

Benchmarking Spatial Big Data

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

Spatial computing is 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 spatial computing. 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. Google Earth is being used in classrooms to teach children about their neighborhoods and the world in a fun and interactive way. 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 using mobility services [52].

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 routing 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.

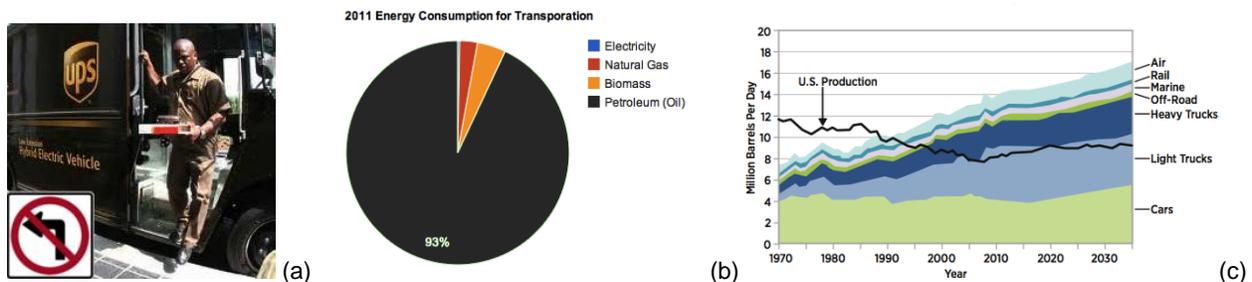


Figure 1: (Left) UPS avoids left-turns to save fuel [26]. (Middle) Petroleum is dominant energy source for US Transportation [49]. (Right) Gap between US petroleum consumption and production levels is large and growing [3,9]. (Best in color)

A domain-specific benchmark (such as a Spatial Big Data benchmark) should address four key criteria: relevance, portability, scalability and simplicity [55]. Related work in spatial database benchmarking [56,57] presents workloads for traditional geographic information systems (GIS) related spatial computing, such as raster and vector datasets. Raster data is used in remote sensing (e.g., Google Earth) whereas vector data represents points, lines and polygons, each with their own library of necessary operators. A key-missing component of these related benchmarks is graph-based datasets, useful for applications such as routing and urban navigation. In addition, the rise of spatio-temporal datasets also requires new workloads, as road networks now come with traffic speeds measured every minute of every day.

This paper makes the following contributions: an up-to-date workload for spatial computing, including four types of data (raster, vector, network and spatio-temporal); a set of summary metrics reminiscent of the historical TPS (transactions per second) metrics [55] for Spatial Big Data (SBD-R, -V, -N, -ST) and requirements, both functional (specific behavior) and non-functional (overall operation of a system), for future spatial computing benchmarks.

2. TRADITIONAL SPATIAL BIG DATA

The data inputs of spatial computing are more complex than the inputs of classical computing, as they include extended objects, such as: points, lines, and polygons in vector representation and field data in regular or irregular tessellation, such as raster data. The data inputs have two distinct types of attributes: non-spatial attributes and spatial attributes. Non-spatial attributes are used to characterize non-spatial features of objects such as name, population and unemployment rate for a city. They are the same as the attributes used in the data inputs of classical data mining. Spatial attributes are used to define the spatial location of extent of spatial objects [59]. The spatial attributes of a spatial object most often include information related to spatial locations, for example, longitude, latitude, and elevation, defined in a spatial reference frame, as well as a shape. There are four basic models to represent spatial data: raster (grid), vector (object), network (graph) and spatio-temporal:

Raster: In its simplest form, a raster consists of a matrix of cells (or pixels) organized into rows and columns (or a grid) where each cell contains a value representing information, such as temperature. A set of operations called Map Algebra was introduced [60] to manipulate representations of continuous variables defined over a common domain. These operations were categorized into three categories: local, focal and zonal; each based on the geographic size of the operation. For example, an elevation raster dataset can be queried with a zonal (large region) operation to derive slope. Raster datasets can be digital aerial photographs, imagery from satellites, digital pictures, or even scanned maps.

Vector: Geographic features in a vector format can be represented by points, lines, or polygons (areas). Vector data over a space is a framework to formalize specific relationships among a set of objects. In Table 1, a relationship between spatial and non-spatial data is described using spatial relations performed on vector datasets. These relations are separated into a number of classifications: topological to describe relationships regardless of projection, directional to describe

Non-spatial Relation	Spatial Relation
Arithmetic	Set-oriented: union, intersection, membership, ...
Ordering	Topological: meet, within, overlap, ...
Instance-of	Directional: North, Left, Above, ...
Subclass-of	Metric: distance, area, perimeter, ...
Part-of	Dynamic: update, create, destroy, ...
Membership-of	Shape-based and visibility

Table 1: Common relationships among spatial and non-spatial data

orientation and metric to describe distances between objects.

Networks: Traditional spatial computing utilizes 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 (10^8) edges for the roads in the U.S.A. [31].

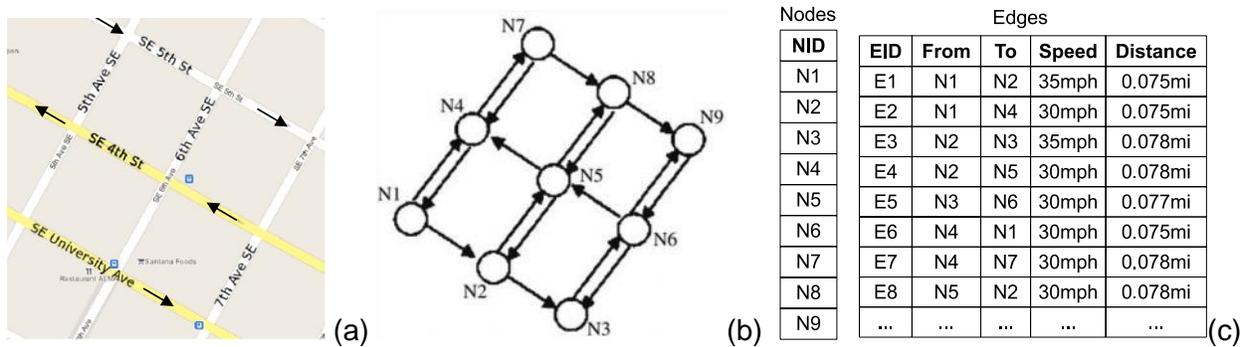


Figure 2: Current representations of road maps as directed graphs with scalar travel time values. (Left) Example road map. [17] (Middle) Graph Representation. (Right) Tabular representation of digital road maps.

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 Dijkstra'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.

OpenLS [58] presents a system (see Figure 3) that incorporates a wide-spectrum of spatial technologies, ultimately reporting to a location-aware client. The location utility performs as a geocoder by determining a geographic

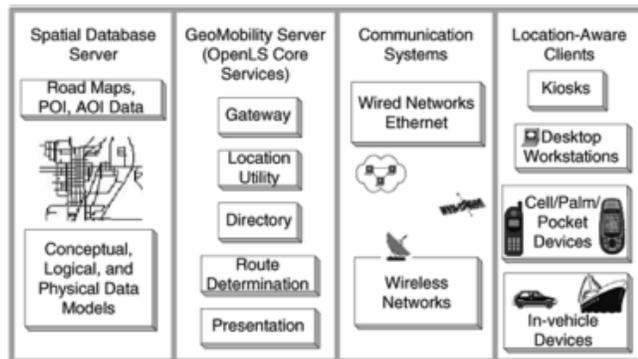


Figure 3: OpenLS Architecture [58]

position, given a place name, street address or postal code. The directory service provides users with access to the nearest, or a specific place, product or service. The presentation layer renders geographic information for display. The route determination component provides routing information between locations.

3. EMERGING SPATIAL BIG DATA

Spatio-temporal datasets 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 NO_x), 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.

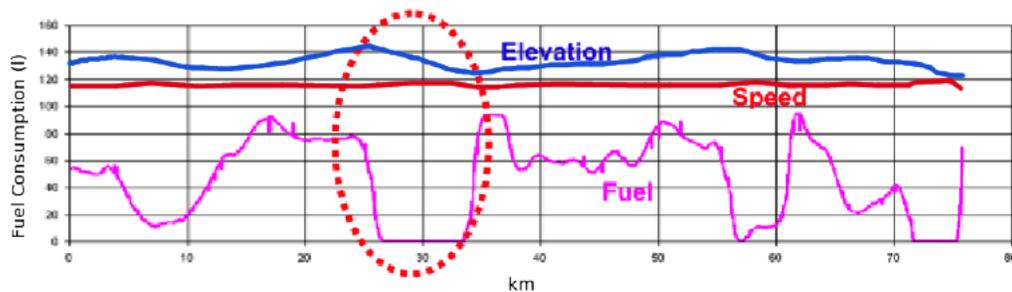


Figure 4: Engine measurement data improves understanding of fuel consumption [5]. (Best in color)

For example, Figure 4 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 5 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.

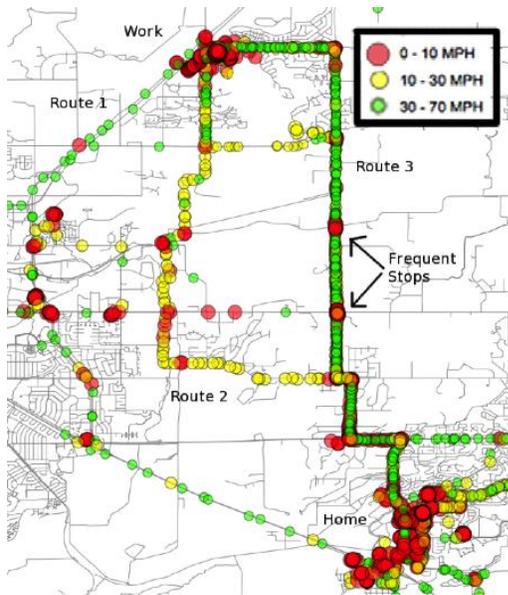


Figure 5: A commuter's GPS tracks over three months reveals preferred routes.

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 3(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 6) 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 6(a), 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)

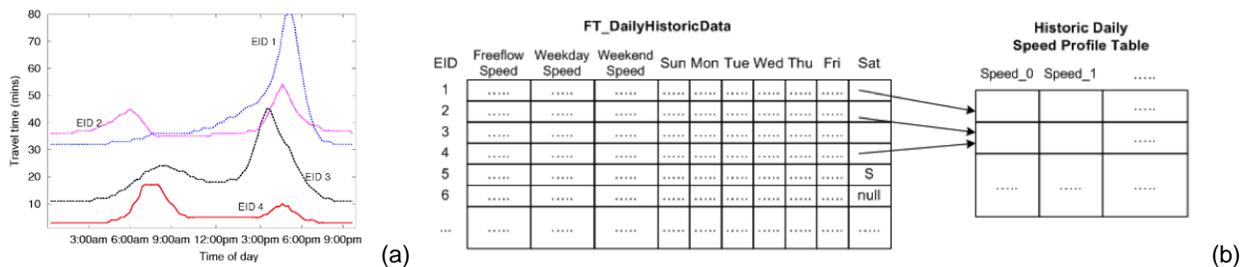


Figure 6: Spatial Big Data on Historical Speed Profiles. (Left) Travel time along four road segments over a day. (Right) Schema for daily historic speed data. (Best in color)

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].

5. METRICS FOR SPATIAL BIG DATA BENCHMARKS

Metrics for spatial big data can be categorized via a classification used in software engineering into functional (specific behaviors) and non-functional requirements (overall operation of a system). In this section, we will describe each and provide examples.

Metrics for Functional Spatial Big Data Requirements: Spatial computing traditionally operates on one of the four data types listed in Table 2:

Table 2: Data Types in Spatial Computing

Data Type	Representation	Operations	Potential Metric
Raster (field)	Geo-Matrix	Map algebra operations on Local, Focal, Zonal regions	Map algebra operations per second
Vector (object)	Points, Lines, Polygons	Intersection Model, Nearest Neighbor, Point Query, Range Query, etc.	Nearest Neighbors per second, Range-query (screen paint) per second
Network	Graphs (nodes, edges)	Shortest Path, Max Flow, Evacuation, etc.	Shortest-Paths per second
Spatio-Temporal	Trajectories, Temporal Networks	Time-dependent shortest path, GPS tracking, logging, etc.	Mobile device interactions per second

SBD-R: Raster datasets are frequently used for remote sensing applications, where large-scale map algebra and matrix operations are used. A helpful performance metric (e.g., map algebra operations per second) would measure how quickly representative operations of this type could be performed on a variety of dataset scales (e.g. terabyte, petabyte, exabyte, etc.).

SBD-V: Processing vector datasets in spatial database systems has historically been computationally expensive, with many key features (e.g., nearest neighbor queries) not being provided with the system. As newer systems are developed with these features, performance metrics measuring how quickly range queries and nearest neighbor queries can be computed are needed. Representative metrics include: nearest neighbors per second and range-queries per second.

SBD-N: Mapping services such as Google Maps has demonstrated the popularity of network-based datasets for use-cases such as personal transportation routing. It is not hard to imagine Google has a measure for how many shortest-paths per second it can calculate as it is serving the world driving directions, but universal and public benchmarks in this field will allow comparison between current spatial database systems. Representative metrics include: shortest-paths per second and evacuation planning.

SBD-ST: Spatio-temporal datasets are becoming more and more commonplace with the rise of location-based services and metrics for database systems rating their ability to handle some of these more common complex queries are crucial. For example, many applications currently request the location of a user, and potentially also monitor nearby points of interest to

report back to the user. So a metric that described the number of mobile device interactions (e.g., tracking, local context, location trigger, etc.) per second would be extremely useful for a variety of end-user applications. Representative metrics include: mobile device interactions per second, GPS logs per second, etc.

Metrics for Non-Functional Spatial Big Data Requirements: Many Spatial Big Data use-cases (e.g., emergency services like E911, disaster response, etc.) typically require fault tolerance, where it should be resilient against natural calamities such as earthquakes, hurricanes, etc. Such requirements necessitate “triple-continental redundancy” [62], where the data is replicated on servers spread across multiple continents. This requirement poses several challenges for current cloud-based storage technologies due to performance issues inherent with wide-area replication and access. A potential metric for disaster resilience is a resilience footprint (e.g., 100 mile resilient, 1,000 mile, 10,000 mile), which may indicate the geo-spatial footprint of the disaster (e.g., fire, flood, tornado, earthquake, hurricane) that will not disrupt service.

Privacy of geographic information is an important and timely challenge due to personal information in GPS tracks, Check-in's, tweets, etc. 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]. For example, Ushahidi is a non-profit tech company providing technology for citizen-based reporting used in many countries with controlling regimes where privacy and protecting the reporter is extremely important [61]. Computer science efforts at obfuscating location information to date have largely yielded negative results. Thus, many individuals hesitate to indulge in mobile commerce and citizen reporting due to concern about privacy of their locations, trajectories and other spatio-temporal personal information [25]. It may be premature to provide specific metrics. However, Spatial Big Data benchmarks and metrics are 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? Crucial to widespread adoption will be comforting the public, where a easy-to-understand metric describing the loss of privacy given information surrendered (e.g., adversary information gain per piece data submitted) will help people understand and compare various services against their privacy concerns.

6. CONCLUSION

Increasingly, location-aware 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., transportation 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.

Current benchmarks for spatial computing remain limited to small data sizes and only a portion of current popular data types. New benchmarks need to be built around Spatial Big Datasets, incorporating all four data types (raster, vector, network, spatio-temporal), while covering a wide variety of use-cases from emergency management, location-based services, advanced routing services, etc. New performance metrics, both functional (e.g., mobile interactions per second) and non-functional (e.g., disaster resilience footprint), will facilitate comparison between new systems being created and promoted by various spatial computing vendors.

7. 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. This work was supported by the NSF (0713214, DGE-0504195) and the US-DoD.

8. 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 Algorithms (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. 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/cGI3B>, 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 Administration. 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.
55. Jim Gray. The Benchmark Handbook for Database and Transaction Systems (2nd Edition). Morgan Kaufmann. 1993.
56. Suprio Ray, Bogdan Simion, Angela Demke Brown: Jackpine: A benchmark to evaluate spatial database performance. ICDE 2011:1139–1150.
57. M. Stonebreaker, J. Frew, K. Gardels, and J. Meredith. “The SEQUOIA 2000 Storage Benchmark,” in SIGMOD '93: Proc. Of the 1993 ACM SIGMOD Intl conference on Management of Data. May 1993. pp. 2–11.
58. M. Mabrouk, T. Bychowski, H. Niedzwiadek, Y. Bishr, JF Gaillet, N. Crisp, W. Wilbrink, M. Horham-mer, G. Roy, and S. Margoulis. Opengis location services (openls): Core services. OGC Implementation Specification, 5:016, 2005. Citeseer.
59. Shashi Shekhar, Michael R. Evans, James M. Kang, Pradeep Mohan: Identifying patterns in spatial information: A survey of methods. Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery (WIDM) 1(3):193–214 (2011)
60. C. Dana Tomlin. Geographic Information Systems and Cartographic Modeling. Prentice Hall. 1990.
61. Ushahidi. <http://www.ushahidi.com>.
62. Eric Bauer, Randee Adams, Daniel Eustace. Beyond Redundancy: How Geographic Redundancy Can Improve Service Availability and Reliability of Computer-Based Systems. Wiley-IEEE Press. 2011.