

A Vision for Micro and Macro Location Aware Services

Abdeltawab M. Hendawi¹ Mohamed Khalefa² Harry Liu³ Mohamed Ali³ John A. Stankovic¹

¹University of Virginia, VA, USA
{hendawi, stankovic}@virginia.edu

²University of Alexandria, Egypt
²khalefa@alexu.edu.eg

³University of Washington, Tacoma WA, USA
³{hongyliu, mhali}@uw.edu

ABSTRACT

A few decades ago, the Internet was created. Since then, searching for information and services has increased exponentially. With the introduction of GPS-enabled devices, a special type of search appeared offering location-aware services. These services customize search results based on users' location. This includes, but not limited to, (1) service finding, e.g., "find the nearest pizza restaurant", (2) routing, e.g., "obtain the shortest path from a user's home to the airport", (3) transportation, e.g., "what are the bus links to get a user from downtown to the mall", and (4) monitoring, e.g., "alert a parent if their child school-bus deviates from its regular route". Though new hardware and software technologies such as smart watches, voice search, and big-data platforms have been introduced and widely used, each single type of the above services has benefited very little from these technologies. On the local level of each service (*the micro level*), a full-fledged view is still missing. On the global level of all service types (*the macro level*), all Location-aware services are still acting as isolated islands and a global optimized service is not available. This paper presents our vision of how to provide an integrated macro location-aware service that acts harmoniously, and how each micro service can be further improved by better incorporation of novel technologies. We also overview the key challenges associated with these suggested improvements. Then, we highlight the potential value-added by the application of our vision.

1. VISION

Location aware services encompass applications that are sensitive to user's locations. These services span many domains that include service finding [10, 13, 26], routing [7, 16], transportation [3, 4, 19], and monitoring based [5, 20] systems.

In this paper, we present a vision for realizing location aware services that span the local micro scale as well as a global macro scale. We provide an architecture that integrates common services to reflect our vision. We argue that the benefits of developing such an integrated system of services offers better user satisfaction by providing a well-personalized service that takes users' profiles (e.g., health status), behavior (e.g., daily movement patterns), and other related factors (e.g., environmental conditions) into account, and a wiser resource management by applying a global optimization that

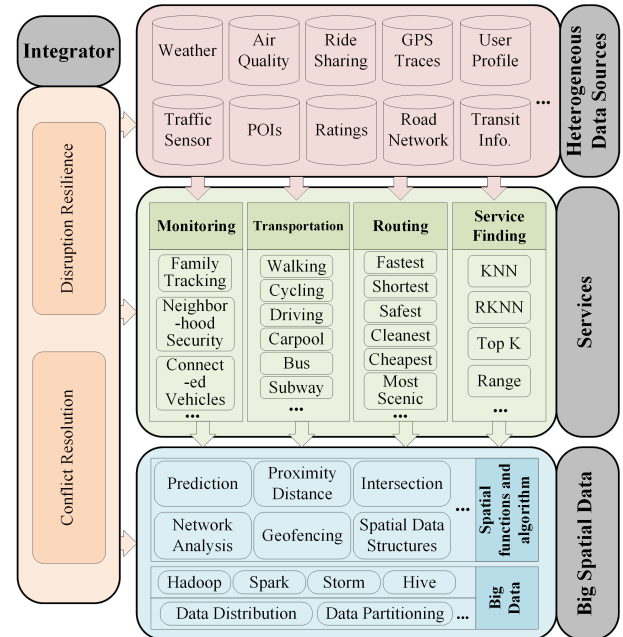


Figure 1: Architecture For Future Integrated Location Aware Services.

assimilates the local optimization at the micro level and resolves possible conflicts, and is resilient to potential disruptions. Figure 1 shows the architecture.

The rest of this paper is organized as follows. Section 2 describes the opportunities and corresponding challenges to improve the conventional location-aware services on the micro scale, i.e., locally within each service. Section 3 provides the opportunities and challenges for the integrated macro scale service that cross-references all service types. Section 4 briefly highlights the added value and potential benefits from applying this vision. Finally, Section 5 concludes the discussion.

2. OPPORTUNITIES ON MICRO LEVEL

In this section, we envision the opportunities to improve each of the service finding, routing, transportation, and monitoring categories and give a glimpse on the core challenges. This section represents the local (micro) view by looking at each service type in isolation. On the other hand, Section 3 gives the global (macro) view by looking at the big picture resulting from integrating the pieces of the puzzle.

2.1 Service Finding

In conventional service finding systems, mobile users search for stationary or mobile entities that have spatial attributes. The users' current location and/or the entities' locations are considered to as-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

sess the relevance of the search results. In our vision, we believe that the users' future locations are also important to consider along with the current location to improve the quality of the provided service. ABIresearch [22] expects that the smart watches market will reach 50 million watches per year. Smart watches are GPS-enabled devices and provide voice search facilities. Users tend to seek services while driving, e.g., "find the next nearest gas station". To answer this inquiry, unfortunately, current services do not look at the future. Rather, they respond by getting the nearest gas station to the user's location at query time, even if the mobile user has already passed it. In addition to augmenting users' predicted locations, we argue that predicting users' behavior would further improve the relevance of the recommended services. However, these improvements come with challenges. In the following, we briefly highlight a couple of these challenges.

2.1.1 Predicting Future Locations

Users' movements involve complex patterns that make it challenging to predict their paths. Therefore, prediction models that assume movements to be in straight lines are not practical. For predictive functions relying on deeper analysis using historical traces can achieve higher quality prediction. However, they severely suffer from efficiency and scalability issues. They can work fine for one user at a time, but in reality there are millions of concurrent users. In addition, the accuracy of these prediction functions drops significantly after the next couple of road segments. Hence, they can not be employed for longer term prediction that looks at tens of minutes into the future [10, 11]. In summary, it is still an open challenge to invent an accurate, efficient, and scalable prediction model that is able to dynamically anticipate users future behavior as well as to enable accurate short and long-term predictions.

2.1.2 Uncertainty Handling

Whether we are introducing services for present or future needs, dealing with uncertainty is important and crucial. Location and time uncertainty are intuitively expected as a result of inaccurate localization from GPS devices, communication delays, errors in speed calculation, to name a few. Consequently, it is critical not to ignore the imprecise nature of data while processing users' requests. Not too much work has been introduced to handle the uncertainty issue while processing service finding requests. For instance, the concept of u -bisector in [26] is utilized to process variations of KNN queries, named, possible nearest neighbor query (PNNQ) and trajectory possible nearest neighbor query (TPNNQ). This work tries to find services nearest to user's imprecise position which is represented as regions and not precise points. This work is limited to these two query types and a target service at the present time. Therefore, we believe that supporting a wider variety of service finding queries for current and future times with consideration of uncertainty remains an ongoing challenge.

2.2 Routing

Routing is a central service for modern life across the world [2]. Previously, commuters used to carry paper maps and check them as needed on their trips. Nowadays, GPS-enabled devices, and on-line and mobile digital maps are widely available and much easier to use. Traditionally, existing routing providers such as Google maps, Bing maps, TomTom, Garmin and many others focus on providing the shortest (distance or travel time) path between two given locations. As these services become more sophisticated, users look for preference based routing that goes beyond distance and travel time. Though there are research trials to provide preference routing [7, 16, 24] that could be considered a step in the right direction, however, they consider GPS trajectories as a sole source, consequently limited personalization is offered. This means, there are

still many opportunities for improvement. In fact, preference needs should not be defined narrowly, but should essentially be expansive to encompass safety, scenic routes, services and attractions, air-quality, social proximity and more. Today, we live in an era where users generate large volumes of geo-tagged data. Analysis of such data gives indicators of the traffic status, speed profiles, and drivers behaviors. All of this should be augmented in future routing engines such that it gives smarter path recommendations. Not only the data concerning the current conditions on roads, but also data about weather conditions, air quality, local events, and user profiles including health issues should be considered. It is also important to use these data sources in the present time as well as make good estimates for their future values [15]. For example, using traffic only is misleading because it could be clear traffic on the user's whole route at the trip start time, but when the user moves forward it becomes congested on some of the segments. Therefore, traffic forecasting would help avoid the areas of predicted congestion. In the next few paragraphs, we highlight the associated challenges for these suggested improvements.

2.2.1 Forecasting Traffic and Environmental Conditions

Predicting traffic is a challenging as we have to combine (1) the prediction based on analysis of historical statistics about the traffic in addition to (2) the present snapshot of vehicles on the roads. The later is harder to do as it requires continuous monitoring of vehicles movement and hence updating the traffic prediction frequently. It is challenging to combine the results from these two types of prediction and assign the appropriate weights to each of them. Things become further complicated when we consider forecasting weather and environmental changes and effects.

2.2.2 Enriched Data Structures And Algorithms

It is not clear that the appropriate data structure that can accommodate this wide spectrum of different routing related data exists. Even the recently proposed multi-preference and time-dependent graphs such as ATAG (attributed time aggregated graph) [12] do not provide a way to encompass the environmental events. There is also a lack of path computation algorithms that can navigate the road network graphs, (or its equivalent new data structure), with an eye on the aforementioned data parameters.

2.3 Transportation

Transportation systems move people from place to place. However, the system is far from optimal, and congestion is a normal reality for cities [17]. Multi-mode transport systems exist in cities, Figure 2, and people tend to avoid congestion by various means. Besides, individuals also have preferences for transportation modes when there are more than one available. Current services recommend different modes of transportation individually. For example, Google Maps provides trip advice on driving, municipal transportation, walking, cycling and flights separately. Though the transportation option can make suggestions combined with buses, ferries, subways, and walking, the choice of modes is still limited. With the popularization of GPS-enabled devices, there are tremendous opportunities to optimize travel-mode plans with all transportation modes integrated. For example, a user in a suburban area can get carpool or taxi service information from his home location, park at a transit terminal, then take a bus to downtown. Because of the heavy congestion in the downtown area during peak hours, then he can walk several blocks avoiding the most terrible traffic, take a bus again, and finally walk through a beautiful park with fresh air to the destination with a good mood. In this multi-mode solution, more personal preferences can be satisfied than before, in addition

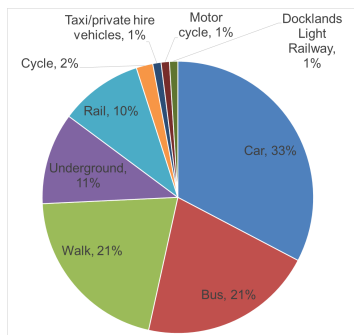


Figure 2: Daily Transport Modality in London, 2012 [9].

to less waiting and travel time, and lower prices. As another example, some users may plan to do some exercise on the way home or travelers may want to choose travel modes that are readily available to attractions from location to location. They can input options such as half an hour walk segment along the most walkable/scenic roads.

2.3.1 Mathematical Modeling

Integrating a variety of transportation modes and location information for users and vehicles make it challenging to provide a multi-mode solution. Mathematical methodologies need to be explored for multi-mode modeling. For example, how to use graph and network theory to model and integrate the heterogeneous transportation networks such as road networks, bus/subway/taxi networks, bike lanes and so on remains unknown. Also, to ensure the reliability of the solution, statistical models are required to validate the level of service. Besides, individuals may have multiple competing preferences, and several people may need the same services, for example, carpooling. How to make tradeoffs that bring the largest utility is a critical question to answer. Game theory can be adopted to make sound tradeoffs.

2.3.2 Incorporating Personal Preferences

When selecting a means of transportation, many users are concerned about factors such as environmental impact, health or having fun in addition to price and travel time/distance. How to measure these various factors and evaluate the utility for each transportation mode are still unclear. We recommend designing a framework for the measurements and evaluations at a high level to guide the implementation and validation for transportation systems that encompass individuals' personalized needs.

2.4 Monitoring

Monitoring is of great importance for several domains including: inventory control, public safety by enforcing guarantee rules, fleet management, ride sharing, smart traffic control, and advertising. These applications can be abstracted as continuous spatio-temporal queries over massive streaming geo-tagged data from numerous devices. While current advertising applications send coupons to all individuals within a certain distance from a store, future advertising applications may send the promotion coupons based on their attributes (gender, age, user categorization, ...) to increase the expected user interaction. Driven by massive geo-tagged data, and linking data across repositories introduces a new breed of complex event processing algorithms for geo monitoring applications. These applications need to process massive amounts of heterogeneous streaming and non-streaming geo data in real-time, while supporting complex events. For example, a parking advisor would not only need to detect the available parking spaces and inform user *A* in real time for her closest parking space, but in addition, it needs to inspect the behavior of nearby vehicles to anticipate if the this parking space would still be free when user *A* reaches this location. Moreover, the application needs to identify areas where parking is

permitted from various geo-data sources, as well as parking cost, and other issues.

To overcome these obstacles and build a scalable monitoring framework which supports thousand of concurrent continuous queries, we highlight two areas of needed improvements:

2.4.1 Efficient handling for frequently updated data

While several research efforts [14, 21] have occurred, their main idea is to cache updates before inserting them into an R-tree data structure. A recent research effort [20] features an efficient library for continuous spatial queries. These techniques are based on R-trees, while the skip quadtree [8] supports frequent updates on quadtrees. Extensive experimental evaluation for these algorithms, and more efficient techniques exploiting in-memory processing are needed.

2.4.2 Hybrid Indexing

While significant research efforts for integrating textual and spatial data, including [27, 23, 6, 18, ?, 1, 25, 5] have occurred, there is a lack of efficient hybrid spatial-temporal indexes, with non-spatial data such numerical data. The index must tolerate having different selectivity estimates over the indexed data. For instance, for a certain range of values the spatial feature may give a higher filtering than textual, and vice versa. In addition, as the non spatial data may be static, e.g., the number of lanes in a certain road segment, or dynamic such as the number of cars on a road segment, updating the dynamic values should be triggered in a lazy manner, and to an appropriate level of approximation. For example, in a congestion monitoring application, the congestion may only develop after certain number of vehicles are present.

3. OPPORTUNITIES ON MACRO LEVEL

In this section, we present our vision for macro level services. While the micro level optimizes query results within the boundary of the modules as discussed in the previous section, in this section, we present our vision for globally optimized results across different micro level modules. To achieve the stated goal, we use a conflict resolution module, where some modules invalidate results of other modules, and support disruption resilience.

3.1 Conflict Resolution

Conflicts would be introduced in our system due to having independent micro level services. Consider the following situation, a building is evacuating due to a fire, a service finding module reports results that are nearby to the on fire building. These search results in the affected area should be removed and not reported to the user; this is a *spatial* conflict. Other spatial conflicts can result from accidents. Please note the first conflict presents a contradiction situation, while in the second scenario, the conflict results in degrading the quality of a service by increasing the delay of the route. Another example for conflict may be due to health issues, as the shortest path returned from the routing micro service may not be suitable to the health conditions of user suffering from asthma. Another possible source of conflict is sparsity of input data, e.g., in a large metropolitan road network some minor road may not have as many sensors as highways. Conflicts are detected in real time, then the action is resolved based on the conflict type: contradiction, or service level degradation.

3.2 Disruption Resilience

Disturbances occur for various planned and unplanned reasons. Planned reasons include road repair and sport events, while unplanned disturbances may be due to traffic incidents, or weather disasters, like tornadoes and floods. It is challenging to predict the behavior of unplanned events and its effect on the road network.

Consider a tornado moving at 30 mph, it is very challenging, yet of high importance, to answer queries such as "find commuters that might be hit by this tornado in the next 30 minutes". The difficulty of this query is due to the highly dynamic environments, as commuters change their locations and velocities and the tornado modifies its shape, size, speed, and direction. Taking all of these factors into consideration makes prediction complex, and challenging.

4. ADDED VALUE

Applying the suggested improvements for both micro and macro levels in and across all types of location aware services is expected to have positive consequences. This includes (1) improving user experience by giving a complete scenario of recommendations that best fit users' profiles within the context of the needed service, and (2) advancing resource utilization (e.g., road network, means of transportation, city air, computational resources, etc). By having a global view over all services we will be able to provide a global optimization of using the resources. This is because depending only on local (micro) optimization does not guarantee the macro optimization across all services. For example, consider the following scenario. A user searches for the nearest restaurant. A traditional service returns the closest restaurant to the user's current location after computing the distance to all restaurants in the vicinity. Some existing service might incorporate other simple factors such as users' ratings. In a future integrated service, Figure 1, the process and the return are completely different. First, the integrated services would check the user's different profiles, (e.g., health, rating, social, etc) and figure out that the user suffers from high blood pressure and the readings of the user's wearable devices indicate blood pressure is really high. Consequently, fast and fatty food restaurants are excluded. Then, the service would examine the environmental conditions to find that air quality is poor is some near by areas. As the user's profile indicates an asthma problem, the service will focus on restaurants with clean air. At the time of the users' query, it is not raining, but the service might predict rain in few minutes, so the service would also recommend taking a taxi and not to bike as the user normally does. In the end, an appropriate restaurant will be returned along with a path to take and a method to ride.

5. CONCLUSION

The take home message from this vision is that we strongly believe that the traditional location-aware services need to evolve in both the *micro* and *macro* levels. In the *micro* level, this can be done by anticipating future conditions including users needs and all environmental conditions. At the *macro* level, future services need to collaborate and work in a more harmonious way where each type of service is sensitive to changes in the other services. Therefore, when a conflict or a disruption occurs, a smart integration mechanism is required to resolve and ease the consequences. By applying this vision, it is expected that the entire quality of location-aware service is improved, and the utilization of resources is more efficient. This will then be reflected in the satisfaction, productivity, and welfare of society.

6. REFERENCES

- [1] Fuzzy keyword search on spatial data. *Lecture Notes in Computer Science*, 5982 LNCS(PART 2):464–467, 2010.
- [2] AllWebsiteStats. Statisits For Websites Usage. <http://allwebsitesstats.com/>, June 2016.
- [3] H. Bast and S. Sabine. Frequency-based search for public transit. In *ACM SIGSPATIAL GIS*, pages 13–22, exas, USA, Nov. 2014.
- [4] J. Benjamin, P. C. Swinkels, G. J. Teeuwen, B. van Antwerpen de Fluiter, and H. A. Fleuren. Operational planning of a large-scale multi-modal transportation system. *European Journal of Operational Research*, 156(1):41–53, 2004.
- [5] L. Chen, G. Cong, C. S. Jensen, and D. Wu. Spatial keyword query processing: an experimental evaluation. In *Proceedings of the VLDB Endowment*, volume 6, pages 217–228. VLDB Endowment, 2013.
- [6] G. Cong, C. S. Jensen, and D. Wu. Efficient retrieval of the top-k most relevant spatial web objects. *Proceedings of the VLDB Endowment*, 2(1):337–348, 2009.
- [7] D. Dellling, A. V. Goldberg, M. Goldszmidt, J. Krumm, K. Talwar, , and R. F. Werneck. Navigation made personal: Inferring driving preferences from gps traces. In *ACM SIGSPATIAL GIS*, Washington, USA, Nov. 2012.
- [8] D. Eppstein, M. T. Goodrich, and J. Z. Sun. The skip quadtree: a simple dynamic data structure for multidimensional data. In *Computational geometry*, pages 296–305. ACM, 2005.
- [9] T. for London. Travel in London Report 6. URL:<http://content.tfl.gov.uk/travel-in-london-report-6.pdf>, 2013.
- [10] A. M. Hendawi, M. Ali, and M. F. Mokbel. A Framework for Spatial Predictive Query Processing and Visualization. In *MDM*, Pennsylvania, USA, June 2015.
- [11] A. M. Hendawi, J. Bao, M. F. Mokbel, and M. Ali. Predictive Tree: An Efficient Index for Predictive Queries On Road Networks. In *ICDE*, Seoul, South Korea, Apr. 2015.
- [12] A. M. Hendawi, A. Rustum, A. A. Ahmadain, D. Oliver, D. H. A. Teredesai, and M. Ali. Dynamic and Personalized Routing in PreGo. In *MDM*, Porto, Portugal, June 2016.
- [13] H. Hu, J. Xu, Q. Chen, and Z. Yang. Authenticating location-based services without compromising location privacy. In *SIGMOD*, pages 301–312, Arizona, USA, May 2012.
- [14] M. L. Lee, W. Hsu, C. S. Jensen, B. Cui, and K. L. Teo. Supporting frequent updates in r-trees: a bottom-up approach. In *VLDB*, pages 608–619. VLDB Endowment, 2003.
- [15] Y. Li, S. George, C. Apfelbeck, A. M. Hendawi, D. Hazel, A. Teredesai, and M. Ali. Routing Service With Real World Severe Weather. In *ACM SIGSPATIAL GIS*, Texas, USA, Nov. 2014.
- [16] W. Ling-Yin, Y. Zheng, and W.-C. Peng. Constructing popular routes from uncertain trajectories. In *KDD*, pages 195–203, Beijin, China, Aug. 2012.
- [17] G. Nikolas and C. F. Daganzo. Macroscopic modeling of traffic in cities. In *TRB 86th annual meeting*, pages 0407–0413, Washington DC., Jan. 2007.
- [18] J. B. Rocha, O. Gkorgkas, S. Jonassen, and K. Nørvg. Efficient processing of top-k spatial keyword queries. *Lecture Notes in Computer Science*, 6849 LNCS:205–222, 2011.
- [19] M. Sabyasachee, T. F. Welch, and M. K. Jha. Performance indicators for public transit connectivity in multi-modal transportation networks. *Transportation Research*, 46(7):066–1085, 2012.
- [20] Y. Shi, A. M. Hendawi, H. Fattah, , and M. Ali. Rxspatial: a framework for real-time spatio-temporal operations. In *SIGMOD*, pages 366–367. ACM, 2016.
- [21] Y. N. Silva, X. Xiong, and W. G. Aref. The rum-tree: supporting frequent updates in r-trees using memos. *The VLDB Journal*, 18(3):719–738, 2009.
- [22] N. Spencer. The Apple Watch and Smart Watch Forecast for 2015. URL:<https://www.abiresearch.com/market-research>, Mar. 2015.
- [23] S. Vaid, C. B. Jones, H. Joho, and M. Sanderson. Spatio-textual indexing for geographical search on the web. In *SSTD, SSTD'05*, pages 218–235, Berlin, Heidelberg, 2005. Springer-Verlag.
- [24] C. Vaida and C. S. Jensen. Vehicle Routing with User-Generated Trajectory Data. In *MDM*, pages 14–23, Pennsylvania, USA, June 2015.
- [25] D. Wu, M. L. Yiu, C. S. Jensen, and G. Cong. Efficient continuously moving top-k spatial keyword query processing. In *ICDE*, pages 541–552. IEEE, 2011.
- [26] X. Xie, R. Cheng, and M. L. Yiu. Evaluating Trajectory Queries over Imprecise Location Data. In *SSDBM*, pages 56–74, Chania Crete, Greece, June 2012.
- [27] Y. Zhou, X. Xie, C. Wang, Y. Gong, and W.-Y. Ma. Hybrid index structures for location-based web search. In *CIKM*, pages 155–162. ACM, 2005.