

Applications of a Transportation Information Architecture

John F. DILLENBURG, Peter C. NELSON, Ouri WOLFSON, Oliver YU, Aravinda Prasad SISTLA, Sue MCNEIL, Aris M. OUKSEL, Bo XU, and Jezekiel BEN-ARIE

Abstract—Basic research in information management, software architecture, human factors, traffic prediction and transportation planning is necessary to manage the integration of millions of disparate computers and sensors, in a high mobility environment, into a collaborative system. This system will then enable solutions to a set of transportation problems ranging from autonomous real-time routing and navigation, to autonomous driving, to inferring driving patterns via data mining. In this paper, this information architecture is discussed together with several innovative applications.

Index Terms—Software, Transportation, Detectors, Vehicles

I. INTRODUCTION

Excess congestion cost the United States economy over \$67 billion in 2000 from fuel and wages alone [1]. Adding more roadways does not appear to be a viable option, either. Only half of the roads necessary to compensate for the increases in congestion were added between 1982 and 2000 [2]. There is not enough space, money, and public approval for all of the roads that would be needed to alleviate congestion.

Steps can be taken to reduce these mobility problems. For instance, an increase of just 4% in the number of ride-sharing vehicles would have been enough to offset 2000's increase in urban congestion [3]. Information technology (IT) providing widely available, real-time, reliable information is the foundation for implementing strategies to improve roadway conditions. This is evident by examining related projects in Japan, Germany, and Singapore.

John F. Dillenburg is with the University of Illinois at Chicago, Chicago, IL 60607 USA (telephone: 312-996-5598, e-mail: dillenbu@uic.edu).

Peter C. Nelson is with the University of Illinois at Chicago, Chicago, IL 60607 USA (telephone: 312-996-3259, e-mail: nelson@cs.uic.edu).

Ouri Wolfson is with the University of Illinois at Chicago, Chicago, IL 60607 USA (telephone: 312-996-6770, e-mail: wolfson@uic.edu).

Oliver Yu is with the University of Illinois at Chicago, Chicago, IL 60607 USA (telephone: 312-996-2318, e-mail: yu@cs.uic.edu).

Aravinda Prasad Sistla is with the University of Illinois at Chicago, Chicago, IL 60607 USA (telephone: 312-996-8779, e-mail: sistla@cs.uic.edu).

Sue McNeil is with the University of Illinois at Chicago, Chicago, IL 60607 USA (telephone: 312-996-9818, e-mail: mneil@uic.edu).

Aris M Ouksel is with the University of Illinois at Chicago, Chicago, IL 60607 USA (telephone: 312-996-0771, e-mail: aris@uic.edu).

Bo Xu is with the University of Illinois at Chicago, Chicago, IL 60607 USA (telephone: 312-996-4881, e-mail: boxu@cs.uic.edu).

Jezekiel Ben-Arie is with the University of Illinois at Chicago, Chicago, IL 60607 USA (telephone: 312-996-2648, e-mail: benarie@ece.uic.edu).

The Vehicle Information and Control System (VICS), a government-sponsored system in Japan with a 10-year track record, is one example of how information technology can be used to improve driving conditions [4]. Via VICS communication networks, drivers receive information on traffic congestion, travel times, accidents, construction, and parking. This allows Japanese drivers to make informed decisions about their travel routes. There are currently over 6 million VICS-enabled cars in Japan.

An inter-vehicle communications system called FleetNet is being developed by a consortium of private companies and universities in Germany [5][6]. FleetNet, funded by Germany's Federal Ministry for Education and Research, focuses on mobile ad hoc radio networks; applications include emergency braking notification, vehicle platooning, and traffic information distribution.

Currently, one of the most advanced information technology-based transportation systems is found in Singapore. It has electronic road pricing, price-based suppression of auto use, real-time surveillance of arterial road speeds, advanced traffic signal control and online traveler information systems, plus an advanced mass transit system.

Intelligent Transportation Systems (ITS) Centers of Excellence are funded by the U.S. Department of Transportation (USDOT). The USDOT also sponsors the Intelligent Vehicle Initiative, which focuses on safety [7], and the National ITS Architecture, a framework for the integration of ITS into existing transportation systems [8].

Computer systems technology is rapidly evolving to a state characterized as "ubiquitous" computing, the situation where processors are so commonplace that few are even aware of their existence. In the future, there will be an almost unimaginable diversity of computational resources available—wireless networks, online databases, embedded processors, powerful sensors, communications satellites, etc.—and "smart" devices of various types will know how to connect with and utilize these diverse resources.

Like today's systems, such devices and their software will need to negotiate the ever-present tradeoff between resource consumption and the quality of results. In other words, as any cellular phone user knows, the availability of resources varies tremendously. Thus, to function effectively, these devices will have to be "tunable"; they need the capability to adjust themselves to the resources at hand, producing better results as more resources become available. And their performance

should not be significantly affected by alternative resource paths, and should degrade gracefully as resources become unavailable.

Computerized transportation systems are the epitome of this ubiquitous computing environment, so it is critical for them to be tunable and adaptable.

II. INFORMATION ARCHITECTURE TOOLKIT

To facilitate the development of a truly general transportation information architecture, a toolkit of very specific advanced capabilities is required. The most important forward-looking capabilities are discussed in this section.

A. *Sensor fusion and computer vision*

This part of the toolkit requires the development of algorithmic and systemic solutions for vehicle and trajectory recognition based upon the fusion of heterogeneous information. Another required objective is to incorporate a capability for machine learning that will gradually improve system performance in accord with experience.

Several theoretical approaches are available for developing methods of information fusion. Alongside traditional Bayes' decision theory [9], which employs conditional probabilities, there are approaches such as Dempster-Shafer theory [10] and Possibility Theory that are founded on fuzzy logic [11] and conjunctive and disjunctive fusion. A comparative evaluation of the utility of these approaches to vehicle recognition applications is needed, however. In addition, a new approach is needed for fusion-based image analysis that differs considerably from conventional approaches that are based on aggregating different but direct measurements of the same image pixels. By contrast, this approach is based on much more advanced representations of image data. The information sources are derived from various advanced image processing and analysis algorithms that yield higher-level descriptions of image and other data sources. For example, employing detailed feature descriptions is more effective than simple gray level data. Such techniques were recently used for face recognition (q.v. [12], where recognition is based on skin color, shape, and Gabor signatures), and for labeling human body parts [13]. It may be possible to attain a much higher degree of statistical independence and therefore much improved recognition from high-level descriptions compared to fusion of low-level pixel data.

Not all of the information sources considered here for vehicle recognition are visual. Visual shape and color information can be augmented with other sensor information, such as the signatures obtained by road-embedded electromagnetic or strain-gauge sensors when vehicles pass over them.

B. *Wireless networking*

Collaborative wireless communications among mobile vehicles and between vehicles and the infrastructure are required by many computerized transportation applications. It is generally perceived that wireless Mobile Ad hoc Networking (MANET) could be adapted to support such

collaborative communication paradigms under which a node (vehicle or fixed node) could send messages to a destination node beyond its wireless transmission range by using other nodes as relays or routers. MANET routing protocols can be either topology-based or location-based. Topology-based routing can be either proactive (e.g., [14]), with each node maintaining updated routing to other nodes through periodic exchange of routing information, or reactive (q.v. [15]), with the routes to destinations being created only when desired by the source node. Location-based routing is characterized by packet forwarding based on the physical locations of nodes (e.g., [16]), so it requires a location-detection mechanism such as GPS, and a location service [17], to enable the source nodes to determine the locations of the destination nodes.

Traditional MANET routing protocols are not suited to the high mobility features of vehicular networks. For deployment in vehicular networks, topology-based routing protocols would require a large number of routing states and incur large routing overheads for updating topology changes. Furthermore, routing is often data-centric with group communications referenced by attribute-based names, rather than being node-centric with communications referenced by global node identifiers. For example, a vehicle would be more interested in querying which locations ahead have average speed below a specific threshold, rather than the average speed of a particular vehicle.

Further research on communication paradigms for vehicular networks that take advantage of emerging standards such as the CALM network architecture [18], and support the mobile point-to-multipoint paradigm via a dynamic adaptive multicast tree based on geographic forwarding is required. A multipoint-to-point paradigm via a different tree with merge points between the source and sink nodes can be supported. The tree-based routing schemes are enabled with reactive localized-rerouting to counteract topology changes. When a node detects that the designated node for tree-based forwarding is inactive, it will switch to the dynamic localized-rerouting (DLR) mode with location-aided forwarding. The use of mobile agent technology (MAT) to support the DLR mode also needs to be investigated.

C. *Traffic prediction*

Travel times and congestion levels on a road link are usually predicted using multiple variables (e.g., day of week, time of day) combined with historical traffic records, real-time loop detector flow and occupancy data, probe vehicle travel times, road construction reports, weather sensors, etc. However, existing work on traffic prediction (see for example [19][20]), routinely ignores or oversimplifies two factors. The first factor is interference between road links. Namely, the conditions on a given link are unlikely to be independent of the conditions on adjacent links. E.g., congestion caused by an accident will block not only the associated link but also the "upstream" links and nearby local streets as well. Initial work in this area includes [21] and [22]. The second factor is traffic caused by special events such as sporting events and rock

concerts.

The effects of interference can be incorporated as follows. For each road link L , the traffic prediction of L is based on a multivariate time series that is merged from the history travel time records of L and those of its neighbors. Here, there are two approaches to predicting the merged multivariate time series. The first utilizes neuro-fuzzy networks (e.g., [23]); the second employs the ARIMA (AutoRegressive Integrated Moving Average) model [24]. For incorporation of special events, a software agent that is able to collect automatically, from web sites, spatial and temporal information on the events is needed. A third approach involves the integration of behavioral information.

In summary, techniques that provide accurate traffic prediction by incorporating road link interference and special events are possible.

D. Cost Models, Auction and Negotiation

Market-based mechanisms will radically alter the economics of resource and service management in the transportation systems of the future. Resources will range from communications bandwidth, to road sensors, to the lanes in a road system, and services will include capacity planning, congestion pricing [25], and/or direct requests like “speed up my travel time for fee or barter.” Traditional centralized approaches (which need complete state information and a common management policy) or decentralized consensus-based policies are inadequate for dealing with the complexity of the transportation system. In contrast with conventional cost models, pricing based on user demand and the supply of resources is the main driver in competitive economic models. Thus, it should be stressed that travelers are in competition with other travelers, and resource and service providers are in competition with other resource and service providers [25].

Robust cost models are needed to estimate demand and trade resources. These models should build on the value pricing experiments of TEA-21 [26], on emissions-trading work, and on basic stock market concepts to develop cost models that account for the costs of obtaining information, travel time, fuel consumption, crashes, and externalities. These in turn can be used to explore more innovative pricing structures that then allow auctions and the trading of resources, plus negotiation among travelers. An example of such negotiation is the trading of lane space or travel time for money. This will lead to a priori high-speed lane designation, and/or the dynamic/spontaneous formation of such lanes. Dynamic high-speed lane formation can be achieved via a centralized, automatic trading system, or via negotiations between neighboring vehicles. Ride sharing cost models should also be analyzed.

Consider travelers who want to transfer to a faster route for the current advertised price. This action could eventually slow down this route, forcing the manager of this resource to increase its price in order to reduce traffic density and to maximize its utility. The new price could be prohibitive for some travelers who could then require compensation, and they

also might transfer to more attractive routes. But the increased demand engendered for these new routes may in turn lead to unacceptable congestion levels, and so on. This chasing around for an equilibrium point lends itself to a whole gamut of game-theoretic approaches [27] and taxing mechanisms, including auctions, bargaining, and arbitration, where various parameters (such as price and current congestion levels) can be used to regulate road system utilization to achieve overall net efficiency for travelers.

Game-theoretic paradigms can use our preliminary work [36][37][38] on managing on-demand resource and service allocation in competitive or priority-differentiated transportation environments. Travelers use a utility model (e.g., a function of the desired congestion level and price), to specify their resource and service requirements and preference parameters, and brokers will provide strategies for choosing appropriate resources that meet user requirements. The brokering approach circumvents the obvious drawbacks of negotiation in real-time situations by requiring that travelers reveal their true interests only to the broker, who then determines solutions with full knowledge of all users’ interests. A broker is part of the distributed mobile service-oriented computing of local brokers, and both broker-broker and broker-traveler interactions follow the Peer-to-Peer (P2P) paradigm [39][40].

The problem of on-demand resource allocation has been studied extensively, and solutions (some Pareto optimal) in both engineering and economics have been proposed [41]. In most proposed solutions, concern is limited to net efficiency measures—measures of the global good or social welfare that accrues to a population of travelers based upon the access policies and cost models of providers, assuming that the travelers reveal their true utility [42][43]. A more realistic assumption is that travelers may elect at times to tilt utility in their favor, out of greed or possible malice. The impact of this type of information asymmetry on overall system performance needs further investigation. In addition, the performance and robustness of our mechanisms over several other dimensions needs to be assessed, including:

- (i) fairness, the extent to which all travelers in symmetric roles are treated equally over repeated interactions;
- (ii) robustness, since a robust arbitration mechanism is impervious to manipulation by a user;
- (iii) privacy and anonymity, where an arbitration scheme does not reveal anybody’s interests;
- (iv) computational efficiency, the time it takes an arbitration scheme to arrive at a decision.

Our earlier mechanisms [36][38] are limited to a specific number of users, and satisfy some or all of these measures, thereby providing a toolkit of broker policies. One challenge now is investigating scaling conditions. E.g., using utility forecasts to provide a priori probabilities to brokering mechanisms planning possible resource allocations. Another challenge is how to use pricing strategies to incorporate supply side and resources allocation [25][26] in a real-time dynamic market place. Resource and service owners need

price generation schemes to rationalize and increase system utilization, and protocols to support competitive price decision-making.

E. Real time maintenance, prediction and generation of spatio-temporal information

A trajectory is fundamental for representing spatio-temporal information about a moving object. A trajectory is modeled as a sequence of (location, time) pairs indicating the location of a moving object at various times. It may represent a future motion plan, or a trace of past motion (e.g., the trace obtained by a GPS receiver). Both the location and the time can be given at different levels of granularity. E.g., for ride sharing the location values could be street intersections, whereas for collision avoidance a location could include the exact position on the lane. Trajectory information can be used in real time as well as for archival and data mining purposes. And sensor values (e.g., readings from shock absorbers) can be associated with each (location, time) pair.

The problem is devising a set of mechanisms and linguistic constructs for easy, natural, and efficient manipulation of trajectories and other spatio-temporal information (such as traffic information generated by static sensors on roads) when developing applications. These mechanisms will include spatio-temporal query languages, indexing, similarity retrieval, and joins for collision course detection and ride sharing. These mechanisms need to be combined into a trajectory data-blade, i.e., an extension of an existing Database Management System (DBMS) that handles the trajectory data type. An example of a data-blade operator is trajectory approximation. Such an operation may be important for data reduction and privacy protection (a user may provide her approximated trajectory, but not the exact one). Line simplification for trajectory approximation is examined in [27], as compared to the wavelet approach (e.g., [28]) and others. Another operator would aggregate trajectories based upon a given similarity function.

An additional functionality is location/trajectory prediction. This is an important operation when a vehicle v needs to know its future trajectory (so that, for example, it can be passed to other vehicles to indicate when/where to deliver information to v). One trivial possibility is for the driver to always enter her destination at the beginning of a trip. However, this is not always feasible, and the driver can often be released of this burden. For example, based upon periodic patterns (e.g., home-to-office travel) the trajectory can be learned automatically. Learning such trajectories necessitates the detection of periodic patterns in GPS traces. One approach to solving this problem is to (i) segment a trace into stationary and moving sections, (ii) translate the GPS points into a finite set of symbols, and (iii) employ an existing symbolic-patterns technique. There has recently been interest in a similar type of learning (e.g., [29][30]), but the problem is still very much open. A different approach to prediction may involve heuristics based on statistical information, and other rules of thumb.

The toolkit will also need to include strategies that address the magnitude of data using data reduction and compression, aggregation, and aging methods. For interoperability with current and future applications the data-blade can be extended to deal with ontology-based semantics for flexible data sharing [31], building upon our recent work.

In summary, the toolkit will need to include a trajectory data-blade that performs spatio-temporal database operations, as well as approximation and prediction. Examples of queries supported in this data-blade include: retrieve the location of the next gas station (which can be sensed by vehicles passing the station), or retrieve the average speed a mile ahead. This data-blade can be built upon the related work [32][33][34]. Spatial databases and temporal databases have been studied separately in the extant literature, but their combination in real-time databases for high mobility is beyond the current state of the art.

III. APPLICATIONS OF THE TOOLKIT

In the sections that follow, three concepts are examined that make use of the information architecture toolkit discussed above. Each of these concepts provides a real-world example of how such a toolkit could be used to improve traffic conditions dramatically.

A. Intelligent Traveler Assistant (ITA)

One application of the information architecture toolkit discussed in the previous sections is the intelligent traveler assistant (ITA) concept. The ITA [44] is a software system running on handheld computers networked to a traffic information center, see Figure 1. It provides the capability to plan multi-modal routes for its user. The traveler will enter a desired destination and any intermediate destinations into the ITA, which will then formulate several plans to get the user to these destinations. These plans will be listed on the ITA, sorted and filtered by the user's saved preferences, together with the associated trip costs, estimated pickup times, and estimated arrival times. Each trip planned may be multi-modal in nature with the following types of trip segments: dynamic ride sharing, taxi sharing, bus, train, ferry, or automobile. The ITA will make use of wireless Internet technology [45] to send and receive real-time traffic information, and to receive up-to-the-minute locations of public transportation vehicles, transit schedules, and arrange for payment of fares. The Global Positioning System (GPS) will be used to track the user's current position for use in route planning and also as a means of assessing traffic conditions.

The ITA will make use of the toolkit for: (i) sensor fusion to merge together travel information, (ii) wireless networking, (iii) traffic prediction to obtain the best routes for the expected traffic conditions, and (iv) cost models. It will be tested within the Gary-Chicago-Milwaukee USDOT priority corridor, which is large and complex enough to ensure that effects of noise, error, and limited resources are properly addressed. The corridor hosts several agencies that provide real-time GPS locations for buses and trains. Its travel

information center is co-located with the Gateway [46] and Ride Match System 21 [47]

The ITA will be distributed as software to users who provide their own devices, wireless connection, and GPS receivers. Several releases will take place, with additional capabilities added with each release. Release 1 will include static routing using a single mode of transport; release 2 will add transit schedules, ride sharing and congestion based routing; release 3 will incorporate dynamic ridesharing, real-time transit information, and use the ITAs as traffic probes. Release 4 will add traffic prediction, car/taxi sharing, and fare negotiation.

Analysis of ITA trajectories and route planning logs can be used to determine success, as measured by (i) the number of positive choices made (e.g., ride sharing vs. a single occupant trip), (ii) the time saved compared to an ad hoc plan, and (iii) anticipated vs. unanticipated waiting time.

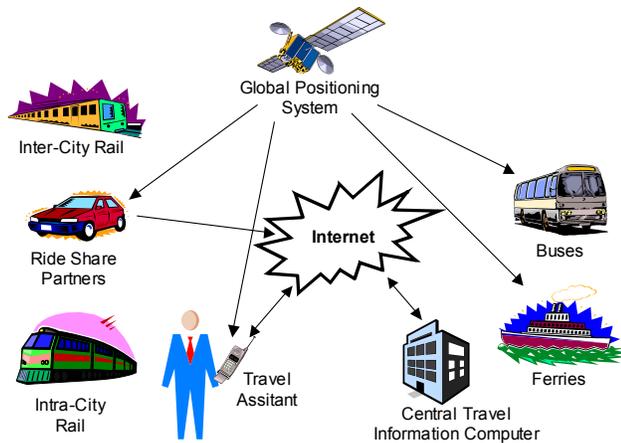


Figure 1. ITA Concept

B. Small Scale Autonomous Driving (SSAD)

The next application of the toolkit is the small scale autonomous driving (SSAD) concept. The problem of autonomous, unmanned, vehicle driving has been studied extensively. There are two basic approaches to solving this problem, both of which are sensor oriented, meaning that the unmanned vehicle is equipped with sensors to detect, understand, and navigate in its environment. In one approach the highways are also instrumented (intelligent highway approach), whereas in the other they are not. The sensors' objective is to detect and navigate around static and moving obstacles. The SSAD concept aims to exploit recent developments in pervasive computing and Mobile Ad hoc Networks (MANETs) to greatly simplify the obstacle detection problem.

Given detailed enough maps that represent the road network, the problem of sensing the environment reduces to that of sensing neighboring vehicles. Neighboring moving objects have the incentive to collaborate in order to avoid each other, and they can communicate via the emerging technology of ad hoc networks. Thus each vehicle can broadcast to its neighbors not only its present location (the information that

the sensors approach aims to detect), but also its motion plan in the future. This enables anticipatory and predictive navigation, which can better avoid congestion and accidents.

Thus, using accurate maps and short distance wireless data communication, each vehicle can navigate based on a detailed database that contains static and dynamic information about the immediate environment. The dynamic information is represented by a trajectory database that models the current location, speed, and motion plan of neighboring vehicles, and is continuously revised based on communication with them. The problem of sensing the environment is greatly simplified, and replaced by the problem of real-time, accurate maintenance of the trajectory database in an unmanned vehicle.

As described, the feasibility of this approach depends on all the vehicles on the road system (or certain designated parts of it) collaborating (i.e., communicating their trajectories). However, if this is not the case, the approach can be employed to enable an advanced-cruise-control (that delegates driving to a human in certain situations). For example, suppose that not all vehicles communicate trajectories, but that all are equipped with a transmitter/receiver of their vehicle ID. This means that each vehicle continuously transmits its ID, and is able to sense every other vehicle that is closer than a certain distance (but not necessarily know its location, speed, or trajectory). In other words, a vehicle can participate in one of two levels of collaboration, high or low. At a high level the vehicle communicates its trajectory, and is capable of autonomous navigation based on trajectories received from neighboring vehicles. At a low level it only communicates its vehicle ID. In this scenario, as long as all the neighboring vehicles are at a high level of collaboration, vehicle navigation is autonomous, as outlined above. However, the receiver in a given vehicle x continuously monitors the environment to detect vehicles at the low level of collaboration. If such is detected, then navigation of x is delegated to the human driver.

The objective here is to determine if and how this database of trajectories can be maintained, in real time, at the level of accuracy that is necessary for autonomous driving. The SSAD concept makes use of the toolkit for wireless networking, mobile data management, real-time control, adaptability, context awareness, resource discovery, and game theory.

C. Multi-Vehicle Cooperation (MVEC)

The last application of the information architecture toolkit to be examined here is the multi-vehicle cooperation (MVEC) concept. The MVEC concept groups together communicating vehicles. The an MVEC equipped vehicle will contain a laptop, a GPS receiver, and vehicle-to-vehicle communications equipment. The objective of the MVEC concept is to demonstrate query processing in a mobile environment. In a group of MVEC vehicles, the vehicles will travel on the highway, in live traffic, in a convoy with an average distance of 200 meters between successive vehicles. They will process queries of the type "what is the average

speed one mile ahead?” Processing such a query involves wireless mobile multi-hop communication, and mobile-database collaboration.

ACKNOWLEDGMENT

We would like to thank David Boyce, Isabel F. Cruz, J. N. Hagstrom, Kazuya Kawamura, Bing Liu, Michael J. Scott, Piyushimita (Vonu) Thakuria, and Clement Yu for their contributions and input to this paper.

REFERENCES

- [1] *Pocket Guide to Transportation*, Bureau of Transportation Statistics, 2003. www.bts.gov/publications/pocketguide.
- [2] David Schrank and Tim Lomax. *The 2001 Urban Mobility Report*. Texas Transportation Institute.
- [3] Anthony Downs. *Stuck in Traffic: Coping With Peak-Hour Traffic Congestion*. Brookings Institution, 1992.
- [4] Vehicle Information and Communication System Center (VICS Center), Japan, www.vics.or.jp/eng/.
- [5] FleetNet—Internet on the Road project, Germany, September 2000–December 2003. www.fleetnet.de.
- [6] W. Franz, R. Eberhardt, T. Luckenbach. *FleetNet—Internet on the Road*. The 8th World Congress on Intelligent Transportation Systems, Sydney, Australia, October 2001
- [7] Intelligent Vehicle Initiative, <http://www.its.dot.gov/ivi/ivi.htm>.
- [8] National ITS Architecture, <http://itsarch.iteris.com/itsarch/index.htm>
- [9] Sonka M., Hlavac, V., and Boyle, R. *Image Processing, Analysis and Machine Vision*. Collection Chapman & Hall Computing. Chapman & Hall, 2-6 Boundary Row, London SE1 8HN, Grande-Bretagne, 1993.
- [10] Shafer, G. *A Mathematical Theory of Evidence*, Princeton University Press, Princeton, New Jersey, 1976.
- [11] Shafer, G. *A Mathematical Theory of Evidence*, Princeton University Press, Princeton, New Jersey, 1976.
- [12] Seshadrinathan M and Ben-Arie J. Detection of Faces by Integration of Evidence. Forthcoming in *4th EURASIP Conf. on Video/Image Processing and Multimedia Communications*, Zagreb, Croatia, July 2003.
- [13] Rajaram .S, Ben-Arie J and Sivalingam, D. Labeling of Human Body Parts. To appear in *4th EURASIP Conf. on Video/Image Processing and Multimedia Communications*, Zagreb, Croatia, July 2003.
- [14] C.E. Perkins and P. Bhagwat. Highly Dynamic Destination-Sequenced Distance-Vector Routing (DSDV) for Mobile Computers. *Comp. Comm. Rev.*, Oct. 1994, 234–244.
- [15] M. Jiang, J. Li, Y.C. Tay. Cluster Based Routing Protocol. IETF Draft, August 1999.
- [16] Y.B. Ko and N.H. Vaidya. Location-Aided Routing (LAR) in Mobile Ad Hoc Networks. *ACM/Baltzer WINET J.*, Vol. 6, No. 4, 2000, 307–321.
- [17] O. Yu. End-to-end Dynamic Adaptive QoS Provisioning over GPRS Wireless Mobile Network. Forthcoming June 2003, *Journal of Mobile Networks and Applications*.
- [18] Knut Evenson. CALM (Continuous Air interface for Long and Medium distance) project, Norway, 2003. <http://www.magcom.no/Tolling/chicago/KEv%20Final%20paper%20154%20CALM%20-%20The%20new%20ITS%20Communication%20concept%20r21.pdf>.
- [19] L. Chen and F.Y. Wang. A neuro-fuzzy system approach for forecasting short-term freeway traffic flows. *Proc. IEEE 5th Intl. Conf. on Intelligent Transportation Systems*, Singapore, September 2002.
- [20] L. Fu and B. Hellinga. Real-time, Adaptive Prediction of Incident Delay for Advanced Traffic Management and Information Systems. *Proc. Ann. Conf. Canadian Inst. Trans. Eng.*, Ottawa, Canada, May 12–15, 2002.
- [21] E. Torres, P. Nelson, and N. Roupail. Estimating Link Delays for Arterial Streets. *Urban and Regional Transportation Modeling: Essays in Honor of David Boyce*, Ed. Der-Hong Lee, Edward Elgar, 2003.
- [22] C. Zhou and P.C. Nelson. Predicting Traffic Congestion Using Recurrent Neural Networks. *The 9th World Congress on Intelligent Transport Systems*, Chicago, October 14–18, 2002.
- [23] L. Chen and F.Y. Wang. A neuro-fuzzy system approach for forecasting short-term freeway traffic flows. *Proc. IEEE 5th Intl. Conf. on Intelligent Transportation Systems*, Singapore, September 2002.
- [24] D. Montgomery, L. Johnson, and J. Gardiner. *Forecasting and Time Series Analysis*. McGraw Hill, 1990.
- [25] W. Vickrey. *Principles of Efficient Congestion Pricing*, www.hhh.umn.edu/centers/slp/projects/concpric/.
- [26] DeCorla-Souza, Patrick. *The Price of Congestion: Managing Our Congested Highways with Value Pricing*. Transportation Management + Engineering, Vol 7, No. 4, 22–25, Aug./Sept. 2002.
- [27] H. Cao, O. Wolfson, and G. Trajcevski. *Spatio-temporal Data Reduction in Moving Objects Databases*. Submitted for publication
- [28] K. Chakrabarti, M. Garofalakis, R. Rastogi, and K. Shim. *Approximate Query Processing Using Wavelets*. Proc. of the 26th Intl. Conf. on Very Large Data Bases, September 2000.
- [29] W. Sethares and T. Staley. *Periodicity Transforms*. IEEE Trans. Sig. Processing, 47(11), 2953–2964, 1999.
- [30] W.G. Aref, M.G. Elfeky, A.K. Elmagarmid. *Incremental, Online and Merge Mining of Partial Periodic Patterns in Time-Series Databases*. Purdue Technical Report, 2001
- [31] A.U. Frank. *A Specification Language for Interoperable GIS*. In *Interoperating Geographic Information Systems*, M. Goodchild, et. al. Eds., Kluwer, 1998.
- [32] O. Wolfson, P. Sistla, S. Chamberlain, and Y. Yesha. *Updating and Querying Databases that Track Mobile Units*. Invited paper, special issue of the Distributed and Parallel Databases J., 7(3), 257–288, 1999.
- [33] G. Trajcevski, O. Wolfson, H. Cao, H. Lin, F. Zhang, and N. Rishe. *Managing Uncertain Trajectories of Moving Objects with DOMINO*. Proc. 4th Intl. Conf. On Enterprise Inf. Sys., 217–224, Spain, April 2002.
- [34] R. H. Gueting, M. H. Boehlen, M. Erwig, C. S. Jensen, N. A. Lorentzos, M. Schneider, and M. Vazirgiannis. *A Foundation for Representing and Querying Moving Objects*. ACM-TODS, Vol. 25, No. 1, 2000
- [35] Gardner, R. *Games for Business and Economics*. John Wiley & Sons, 2003.
- [36] Ouksel, A. M., Y. Babad, and T. Tesch. *Agent-Based Regulatory Mechanisms in B2B Exchanges*. Submitted to Management Science, 2003.
- [37] Tesch, T. *Arbitration Protocols for Software Agents*. Ph.D. dissertation, Feb. 2002, Technischen Universitat Darmstadt, Aris M. Ouksel, Co-Supervisor.
- [38] Thomas, T., P. Fankhauser, and A. M. Ouksel. *Arbitration Protocols for Competing Software Agents*. Intl. J. of Cooperative Information Systems, 10(1), 125–143, 2001
- [39] Moro, G., A. M. Ouksel, and C. Sartori. *Agents and Peer-to-Peer Computing: A Natural Combination*. In *Agents and Peer-to-Peer Computing*, Springer Verlag, Lecture Notes in Computer Science, 2530, 1–12, 2003.
- [40] Sullivan, W. T. III et al.. *A new major SETI project based on Project Serendip data and 100,000 personal computers*. Proc. 5th Intl. Conf. on Bioastronomy, 1999. setiathome.ssl.berkeley.edu/woody_paper.html
- [41] Buyya, R., Abramson, D., and Giddy, J. *A Case for Economy Grid Architecture for Service-Oriented Grid Computing*. Proc. Intl. Parallel and Distributed Processing Symp.: 10th Intl. Heterogeneous Computing Workshop (HCW 2001), San Francisco, April 23, 2001.
- [42] Stonebraker, M. et al.. *An Economic Paradigm for Query Processing and Data Migration in Mariposa*. Proc. 3rd Intl. Conf. on Parallel and Distributed Information Systems, 28–30, 1994.
- [43] Sandholm, T. *Distributed Rational Decision Making*. In *Multi-Agent Systems: A Modern Introduction to Distributed Artificial Intelligence*, G. Weiss (Ed), The MIT Press, USA, 2000.
- [44] J.F. Dillenburg, O. Wolfson, and P.C. Nelson. *The Intelligent Travel Assistant*. The 5th Intl. Conf. on Intelligent Transportation Systems, 691–696, Singapore, September 2002.
- [45] K. Biesecker. *Broadband Wireless, Integrated Services, and Their Application to Intelligent Transportation Systems*. Center for Telecommunications and Advanced Technology, McLean, Virginia, June 2000.
- [46] *The Gary-Chicago-Milwaukee home page*, <http://www.gcmtravel.com>.
- [47] J.F. Dillenburg and José Rodriguez. *Ridematch System 21: Using Vector Based Matching to Direct Ridesharing Activities*. Metropolitan Conf. on Public Transportation Research (MCPTR 2002).