Spatial Big-Data Challenges Intersecting Mobility and Cloud Computing

Shashi Shekhar, Michael R. Evans, Viswanath Gunturi, KwangSoo Yang

{shekhar,mevans,gunturi,ksyang}@cs.umn.edu
Computer Science & Eng. Faculty, University of Minnesota
200 Union Street S.E. #4192, Minneapolis, MN 55455

1. INTRODUCTION

Mobility is efficient, safe and affordable travel in our cities, towns and other places of interest [52]. Mobility services, e.g., routing and navigation, are a set of ideas and technologies that facilitate understanding the geo-physical world, knowing and communicating relations to places in that world, and navigating through those places. The transformational potential of mobility services is already evident. From Google Maps [17] to consumer Global Positioning System (GPS) devices, society has benefited immensely from mobility services and technology. Scientists use GPS to track endangered species to better understand behavior, and farmers use GPS for precision agriculture to increase crop yields while reducing costs. We've reached the point where a hiker in Yellowstone, a biker in Minneapolis, and a taxi driver in Manhattan know precisely where they are, their nearby points of interest, and how to reach their destinations.

Increasingly, however, the size, variety, and update rate of mobility datasets exceed the capacity of commonly used spatial computing and spatial database technologies to learn, manage, and process the data with reasonable effort. Such data is known as Spatial Big Data (SBD). We believe that harnessing SBD represents the next generation of mobility services. Examples of emerging SBD datasets include temporally detailed (TD) roadmaps that provide speeds every minute for every road-segment, GPS trace data from cell-phones, and engine measurements of fuel consumption, greenhouse gas (GHG) emissions, etc. SBD has transformative potential. For example, a 2011 McKinsey Global Institute report estimates savings of "about $600 billion annually by 2020" in terms of fuel and time saved [26,29] by helping vehicles avoid congestion and reduce idling at red lights or left turns. Preliminary evidence for the transformative potential includes the experience of UPS, which saves millions of gallons of fuel by simply avoiding left turns (Figure 1(a)) and associated engine idling when selecting routes [26]. Immense savings in fuel-cost and GHG emission are possible if other fleet owners and consumers avoided left-turns and other hot spots of idling, low fuel-efficiency, and congestion. Ideas advanced in this paper may facilitate '*eco-routing*' to help identify routes that reduce fuel consumption and GHG emissions, as compared to traditional route services reducing distance travelled or travel-time. It has the potential to significantly reduce US consumption of petroleum, the dominant source of energy for transportation (Figure 1(b)). It may even reduce the gap between domestic petroleum consumption and production (Figure 1(c)), helping bring the nation closer to the goal of energy independence.

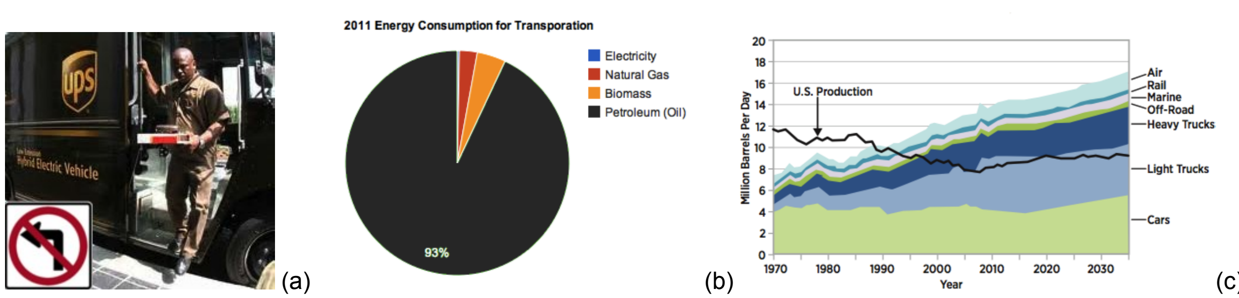

Figure 1: (Left) UPS avoids left-turns to save fuel [26]. (Middle) Petroleum is dominant energy source for US Transportation [49]. (Right) Gap between US petroleum consumption and production levels is large and growing [3,9]. (Best in color)

However, SBD raises new challenges for the state of the art in spatial computing for mobility services such as routing. First, it requires a change in frame of reference, moving from a global snapshot perspective to the perspective of the individual object traveling through a road network. Second, SBD increases the impact of the partial nature of traditional route query specification. It significantly increases computation cost due to the tremendous growth in the set of preference functions beyond travel-distance and travel-time to include fuel consumption, GHG emissions, travel-times for thousands of possible start-times, etc. Third, the growing diversity of SBD sources makes it less likely that single algorithms, working on specific spatial datasets, will be sufficient to discover answers appropriate for all situations. Other challenges include geo-sensing, privacy, prediction, etc.

2. TRADITIONAL MOBILITY SERVICES

Traditional mobility services utilize digital road maps [18, 31, 33, 42]. Figure 2(a) shows a physical road map and Figure 2(b) shows its digital, i.e., graph-based, representation. Road intersections are often modeled as vertices and the road segments connecting adjacent intersections are represented as edges in the graph. For example, the intersection of `SE 5th Ave' and `SE University Ave' is modeled as node N1. The segment of `SE 5th Ave' between `SE University Ave' and `SE 4th Street' is represented by the edge N1-N4. The directions on the edges indicate the permitted traffic directions on the road segments. Digital roadmaps also include additional attributes for road-intersections (e.g., turn restrictions) and road-segments (e.g., centerlines, road-classification, speed-limit, historic speed, historic travel time, address-ranges, etc.) Figure 2(c) shows a tabular representation of the digital road map. Additional attributes are shown in the node and edge tables respectively. For example, the entry for edge E1 (N1-N2) in the edges table shows its speed and distance. Such datasets include roughly 100 million (108) edges for the roads in the U.S.A. [31].

(a)

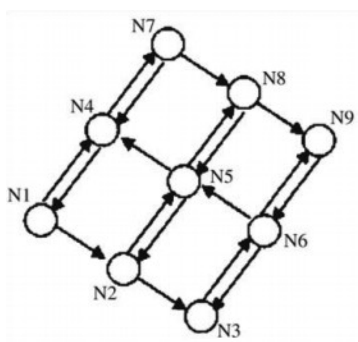

(b)

Nodes

NID
N1
N2
N3
N4
N5
N6
N7
N8
N9

Edges

EID	From	To	Speed	Distance
E1	N1	N2	35mph	0.075mi
E2	N1	N4	30mph	0.075mi
E3	N2	N3	35mph	0.078mi
E4	N2	N5	30mph	0.078mi
E5	N3	N6	30mph	0.077mi
E6	N4	N1	30mph	0.075mi
E7	N4	N7	30mph	0.078mi
E8	N5	N2	30mph	0.078mi
...	...	...	...	...

(c)

Figure 2: Current representations of road maps as directed graphs with scalar travel time values. (Left) Example road map. [17] (Middle) Graph Representation. (Right) Tabular representation of digital road maps.

Route determination services [28, 44], abbreviated as routing services, include the following two services: best-route determination and route comparison [40]. The first deals with determination of a best route given a start location, end location, optional waypoints, and a preference function. Here, choice of preference function could be: fastest, shortest, easiest, pedestrian, public transportation, avoid locations/areas, avoid highways, avoid toll ways, avoid U-turns, and avoid ferries. Route finding is often based on classic shortest path algorithms such as Dijktra's [23], A* [8], hierarchical [19, 20, 41, 43], materialization [37, 39, 41], and other algorithms for static graphs [4, 6, 7, 12, 14, 34, 38]. Shortest path finding is often of interest to tourists as well as drivers in unfamiliar areas. In contrast, commuters often know a set of alternative routes between their home and work. They often use an alternate service to compare their favorite routes using real-time traffic information, e.g., scheduled maintenance and current congestion. Both services return route summary information along with auxiliary details such as route maneuver and advisory information, route geometry, route maps, and turn-by-turn instructions in an audio-visual presentation media.

3. EMERGING SPATIAL BIG DATA

SBD are significantly more detailed than traditional digital roadmaps in terms of attributes and time resolution. In this subsection we describe three representative sources of SDB that may be harnessed in next generation routing services.

Spatio-Temporal Engine Measurement Data: Many modern fleet vehicles include rich instrumentation such as GPS receivers, sensors to periodically measure sub-system properties, and auxiliary computing, storage and communication devices to log and transfer accumulated datasets [21, 22, 27, 30, 46, 47]. Engine measurement datasets may be used to study the impacts of the environment (e.g., elevation changes, weather), vehicles (e.g., weight, engine size, energy-source), traffic management systems (e.g., traffic light timing policies), and driver behaviors (e.g., gentle acceleration/braking) on fuel savings and GHG emissions.

These datasets may include a time-series of attributes such as vehicle location, fuel levels, vehicle speed, odometer values, engine speed in revolutions per minute (RPM), engine load, emissions of greenhouse gases (e.g., CO2 and NOX), etc. Fuel efficiency can be estimated from fuel levels and distance traveled as well as engine idling from engine RPM. These attributes may be compared with geographic contexts such as elevation changes and traffic signal patterns to improve understanding of fuel efficiency and GHG emission.

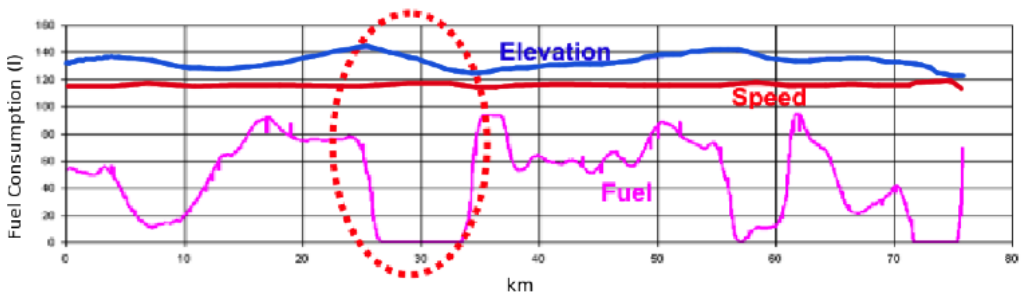

Figure 3: Engine measurement data improves understanding of fuel consumption [5]. (Best in color)

For example, Figure 3 shows heavy truck fuel consumption as a function of elevation from a recent study at Oak Ridge National Laboratory [5]. Notice how fuel consumption changes drastically with elevation slope changes. Fleet owners have studied such datasets to fine-tune routes to reduce unnecessary idling [1, 2]. It is tantalizing to explore the potential of this dataset to help consumers gain similar fuel savings and GHG emission reduction. However, these datasets can grow big. For example, measurements of 10 engine variables, once a minute, over the 100 million US vehicles in existence [11, 45], may have 10^{14} data-items per year.

GPS Trace Data: A different type of data, GPS trajectories, is becoming available for a larger collection of vehicles due to rapid proliferation of cell-phones, in-vehicle navigation devices, and other GPS data logging devices [15, 54] such as those distributed by insurance companies [51]. Such GPS traces allow indirect estimation of fuel efficiency and GHG emissions via estimation of vehicle-speed, idling and congestion. They also make it possible to make personalized route suggestions to users to reduce fuel consumption and GHG emissions. For example, Figure 4 shows 3 months of GPS trace data from a commuter with each point representing a GPS record taken at 1 minute intervals, 24 hours a day, 7 days a week. As can be seen, 3 alternative commute routes are identified between home and work from this dataset. These routes may be compared for idling, which are represented by darker (red) circles. Assuming the availability of a model to estimate fuel consumption from speed profile, one may even rank alternative routes for fuel efficiency. In recent years, consumer GPS products [15, 48] are evaluating the potential of this approach.

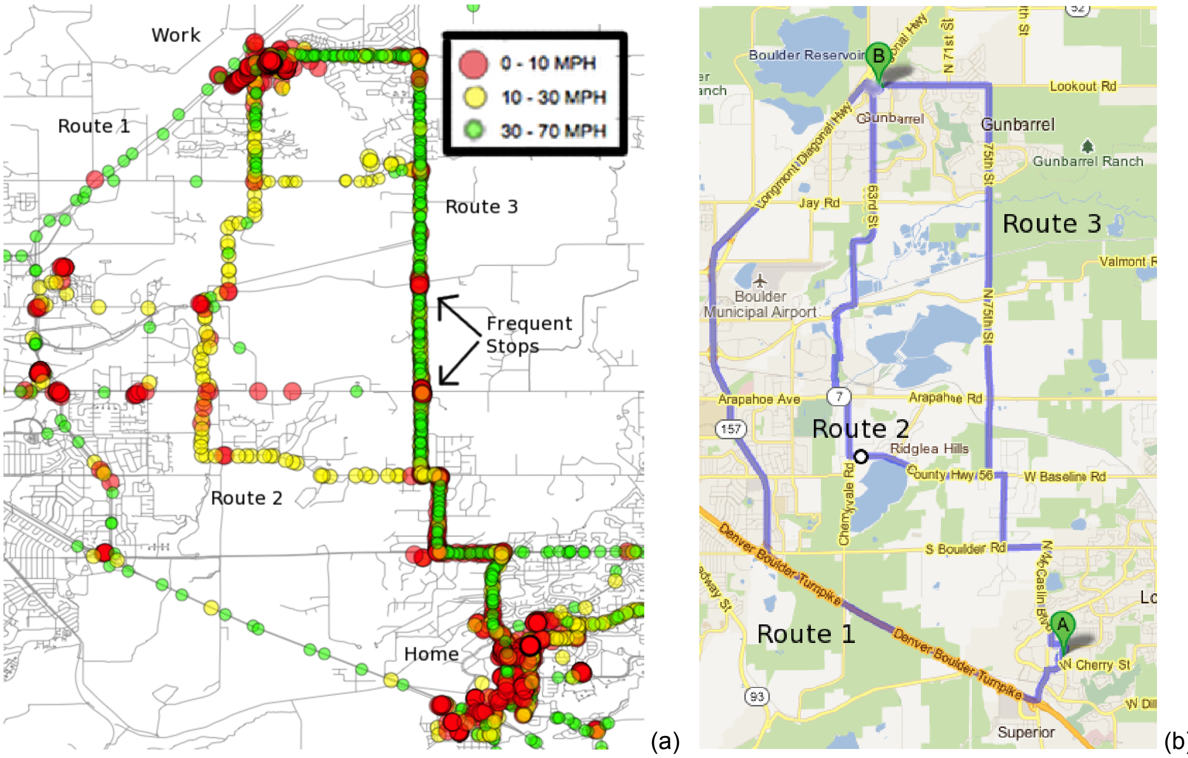

Figure 4: A commuter's GPS tracks over three months reveals preferred routes. (Best in color)

Historical Speed Profiles: Traditionally, digital road maps consisted of centerlines and topologies of the road networks [16, 42]. These maps were used by navigation devices and web applications such as Google Maps [17] to suggest routes to users. New datasets from companies such as NAVTEQ [31] use probe vehicles and highway sensors (e.g., loop detectors) to compile travel time information across road segments for all times of the day and week at fine temporal resolutions (seconds or minutes). This data is applied to a profile model, and patterns in the road speeds are identified throughout the day. The profiles have data for every five minutes, which can then be applied to the road segment, building up an accurate picture of speeds based on historical data. Such TD roadmaps contain much more speed information than traditional roadmaps. Traditional roadmaps (Figure 2(a)) have only one scalar value of speed for any given road segment (e.g., EID 1). In contrast, TD roadmaps may potentially list speed/travel time for a road segment (e.g., EID 1) for thousands of time points (Figure 5(a)) in a typical week. This allows a commuter to compare alternate start-times in addition to alternative routes. It may even allow comparison of (start-time, route) combinations to select distinct preferred routes and distinct start-times. For example, route ranking may differ across rush hour and non-rush hour and in general across different start times. However, TD roadmaps are big and their size may exceed 10^{13} items per year for the 100 million road-segments in the US when associated with per-minute values for speed or travel-time. Thus, industry is using speed-profiles, a lossy compression based on the idea of a typical day of a week, as illustrated in Figure 5(b), where each (road-segment, day of the week) pair is associated with a time-series of speed values for each hour of the day.

In the near future, values for the travel time of a given edge and start time will be a distribution instead of scalar. For example, analysis of GPS tracks may show that travel-time for a road-segment is not unique, even for a given start-time of a typical week. Instead, it may consist of different values (e.g., 1, 2, 3 units), with associated frequencies (e.g., 10, 30, 20). Emergence of such SBD may allow comparison of routes, start-times and (route, start-time) combinations for statistical distribution criteria such as mean and variance. We also envision richer temporal detail on many preference functions such as fuel cost. Other emerging datasets include those related to pot-holes [35], crime reports [36], and social media reports of events on road networks [50].

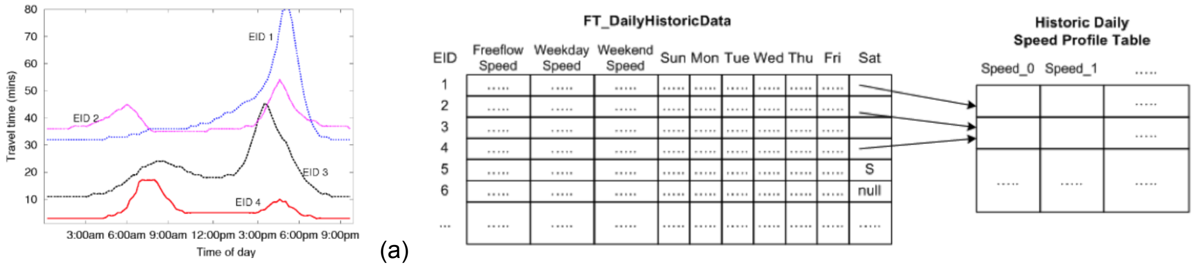

(a) (b)

Figure 5: Spatial Big Data on Historical Speed Profiles. (Left) Travel time along four road segments over a day. (Right) Schema for daily historic speed data. (Best in color)

4. NEW CHALLENGES

New generation of mobility services (e.g., Eco-routing) leveraging SBD raise significant new challenges for state of the art spatial computing. First, it requires a change in frame of reference from a snapshot perspective to the perspective of the individual traveling through a transportation network [53]. For instance, consider the new temporally detailed (TD) roadmaps providing historical travel-time (or speed) for each road-segment for every distinct minute of a week. Consider a person sitting in a vehicle and moving along a chosen path in a TD roadmap. She would experience a different road-segment and its historical speed as well as traversal-time at different time-intervals, which may be distinct from the start-time.

Second, the growing diversity of SBD significantly increases computational cost because it magnifies the impact of the partial nature and ambiguity of traditional routing query specification. Typically, a routing query is specified by a starting location and a destination. Traditional routing services would identify a small set of routes based on limited route properties (e.g., travel-distance, travel-time (historical and current)) available in traditional digital roadmap datasets. In contrast, SBD face orders of magnitude richer information, more preference functions (e.g., fuel efficiency, GHG emission, safety, etc.) and correspondingly larger sets of choices. New questions thus arise in context of eco-routing: What is the computational structure of determining routes that minimize fuel consumption and GHG emissions? Does this problem satisfy the assumptions behind traditional shortest-path algorithms (e.g., stationary ranking of alternative routes assumed by a dynamic programming principle)? For example, temporally detailed roadmaps can potentially provide a distinct route for every possible start-time, even when we just consider travel-time. This raises an optimality challenge of correctly determining the fastest route corresponding to each start-time, since ranking of candidate routes might vary with time of day (rush hour vs. non-rush hour). It also raises a representation challenge to summarize potentially large sets of routes in the result. There is also the computational challenge of efficiently determining a large collection of routes (e.g., one for each start time and preference function) by identifying and reducing unnecessary computations perhaps leveraging current cloud computing paradigm (e.g., map reduce) or via novel custom cloud computing paradigms, tentatively called spatial cloud computing.

Third, the tremendous diversity of SBD sources substantially increases the need for diverse solution methods. For example, methods for determining fuel efficient routes that leverage engine measurement and GPS track datasets may be quite different from algorithms to identify minimal travel-time routes for a given start-time exploiting TD roadmaps. In addition, SBD data (e.g., TD roadmaps, GPS-tracks and engine-measurement datasets) differ in coverage, roadmap attributes and statistical details. For example, TD roadmaps cover an entire country, but provide mean travel-time for a road-segment for a given start-time in a week. In contrast, GPS-track and engine-measurements have smaller coverage to well-travelled routes and time-periods, but may provide a richer statistical distribution of travel-time for each road-segment, perhaps revealing newer patterns such as seasonality. New algorithms are likely to emerge as new SBD become available and as a result, a new, extensible, architecture will be needed to rapidly integrate new datasets and associated algorithms.

Fourth, another challenge area is in the use of geospatial reasoning and SBD in sensing and inference across space and time. Multiple tradeoffs (including those arising in privacy considerations) can come to the fore with attempts to sense and draw inferences from stable or mobile sensors. New challenges arise from crowd-sourced sensors. For example, the ubiquity of mobile phones presents an incredible opportunity for gathering information about all aspects of our world and the people living in it [24]. Already research has shown the potential for mobile phones with built-in motion detectors carried by everyday users to detect earthquakes mere seconds after they begin [13]. Navigation companies frequently utilize mobile phone records to estimate traffic levels on busy highways [50]. How can computers efficiently utilize this prevalent sensing power of mobile phones without drastically impacting battery life or personal privacy concerns? This raises many computer science questions related to sensor placement, configuration, etc.

Fifth, privacy of geographic information inside SBDs is an important challenge. While location information (GPS in phones and cars) can provide great value to users and industry, streams of such data also introduce spooky privacy concerns of stalking and geo-slavery [10]. Computer science efforts at obfuscating location information to date have largely yielded negative results. Thus, many individuals hesitate to indulge in mobile commerce due to concern about privacy of their locations, trajectories and other spatio-temporal personal information [25]. Spatio-temporal computing research is needed to address many questions such as the following: "whether people reasonably expect that their movements will be recorded and aggregated..."? [32]. How do we quantify location privacy in relation to its spatio-temporal precision of measurement? How can users easily understand and set privacy constraints on location information? How does quality of location-based service change with variations in obfuscation level?

Sixth, SBD can also be used to make predictions about a broad range of issues including the next location of a car driver, the risk of forthcoming famine or cholera, or the future path of a hurricane. Such predictions would challenge the best of machine learning and reasoning algorithms, including directions with geospatial time series data. Many current techniques assume independence between observations and stationarity of phenomena. Novel techniques accounting for spatial auto-correlation and non-stationarity may enable more accurate predictions. How can new techniques remain computationally efficient while incorporating auto-correlation and non-stationarity while remaining computationally efficient?

5. CONCLUSIONS

Increasingly, mobility datasets are of a size, variety, and update rate that exceed the capability of spatial computing technologies. This paper addresses the emerging challenges posed by such datasets, which we call Spatial Big Data (SBD), specifically as they apply to mobility services (e.g., transporation and routing). SBD examples include trajectories of cell-phones and GPS devices, vehicle engine measurements, temporally detailed (TD) road maps, etc. SBD has the potential to transform society. A recent McKinsey Global Institute report estimates that personal location data could save consumers hundreds of billions of dollars annually by 2020 by helping vehicles avoid congestion via next-generation mobility services such as eco-routing. Eco-routing may leverage various forms of Spatial Big Data to compare routes by fuel consumption or greenhouse gas (GHG) emissions rather than total distance or travel-time.

However, the envisaged SBD-based next-generation mobility services pose several challenges for current routing techniques. First, SBD requires a change in frame of reference, moving from a global snapshot perspective to the perspective of an individual object traveling through a transportation network. Second, SBD magnifies the impact of partial information and ambiguity of traditional routing queries specified by a start location and an end location. For example, traditional routing identifies a unique (or a small set of) route(s), given historical and current travel-times. In contrast, SBD may identify a much larger set of solutions, e.g., one route each for thousands of possible start-times in a week, significantly increasing computational costs. Third, SBD challenges the assumption that a single algorithm utilizing a specific dataset is appropriate for all situations. The tremendous diversity of SBD sources substantially increases the diversity of solution methods. For example, methods for determining fuel efficient routes leveraging engine measurement and GPS track datasets may be quite different from algorithms used to identify minimal travel-time routes exploiting temporally detailed roadmaps. Newer algorithms will be needed as new SBD becomes available, creating demand for a flexible architecture to rapidly integrate new datasets and associated algorithms. Other challenges include geo-sensing, privacy, prediction, etc.

6. ACKNOWLEDGEMENTS

We would like to thank Eric Horvitz (Microsoft), the Computing Community Consortium (CCC), Hillol Kargupta (UMBC), Erik Hoel (ESRI), Oak Ridge National Labs (ORNL) and the US-DoD for their helpful comments and support.

7. REFERENCES

1. American Transportation Research Institute (ATRI). ATRI and FHWA release bottleneck analysis of 100 freight significant highway locations. http://www.atri-online.org/index.php?option=com_content&view=article&id=248&Itemid=75, 2010.
2. American Transportation Research Institute (ATRI). Fpm congestion monitoring at 250 freight significant highway location: Final results of the 2010 performance assessment. http://www.atri-online.org/index.php?option=com_content&view=article&id=303:250-freight-significant-locations, 2010.
3. Austin Brown. Transportation Energy Futures: Addressing Key Gaps and Providing Tools for Decision Makers. Technical report, National Renewable Energy Laboratory, 2011.
4. J. Booth, P. Sistla, O. Wolfson, and I.F. Cruz. A data model for trip planning in multimodal transportation systems. In Proceedings of the 12th International Conference on Extending Database Technology: Advances in Database Technology, pages 994–1005. ACM, 2009.
5. G. Capps, O. Franzese, B. Knee, MB Lascurain, and P. Otaduy. Class-8 heavy truck duty cycle project final report. ORNL/TM-2008/122, 2008.
6. E. Chan and J. Zhang. Efficient evaluation of static and dynamic optimal route queries. Advances in Spatial and Temporal Databases, pages 386–391, 2009. Springer. LNCS 5644.
7. T.S. Chang. Best routes selection in international intermodal networks. Computers & operations research, 35(9):2877–2891, 2008. Elsevier.

8. T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. Introduction to Algorithms. MIT Press, 2001.
9. Davis, S.C. and Diegel, S.W. and Boundy, R.G. Transportation energy data book: Edition 28. Technical report, Oak Ridge National Laboratory, 2010.
10. J. Dobson and P. Fisher. Geoslavery. Technology and Society Magazine, IEEE, 22(1):47–52, 2003.
11. Federal Highway Administration. Highway Statistics. HM-63, HM-64, 2008.
12. D. Frigioni, M. Ioffreda, U. Nanni, and G. Pasqualone. Experimental analysis of dynamic algorithms for the single. Journal of Experimental Algorithmics (JEA), 3:5, 1998. ACM.
13. M. Faulkner, M. Olson, R. Chandy, J. Krause, K. Chandy, and A. Krause. The next big one: Detecting earthquakes and other rare events from community-based sensors. In Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on, pages 13–24. IEEE, 2011.
14. Daniele Frigioni, Alberto Marchetti-Spaccamela, and Umberto Nanni. Semidynamic algorithms for maintaining single-source shortest path trees. Algorithmica, 22(3):250–274, 1998.
15. Garmin. http://www.garmin.com/us/.
16. Betsy George and Shashi Shekhar. Road maps, digital. In Encyclopedia of GIS, pages 967–972. Springer, 2008.
17. Google Maps. http://maps.google.com.
18. Erik G. Hoel, Wee-Liang Heng, and Dale Honeycutt. High performance multimodal networks. In Advances in Spatial and Temporal Databases, pages 308–327, 2005. Springer. LNCS 3633.
19. G.R. Jagadeesh, T. Srikanthan, and K.H. Quek. Heuristic techniques for accelerating hierarchical routing on road networks. IEEE Transactions on Intelligent Transportation Systems, 3(4):301–309, 2002.
20. Ning Jing, Yun-Wu Huang, and Elke A. Rundensteiner. Hierarchical optimization of optimal path finding for transportation applications. In Proceedings of the fifth international conference on Information and knowledge management (CIKM), pages 261–268, 1996. ACM.
21. H. Kargupta, J. Gama, and W. Fan. The next generation of transportation systems, greenhouse emissions, and data mining. In Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining, pages 1209–1212. ACM, 2010.
22. H. Kargupta, V. Puttagunta, M. Klein, and K. Sarkar. On-board vehicle data stream monitoring using minefleet and fast resource constrained monitoring of correlation matrices. New Generation Computing, 25(1):5–32, 2006. Springer.
23. Jon Kleinberg and Eva Tardos. Algorithm Design. Pearson Education, 2009.
24. A. Krause, E. Horvitz, A. Kansal, and F. Zhao. Toward community sensing. In Proceedings of the 7th international conference on Information processing in sensor networks, pages 481–492. IEEE Computer Society, 2008.
25. J. Krumm. A survey of computational location privacy. Personal and Ubiquitous Computing, 13(6):391– 399, 2009.
26. Joel Lovell. Left-hand-turn elimination. http://goo.gl/3bkPb, December 9, 2007. New York Times.
27. Lynx GIS. http://www.lynxgis.com/.
28. M. Mabrouk, T. Bychowski, H. Niedzwiadek, Y. Bishr, JF Gaillet, N. Crisp, W. Wilbrink, M. Horhammer, G. Roy, and S. Margoulis. Opengis location services (openls): Core services. OGC Implementation Specification, 5:016, 2005. Citeseer.
29. J. Manyika et al. Big data: The next frontier for innovation, competition and productivity. McKinsey Global Institute, May, 2011.
30. MasterNaut. Green Solutions. http://www.masternaut.co.uk/carbon-calculator/.
31. [76] NAVTEQ. www.navteq.com.

32. New York Times. Justices Say GPS Tracker Violated Privacy Rights. http://www.nytimes.com/2012/01/24/us/police-use-of-gps-is-ruled-unconstitutional.html, 2011.

33. OpenStreetMap. http://www.openstreetmap.org/.

34. Michalis Potamias, Francesco Bonchi, Carlos Castillo, and Aristides Gionis. Fast shortest path distance estimation in large networks. In Proceedings of the 18th ACM conference on Information and knowledge management, CIKM '09, pages 867–876, 2009.

35. Pothole Info. Citizen pothole reporting via phone apps take off, but can street maintenance departments keep up? http://goo.gl/cGl3B, 2011.

36. SafeRoadMaps. Envisioning Safer Roads. http://saferoadmaps.org/.

37. Hanan Samet, Jagan Sankaranarayanan, and Houman Alborzi. Scalable network distance browsing in spatial databases. In Proceedings of the 2008 ACM SIGMOD international conference on Management of data, SIGMOD '08, pages 43–54, 2008.

38. P. Sanders and D. Schultes. Engineering fast route planning algorithms. In Proceedings of the 6th international conference on Experimental algorithms, pages 23–36. Springer-Verlag, 2007.

39. J. Sankaranarayanan and H. Samet. Query processing using distance oracles for spatial networks. Knowledge and Data Engineering, IEEE Transactions on, 22(8):1158–1175, 2010. IEEE.

40. J.H. Schiller and A. Voisard. Location-based services. Morgan Kaufmann, 2004.

41. S. Shekhar, A. Fetterer, and B. Goyal. Materialization trade-offs in hierarchical shortest algorithms. In Advances in Spatial Databases, pages 94–111. Springer, 1997.

42. S. Shekhar and H. Xiong. Encyclopedia of GIS. Springer Publishing Company, Incorporated, 2007.

43. Shashi Shekhar, Ashim Kohli, and Mark Coyle. Path computation algorithms for advanced traveller information system (atis). In Proceedings of the Ninth International Conference on Data Engineering, April 19-23, 1993, Vienna, Austria, pages 31–39. IEEE Computer Society, 1993.

44. Shashi Shekhar, Ranga Raju Vatsavai, Xiaobin Ma, and Jin Soung Yoo. Navigation systems: A spatial database perspective. In Location-Based Services, pages 41–82. Morgan Kaufmann, 2004.

45. Daniel Sperling and D. Gordon. Two billion cars. Oxford University Press, 2009.

46. TeleNav. http://www.telenav.com/.

47. TeloGIS. http://www.telogis.com/.

48. TomTom. TomTom GPS Navigation. http://www.tomtom.com/, 2011.

49. U.S. Energy Information Adminstration. Monthly Energy Review June 2011. http://www.eia.gov/totalenergy/data/monthly/.

50. Waze Mobile. http://www.waze.com/.

51. Wikipedia. Usage-based insurance — wikipedia, the free encyclopedia. http://en.wikipedia.org/w/index.php?title=Usage-based_insurance&oldid=464698702, 2011. [Online; accessed 15-December-2011].

52. Shrank, D.L., Lomax, T.J., The 2011 Urban Mobility Report, Texas Transportation Institute, September 2011.

53. Venkata M.V. Gunturi, Ernesto Nunes, KwangSoo Yang and Shashi Shekhar, A Critical-Time-Point Approach to All-start-time Lagrangian Shortest Paths: A Summary of Results. Proc. of 12th International Symposium on Spatial and Temporal Databases (SSTD '11), Minneapolis, MN. 2011.

54. Changqing Zhou, Dan Frankowski, Pamela J. Ludford, Shashi Shekhar, Loren G. Terveen: Discovering personal gazetteers: an interactive clustering approach. 12th ACM International Workshop on Geographic Information Systems, ACM-GIS 2004, November 12-13, 2004, Washington, DC, USA. GIS 2004:266-273.