and increasing data accuracy is to automatically detect activities from GPS traces, and prompt them for user verification. This approach is called *prompted recall surveying* [1]. Activities may be detected using sequential association mining techniques (see e.g., [23]). Specifically, a respondent's activity schedule over a time period (e.g., each day) is modeled as a sequence. Association rules are learned from historical sequences, e.g., "with 80% probability the traveler goes to a restaurant after shopping on Friday". Using these, and given some observations that have been filled up in a new sequence, the successive activities can be predicted.

3.2.4 Uncertainty in moving objects databases

A moving objects database often tracks the location of vehicles and travelers. The objective of uncertainty management is to assist the user in accounting for spatio-temporal uncertainty, and in expressing imprecise queries/triggers. Uncertainty management is often the first step in abstracting sensor-data.

For example of an uncertain query, a trucking company may ask: "Retrieve the current location of the delivery trucks that will possibly be inside a region R, sometime between 3:00PM and 3:15PM". In [22] we proposed operators for uncertain spatio-temporal range queries, and their processing algorithms. Computational geometry played an important role in these algorithms. In [15], the uncertainty is modeled quantitatively by probabilistic values. This enables answering the queries such as: "What is the probability that a given object will be inside a given region R sometime between t_1 and t_2?"

3.2.5 Compression/abstraction of spatio-temporal data

A key observation that lies at the foundation of spatio-temporal data compression is that a GPS point (x, y, t) can be eliminated if (x, y, t) can be approximated with a reasonable accuracy by interpolating the adjacent (i.e., before and after) GPS points. In [6] we formalized this intuition by employing a mechanism based on line simplification, which has been studied in computational geometry, cartography and computer graphics. Basically, line simplification approximates a polygonal line by another that is "sufficiently close", and has less straight-line segments (or points). The attractiveness of line simplification (compared to other lossy data compression techniques such as wavelets) stems from the fact that the approximation carries a given error bound. However, we discovered that although the approximation error is bounded, the error of the answers to queries may not be bounded. Whether or not it is bounded, depends on the combination of the distance function (or distance for short) used in the approximation, and the spatio-temporal query type. Furthermore, we considered an aging mechanism by which a trajectory is represented by increasingly coarser abstractions as time progresses.

4 Mobile P2P Data Management

4.1 Applications

A mobile peer-to-peer (MP2P) database resides on a set of mobile peers that communicate with each other via short-range wireless protocols, such as IEEE 802.11 and Bluetooth. The peers communicate reports (data) and queries to neighbors directly, and the reports and queries propagate by transitive multi-hop transmissions. The mobile P2P database provides various types of real-time traveler information, including safety alerts (e.g., a car in front with a malfunctioning brake light), traffic conditions (possibly represented by multimedia clips), accidents, transit-vehicle on-time performance, parking availability, ride share opportunities, audio and video clips of traffic conditions, etc. These applications are more amenable to mobile P2P dissemination than client/server because of real-time safety requirements. Also, since the network id's of the destination nodes are often unknown, the limited transmission range of the network is used instead.

Observe that the MP2P database stores spatio-temporal information, and that the query/trigger processing algorithm does not guarantee to produce all the answers. However, as more data and queries are being disseminated, the throughput and response time improve.

4.2 Research Issues and Approaches

Relevant research issues include: 1. query processing in a distributed and mobile ad hoc environment without data-directory services; 2. query processing with storage, bandwidth, and energy constraints, particularly in managing multimedia data; 3. integration of mobile P2P communication and infrastructure communication; 4. analysis of how reports are propagated in space and time.

4.2.1 Query processing with resource constraints

Given energy and bandwidth constraints, prioritization of reports is critical, so that maximum useful traveler information is disseminated. A useful report is one that has an impact on the decision making process of the receiver. For example, in a travel-time dissemination application, a report is useful if it changes the shortest-route of a vehicle. In a parking space discovery application, a report is useful if it leads to a successful occupation of the reported parking space. The usefulness of a report can be estimated by a mobile node based on its characteristics such as the age, distance, and application. In [19], we demonstrated the feasibility and advantages of using online machine learning algorithms to determine the relevance (i.e. probability of being useful) of a report. The general idea is to use the received reports as an input to a supervised machine learning process. The interesting aspect is that the supervision can be carried out automatically, without user intervention. For example, an algorithm running on a vehicle can make judgments regarding the usefulness of a report. Over time, each vehicle learns a model that can estimate the relevance of a report, and the model can then be used as a ranking function.

4.2.2 Querying blobs in mobile P2P databases

When data reports include binary large objects (blobs) such as video/voice clips, query processing in MP2P databases becomes even more challenging. However, for the purpose of matchmaking, namely finding the reports that satisfy a query, only the metadata descriptions of the reports are needed. For example, the metadata of a multimedia clip may simply include the time and location at which the clip was produced. In such an environment, the design choices for query processing can be made along multiple dimensions. One dimension considers that the mobile P2P communication may use purely short-range or it may use both short-range and cellular communication. Another dimension considers that, due to size-differences, the metadata and blob sub-reports of a given report may be disseminated independently, and by different means. In [26] we analyzed the blob query-processing strategies along these dimensions.

4.2.3 Propagation Analysis

The fundamental question that the propagation analysis answers is: given a mobile P2P database and its properties in terms of peer density, mobility, dissemination method, and storage/bandwidth/energy constraints, what is the probability that a peer at location (x, y) has received a report R that was generated at location (x_0, y_0) t time units ago. This probability is useful for multiple purposes. For example, it helps ranking because if this probability is high then the rank of R should be lowered. It also enables the comparison of various dissemination methods without having to conduct simulations.

Epidemiology renders some methodology to propagation analysis since mobile P2P dissemination is to some extent similar to the spreading of epidemic diseases. However, mobile P2P dissemination is much more complex because many diseases (reports) are spread at the same time, and the spreading of one interferes with the other (competing for bandwidth and storage). In [20] we conducted an analysis with bandwidth and storage constraints taken into account. The analysis is based purely on the age of a report, and extension to both age and distance remains an open problem.

5 Social Networks

Social networks can be used to crowd-source information and to disseminate a variety of valuable spatio-temporal information to the traveler and the traffic administrator. For example, traffic congestion, parking slot availability, ride sharing opportunities, and cooperative driving information can be obtained and disseminated via (possibly ad hoc) social networks.

Relevant research issues include: 1. providing incentives for participation; 2. quality of information supplied by non-liable crowds; 3. mitigating privacy loss when crowds have to disclose their locations [7].

Now consider incentives. In general, there are three approaches to incentivizing users to participate in crowd-sourcing. The first approach, called *pricing incentive*, remunerates participating users using a virtual currency scheme (see e.g., [27]). The decisions to make are who pays, who charges, and how much is paid or charged in various transactions. When crowd-sourcing is implemented in a mobile P2P fashion, an open issue is the atomicity problem, which occurs when a trading transaction between two users needs to be settled even if the users are disconnected before the transaction completes [27]. The second approach, called *soft-incentive* or *non-pricing incentive*, denies services to non-participating users using a reputation scheme (see e.g., [5]. The third approach used in ad hoc networks assumes, based on game theory, that rational users are incentivized to participate (see e.g., [9]). It is not clear whether this approach is feasible in a transportation environment, or can be made feasible without resorting to some form of reputation management.

References

[1] J. Auld, Chad A. Williams, Abolfazl Mohammadian, and Peter C. Nelson. An Automated GPS-Based Prompted Recall Survey With Learning Algorithms. *Transportation Letters*, 1(1), Jan. 2009, pp. 59-79.

[2] D. Ayala, J. Lin, O. Wolfson, N. Rishe, and M. Tanizaki. Communication Reduction for Floating Car Data-based Traffic Information Systems. *GeoProcessing*, 2010.

[3] J. Biagioni, A. Agresta, T. Gerlich, and J. Eriksson. Transitgenie: A real-time, context-aware transit navigator (demo abstract). In *SenSys*, pages 329-330, ACM, 2009.

[4] J. Booth, A. P. Sistla, O. Wolfson, and I. Cruz. A Data Model and Query Language for Urban Transport Systems. *12th Int. Conf. on Extended Database Technology*, Mar. 2009, pages 994-1005.

[5] S. Buchegger and J. Boudec. A Robust Reputation System for Peer-to-Peer and Mobile Ad-hoc Networks. *Proceedings of P2PEcon*, Harvard University, Cambridge MA, U.S.A., June 2004.

[6] H. Cao, O. Wolfson, G. Trajcevski. Spatio-Temporal Data Reduction with Deterministic Error Bounds. *The VLBD Journal*, Vol.15(3), Sept. 2006, pp. 211-228.

[7] C. Cottrill. Approaches to Privacy Preservation in Intelligent Transportation Systems and Vehicle-Infrastructure Integration Initiative. In *Transportation Research Record*, 2129:9-15, 2010.

[8] S. Dashtinezhad, T. Nadeem, C. Liao, and L. Iftode. Trafficview: A scalable traffic monitoring system. *IEEE Int. Conf. on Mobile Data Management*, January 2004.

[9] M. Felegyhazi, J. Hubaux, and L. Buttyan. Nash Equilibria of Packet Forwarding Strategies in Wireless Ad Hoc Networks. *IEEE Transactions on Mobile Computing*, 5(5), 2006, pp. 463-476.

[10] B. Hull, V. Bychkovsky, Y. Zhang, K. Chen, M. Goraczko, E. Shih, H. Balakrishnan, and S. Madden. CarTel: A Distributed Mobile Sensor Computing System. In *Proc. ACM SenSys*, Nov. 2006.

[11] C. Jensen and N. Tradiauskas. Map Matching. *Encyclopedia of Database Systems*, 2009, pp. 1692-1696.

[12] S. Kim, B. George, S. Shekhar. Evacuation route planning: scalable heuristics. *ACMGIS*, 2007.

[13] H. Lin, R. Zito, and M. Taylor. A Review of Travel-time Prediction in Transport and Logistics. Eastern Asia Society for Transportation Studies, 2005.

[14] J. Liu, and O. Wolfson, H. Yin. Extracting Semantic Location from Outdoor Positioning Systems. *Int. Workshop on Managing Context Information and Semantics in Mobile Environments (MCISME)*, 2006.

[15] H. Mokhtar, J. Su. Universal Trajectory Queries for Moving Object Databases. *MDM* 2004.

[16] K. Perumalla. A systems approach to scalable transportation network modeling. *Proceedings of the 38th conference on Winter simulation*, 2006.

[17] D. Pfoser, N. Tryfona, and A. Voisard. Dynamic Travel Time Maps – Enabling Efficient Navigation. *Proc. 18th SSDBM*, 2006, pp. 369-378.

[18] Research and Innovative Technology Administration. *ITS Strategic Research Plan*, 2010-2014. Executive Summary. http://www.its.dot.gov/strat_plan/pdf/ITSStrategicResearch_Jan2010.pdf

[19] P. Szczurek, B. Xu, J. Lin, O. Wolfson. Machine Learning Approach to Report Prioritization with an Application to Travel Time Dissemination. *IWCTS*, 2009.

[20] P. Sistla, O. Wolfson, B. Xu. Opportunistic Data Dissemination in Mobile Peer-to-Peer Networks. *Proceedings of the 9th International Symposium on Spatial and Temporal Databases*, 2005, pp. 346-363.

[21] M. Schubert, M. Renz, and H. Kriegel. Route Skyline Queries: A Multi-Preference Path Planning Approach. *ICDE*, 2010.

[22] G. Trajcevski, O. Wolfson, K. Hinrichs, S. Chamberlain. Managing Uncertainty in Moving Objects Databases. *ACM Transactions on Database Systems (TODS)*, 29(3), Sept. 2004, pp. 463-507.

[23] C. A. Williams, P. C. Nelson, and A. Mohammadian. Attribute Constrained Rules For Partially Labeled Sequence Completion. *Advances in Data Mining - Applications and Theoretical Aspects*, 2009.

[24] L. Wischoff, A. Ebner, H. Rohling, M. Lott, and R. Halfmann. SOTIS - a self-organizing traffic information system. *Vehicular Technology Conference*, 2003.

[25] O. Wolfson, H. Yin. Accuracy and Resource Consumption in Tracking Moving Objects. *SSTD*, 2003.

[26] O. Wolfson, B. Xu, H. Cho. Multimedia Traffic Information in Vehicular Networks. *ACMGIS*, 2009.

[27] B. Xu, O. Wolfson, and N. Rishe. Benefit and Pricing of Spatio-temporal Information in Mobile Peer-to-Peer Networks. *HICSS-39*, 2006.

[28] H. Yin, O. Wolfson. A Weight-Based Map Matching Method in Moving Objects Databases. *The 16th International Conference on Scientific and Statistical Database Management (SSDBM)*, 2004.

[29] B. van Arem, C. van Driel, and R. Visser. The Impact of Cooperative Adaptive Cruise Control on Traffic-Flow Characteristics. *IEEE Transactions on Intelligent Transportation Systems*, (7)4:429-436, 2006.