

When the Internet Hits the Road

Christian S. Jensen

Department of Computer Science, Aalborg University, Denmark
<http://www.cs.aau.dk/~csj>

Abstract: The Internet has recovered from the dot-com crash of the early 2000's and now features an abundance of new, innovative technologies and services. We are also witnessing the emergence of a communication and computing infrastructure that encompasses millions of people with mobile devices, such as mobile phones, with Internet connectivity. This infrastructure will soon enable the Internet to go mobile.

This paper describes the background and aspirations of a new research project that is concerned with data management aspects of innovative mobile Internet services. It is argued that mobile services will be context aware, and the project devotes particular attention to geographical context awareness. The project will adopt a prototyping approach where services are built and exposed to users, and where data management challenges are identified and addressed. The paper describes the evolving service platform that supports the approach chosen, it describes some of the data management techniques being integrated into the service platform, and it describes research guidelines that the project aims to follow.

1 Introduction

The Internet is teeming with new and innovative technologies and services. Many of these are fueled by Google-like business models where services are made available on a global scale and are free to their users, being paid for instead by other means such as advertisement. Many services build on community concepts. Blogs and RSS feeds are everywhere.

Services that cater to our needs for factual information include Wikipedia. For images, there are services such as flickr.com, photo.net, and plazes.com. For music, there is last.fm, pandora.com, and The Hype Machine. For video, youtube.com is probably the most visible.

To give an impression of the richness of the services currently available, here are some additional video-related services: Angry Alien, AnimeEpisodes.Net, Badjojo, Blastro, Blennus, Blip.tv, Bofunk, Bolt, Break.com, Castpost, CollegeHumor, Current TV, Dachix, Dailymotion, Danerd, DailySixer.com, DevilDucky, Double Agent, eVideoShare, ETVTV1, FindVideos, Free Video Blog, Google Video, Grinvi, Hiphopdeal, iFilm, Keiichi Anime Forever, Kontraband, Lulu TV, Metacafe, Midis.biz, Music.com, MusicVideoCodes.info, MySpace, MySpace Video Code, Newgrounds, NothingToxic, PcPlanets, Pixparty, PlsThx, Putfile, Revver, Sharkle, SmithHappens, StreetFire, That Video Site, Totally Crap, VideoCodes4U, VideoCodesWorld, VideoCodeZone, vidiLife, Vimeo, vSocial, Yikers, and Zip-

pyVideos (visit these at your own peril; I have only visited a few of them).

Second Life at secondlife.com features a virtual world that is being built and owned by its residents, of which there are currently more than 2 million. While the world is virtual, it is also real: it costs real money to “live” in this virtual world; it has its own currency that can be exchanged for real money; and one can make and spend money in this virtual world. Of course, there has already been a lawsuit in the real world about a property deal in the virtual world.

The above is just a small snapshot of what is currently available. While many other services could be singled out, these services collectively represent a good view of how a range of technologies are being used for the creation of innovative and, seemingly, commercially viable services. In parallel with these developments on the conventional Internet, an infrastructure is emerging that will enable the mobile Internet.

In particular, driven in large part by swift and sustained advances in computing and communication technologies, an infrastructure is emerging that contains vast quantities of computing and sensing devices that are Internet-worked by means of wireless communication technologies.

Notably, we are witnessing continued improvements in the capabilities of consumer electronics such as mobile phones, personal digital assistants, laptop computers, cameras, mp3 players, watches, navigation systems, and driver assistance systems. Existing types of systems and devices combine and new types with new functionalities emerge. The performance and performance/price ratios associated with key technologies utilized by such systems and devices continue to improve, which promises an increased proliferation.

Mobile phones are of particular interest. Most new phones are Internet-enabled, i.e., they have built-in data communication capabilities. The bandwidth is currently limited to well below 100 kbit/s for GPRS and EDGE [C06] networks and below 400 kbit/s for 3G [N03] networks. With the introduction of HSDPA [N06], with up to 14.4 Mbit/s, and Wimax [W06], with speeds of up to 10 Mbit/s, this is slated to change.

The current pricing for mobile Internet access, often in excess of 1€/MB, remains quite high. Flat-rate subscriptions are possible with some mobile providers, but again at a relatively high price. This is expected to change in the future—flat-rate subscriptions may be priced near today’s wireline broadband subscriptions. However, substantially cheaper mobile Internet access may still be some years away.

Geo-positioning is also becoming increasingly available. For example, network assisted GPS promises to eliminate the excessive power consumption of GPS receivers, thus rendering GPS practical for outdoor, battery-powered devices. Support for navigation will thus increasingly find its way into mobile phones. Some such devices already exist, e.g., the HP iPAQ hw6915, and others have been announced, e.g., the Nokia E95 which is expected to be available during the first quarter of 2007.

The first satellite of the Galileo positioning system has already been launched, and Galileo is expected to be operational around the turn of the decade [W07]. Galileo will offer better positioning than does GPS with respect to several aspects, including the accuracy, penetration, and time to fix [B06]. For example, the best-case accuracy (without the use

of ground stations) of Galileo is 45 cm as opposed to 2 m for GPS. Next generation GPS will also offer better positioning, and Galileo and GPS are expected to be interoperable.

This emerging infrastructure has the potential for enabling entirely new, geo-enabled services that were either not relevant or of little use in fixed desktop computing settings.

The range of possible applications and services is virtually limitless. For example, it includes services related to traffic and transportation such as “fleet” management, including emergency vehicle dispatching and hazardous cargo and traffic offender tracking; road-pricing where payment is dependent on where, when, and how much a vehicle drives; and other “metered” services, such as insurance and parking. It includes services that warn drivers about accidents, slow-moving vehicles, and icy and slippery road conditions ahead. It also includes a wider range of safety-related services, such as services that track senile senior citizens, tourists traveling in potentially dangerous environments, and prisoners serving time at home. Next, it includes the oft-mentioned point-of-interest services that offer information concerning gas stations, hospitals, etc. It also includes the emerging and challenging area of games and “-tainment” (edu-, info-, enter-) services. One theme is to move games from occurring in a virtual world behind a small computer screen to instead occurring in reality. Virtual objects, e.g., treasures (or caches, cf. geocaching [G07]), monsters, and bullets, are given geographical coordinates, and the coordinates of the real participants are also known by the system. This arrangement then enables games that aim to find treasures, catch or escape monsters, and hit with (virtual!) bullets.

This paper covers aspects of the *streamspin* project [S07] that has as its objective to identify and provide solutions to data management challenges on the mobile Internet. In particular, the project aims to identify fundamental data management services that will be needed by many mobile services; the project aims then to subsequently provide solutions.

The project will adopt a prototyping approach where innovative mobile services are prototyped and made available to users. In doing so, it is an objective to understand and learn from the services and technologies found on the conventional Internet—we hope this will enable us to prototype novel and viable mobile services. To facilitate the approach, we will build a testbed platform that will allow us to rapidly develop novel services and to experiment with data management aspects of these. The platform will fuel an Internet portal through which services are made available to users.

Section 2 argues that the mobile Internet is different from the conventional Internet and describes some of the differences. It also points to consequences of these differences. Section 3 then covers the *streamspin* Internet portal and service platform, including the functionality of the portal and the architecture of the platform. Section 4 adopts a geographical focus and describes the geo-related content and context that are relevant for mobile services. Next, Section 5 discusses self-imposed guidelines for the development of new data management technologies in the research project. Finally, Section 6 summarizes the paper.

2 The Mobile Internet is Different!

Users of the conventional Internet by far and large sit in front of a desktop or laptop computer with a relatively large screen and a convenient qwerty keyboard. And they are in a controlled environment, typically either at work or at home.

The mobile Internet is markedly different. Here, the services are typically delivered to mobile devices with small screens and either without a qwerty keyboard or with a very inconvenient one. The user is out and about, so the use situation is very varied.

The user may be sitting in a cafe or may be in a meeting. Or the user is on the move: by foot, using a collective means of transport, on bicycle, or driving a car. The user may be traveling in a familiar environment or may be in unfamiliar surroundings and alone. In many use situations, the user's main focus of attention is not the service, but rather something else, e.g., navigating safely in traffic. The user is often engaged in a particular, primary activity.

These characteristics of the use situations are likely to have several consequences. In particular, many mobile services will aim to assist the user with the user's primary activity. It thus becomes important to be able to anticipate the user's needs. The varied use situations offer important clues as to what the primary activities are. The use situations and likely activities, and thus user needs, may be captured in a user's context.

Also, it is important to utilize as best as possible the more constrained interaction between the mobile device and its user, as caused by the small screen, inconvenient keyboard, and, at times, disruptive surroundings. For example, doing a conventional Google search on a mobile phone is possible, but far from convenient in all use situations. Put differently, it is important to deliver the right service at the right time, with minimal system interaction. Again, the notion of context is useful.

Although context awareness is particularly important for mobile services, there is no general agreement on what mobile context is—rather, the notion of context is open ended, and its substance should probably depend on its intended use. It may include relatively static elements such as the user's demographic data (e.g., age, education, gender) and interests and preferences (e.g., I like sushi and Thai food, and I prefer outdoor products from Arc'teryx and Falke). It may include a collection of profiles that the user has defined and may activate and deactivate. For example, a user may define a "I am hungry" profile, a "I want to shop" profile, and a "Keep me informed of sports news" profile.

The user's current location-related data is also a possible element of a user's context. As discussed earlier, accurate geo-positioning is becoming increasingly feasible. The location is an example of an attractive type of context, namely the type of context that can be acquired automatically, without requiring any activity on the part of the user. A GPS device will report not only a location, but also a heading and a speed.

A related kind of geo-context is the user's destination, and yet another is the route that takes the user from the current location to the destination. Routes are interesting for two reasons. First, a mobile user typically travels towards a destination (rather than moving around, aimlessly), and a user often follows the same route when going from one location

to another. For example, a user typically travels on the same route from home to work. Second, routes are significant as context for many services. For example, a service that knows the route of a user may alert the user about road conditions, e.g., congestion and accidents, on the route ahead, while not bothering the user with conditions that do not relate to the user's route. Routes may also be used when a user requests the locations of "nearby" points of interest. Thus, a service may suggest gas stations to the user that are near to the user's route, rather than merely to the user's current location.

Another consequence of the characteristics of the mobile Internet is that many mobile services will be so-called push services, i.e., services that a user registers for and that then automatically are provided to the user when certain conditions are met (e.g., [FZ98]).

An example illustrates this. Assume that a user has registered for a traffic-condition service. The service knows the user's current location, destination, and route to the destination. These elements are part of the user's context. The location is provided by a tracking system that utilizes a GPS device. The destination and route are provided by a navigation system or predicted by a route-monitoring system that has accumulated a record of the user's past travel behavior. When travel conditions are normal, the service is inactive—only when an accident or some other event has led to an anomalous condition that affects the user's current travel does the service become active and supplies the user with, e.g., re-routing information or an updated arrival time prediction.

3 The *streamspin* Portal and Testbed Platform

We describe first the functionality of the portal, then describe the testbed platform that enables the functionality made available via the portal. Both the portal and platform exist in early prototype versions, and they will evolve throughout the course of the project.

3.1 Portal Functionality

The *streamspin* Internet portal is an open, context-aware, and regulated portal for mobile content and services. The portal is open in the sense that any registered user may supply and use content and services.

Services may be developed using an application programming interface that is made available to developers, but services may also be created by completing service creation templates. For example, parents that need to transport their kids in connection with soccer practice and matches may complete a web form at home that will then generate a coordination service. The templates also allow easy setup of widely usable streams based on personal content. Examples include email, RSS and blog push. Finally, the portal also allows service providers to plug in their own templates that then allow users to create, and subsequently subscribe to, services that are delivered by the service providers. For example, it might be natural for a company to supply a customized email service to its employees by using this functionality.

The portal is context aware in the sense that users can enter demographic data and can create and activate profiles, as described briefly in the previous section. In addition, the users can supply dynamic context data, such as the current location, to the portal as well as to services subscribed to via the portal.

The portal is regulated in the sense that only legal content and services are allowed. In addition, the services registered at the portal must respect the contexts of the users. Thus, services should not push content to a user who has activated the “Do not disturb me” profile.

The portal acts as an intermediary between service and content providers, on the one side, and the users, on the other. The former then submit and register content and services, while the latter supply their context and subscribe to services. A main task of the platform underlying the portal then is to find matches between contexts and services and content. In our example from earlier, the platform must find matches between the route ahead for a user and events that occur. A match exists when an event is likely to affect the driving conditions of the user.

The platform is intended to serve as a convenient testbed for the exploration of new concepts in relation to mobile service delivery. Thus, the platform is designed to be open and extendable.

From a technical perspective, the platform supports two quite different kinds of services, so-called portal services and third-party services.

Portal services are the simplest kind. The platform receives streams of content from content providers. These may include RSS-feeds with traffic information, weather information, news, etc. Content broadcast via TMC [T07] also fits into this category. For each user who has registered for content, the platform matches all incoming content against the user’s context and then passes matches on to the user. To control the matching, the content providers are offered the opportunity to tag all content with metadata obtained from a hierarchy provided by the platform. With portal services, no third parties have access to the context of a user.

The third-party services differ from the portal services in that they by default pass contents directly to the users. This is natural for communication services, e.g., e-mail, and for high-quality services, e.g., different location-based services. The providers of third-party services also have access to the users’ contexts. Thus, a third-party service can provide a “guided” city tour for tourists or provide points of interest that are near to the service users.

3.2 Testbed Platform Architecture

The overall four-layer system architecture underlying the *streamspin* portal is shown in Fig. 1. Each box is a software component and an upper layer uses the layers below to provide a given functionality.

The information layer at the bottom has four main components. The context encompasses

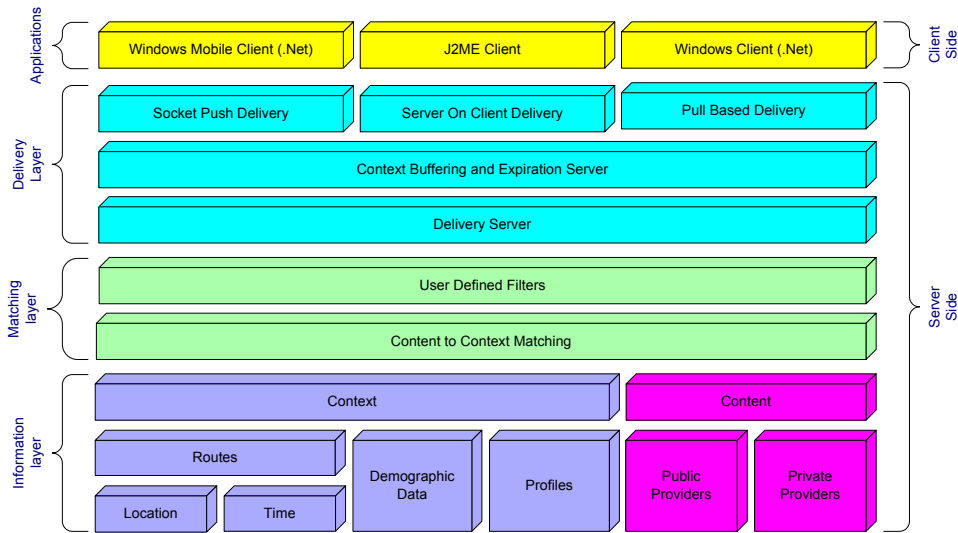


Figure 1: Testbed Platform Architecture

location and time (which may be aggregated into routes). The layer also contains software components for demographic data and profiles. As discussed earlier, a user may have a number of profiles, one or more of which are active. Finally, content is provided by public or private content and service providers.

Next, the matching layer accomplishes the matching of the content against the user's context, including the user's demographic data and active profiles. The platform allows the users to specify additional user-defined filters. This is to enable further personalization of the content provided to a user.

The third layer is the delivery layer, which is responsible for the actual delivery of content to the mobile devices. In this layer, the buffering and expiration component is quite important. It ensures that content for a device that is offline is delivered next time the device becomes online if the content remains relevant. This component is also responsible for ensuring that old (expired) content is not delivered to a device.

The delivery layer offers multiple ways of delivering content to the devices. Overall, it applies both push or pull principles. Each user can specify a preferred means of obtaining content based on, e.g., the technical capabilities of the user's terminal and the mobile services available.

The three layers discussed so far all reside on the server side. The fourth and final layer, the application layer, resides on the client side, i.e., on the mobile devices. This layer's responsibility is to receive the content from the server side and display it on the device. To enable this functionality, a user only has to download a single program from the server and install this on the device. This program is capable of receiving and presenting the content delivered by all services. The addition and deletion of services on the server thus does not affect the software installed on the clients.

The platform currently supports devices that use the Microsoft Mobile 5.0 operating system (with the .Net compact Framework 2.0) as well as Java enabled devices (with MIDP 2.0).

4 Data Management Foundations with a Geographical Focus

This section adopts a data centric perspective on mobile services. It discusses the various types of geographical content and context of relevance for mobile services, including business content (e.g., point-of-interest data), generic geo-content, also termed infrastructure, and user-specific geo-context.

We are working to integrate support into the *streamspin* platform for mobile services that involve these types of data. This includes support for positioning by other means than GPS, it includes tracking services, and it includes services that enable the prediction of a user's geographical context such as the user's destination and route. The section borrows from an earlier paper by the author [Je06]. Two recent books are recommended as further readings [GS05, VS04].

4.1 A Data Centric Perspective

The data centric perspective may be formulated as follows: By capturing pertinent aspects of reality in digital form—in semantically rich and appropriately organized structures, and with powerful update and retrieval techniques available—an ideal foundation for delivering a wide range of mobile services is obtained.

The idea is that a service request by a mobile user translates into queries against the database envisioned in the statement. A key challenge in the delivery of mobile services then becomes a data processing problem.

The phrase “digital mirror of reality” has been used for describing the envisioned database. While this phrase is intuitive and paints the right kind of mental picture, it only partially reflects the desirable capabilities of this database, which go well beyond simply being a mirror. In particular, the database may capture past states of reality and one or more perceived future states, in addition to the current state. In more technical terms, the database supports the valid-time aspect of data. Further, if accountability is a concern, the database may include an incorruptible record of its past states. In technical terms, this is called transaction-time support.

Next, the database and the database management system used may not be a single relational or object-relational database stored in a centralized system. Rather, the database and system may well be distributed and heterogeneous in a number of respects. For example, the data may be physically distributed and may not adhere to the same common schema or data model. The control and data processing may also be distributed.

4.2 Infrastructure and Business Content

The delivery of geo-enabled mobile services in practice is dependent on relevant content being available. Examples of content include weather data; traffic condition data, including information about accidents and congestion; information about sights and attractions, e.g., for tourists; information about hotel rooms, etc. available for booking; and information about the current locations of populations of service users.

The management of such content includes several aspects. An information technology infrastructure must exist that is capable of capturing the content and capable of absorbing the content as it is made available, while being able simultaneously to make the content available to services.

We may distinguish between two types of content: the geographical infrastructure itself and all the other, “real” content that may be given geographical references and that must reference the infrastructure. Points of interest exemplify real content.

The geographical infrastructure, or geo-content, concerns the geographical space itself, with hills, lakes, rivers, fjords, etc. It also concerns the road networks for use by vehicles and the transportation infrastructures for, e.g., pedestrians, trains, aircraft, and ships. The infrastructure for vehicles is of high interest because users may frequently be either constrained to, or at least using, this infrastructure.

Geo-content is important. Users think of the real content as being located in a transportation infrastructure, and they access the content via the infrastructure. For example, the location of a point of interest is typically given in terms of the road along which it is located, and directions for how to reach the location are given in terms of the transportation infrastructure.

For the delivery of a range of geo-enabled mobile services, it is particularly important that a representation of the road infrastructure is available that supports multiple functions, including content capture; content update and querying, including route planning and way finding; and user display. This representation may be composed of several constituent representations [Je04].

It is common practice to specify the location of some content relative to the nearest kilometer post along a specific road. For example, the entry to a parking lot may be indicated by a road, a kilometer post on that road, and an offset. So-called linear referencing may be used when capturing such content.

A weighted, directed-graph representation may also be used that represents a quite abstract view of an infrastructure. This representation ignores geographical detail, but preserves the topology, and it may be used for connectivity-type queries, such as route guidance and way finding.

Next, a geo-representation is also needed that captures the geographical coordinates of the road infrastructure. With this representation, it is possible to map a location given in terms of geographical coordinates, e.g., from GPS receivers or point-on-a-map-and-click interfaces, to a location in the infrastructure. Finally, these three representations must be connected.

All the “real” content encompasses any content that may reference, directly or indirectly, the geographical infrastructure. A museum, a store, or a movie theatre may have both a set of coordinates and a location in the road network. This type of content is open-ended and extremely voluminous. For example, it may include listings of movies currently running in the movie theatre; it may include seat availability information for the different shows; and it may include reviews of the movies. Often, the real content is the primary interest of the users.

Content is generally dynamic. This applies to road networks, where road construction and accidents change the characteristics of the networks with varying degrees of permanence. Other content is also dynamic. Examples abound. New stores open and existing stores relocate or close. The opening hours of a facility may change. The program of a movie theater changes. The sales available in a store change. This dynamicity of content implies that a representation of content must be designed to accommodate updates.

Content is more or less dynamic. The content that derives from the sampling of the positions of moving objects belongs at the highly dynamic end of the spectrum. Capturing the present positions, and possibly the past as well as anticipated future positions, of a large population of mobile users requires special techniques, as discussed next.

4.3 Geo-Context—Locations, Destinations, Routes, and Trajectories

User-specific geo-context is another kind of content. Among such content, the current position of a service user is the traditional geo-context used in location-based services.

To maintain an up-to-date record of the current position of a service user, we may envision a scenario where a central server maintains a representation of the user’s movement and where the local client, e.g., a mobile phone, is aware of the server-side representation. The client frequently compares its GPS position to the server-side position, and when the two differ by a threshold slightly smaller than the accuracy required, an update is issued to the server, which then revises its representation of the client’s movement and sends this new representation to the client [CJP05]. This arrangement, termed shared-prediction-based tracking, aims to reduce the number of updates needed in order to main a current position at a given accuracy.

Different representations of a user’s movement result in different rates of update. We consider several possible representations in turn. First, we may represent the movement of a user as a constant function, i.e., as a point. With this representation, an update is needed every time the user has moved a (Euclidean) distance equal to the threshold away from the previous position. This is a simple representation, and it may be useful when the user is barely moving or is moving erratically within an area that is small in comparison to the area given by the threshold used.

Second, we may represent the movement by a linear function, i.e., by a vector. When the user exceeds the threshold, the user sends the current GPS location and the current speed and direction (which GPS receivers also provide) to the server. The server then uses this information to predict the user’s to-be-current positions.

Third, we may utilize the infrastructure in the representation of a user's movement. This requires that we are able to locate the user with respect to the infrastructure. One possibility is to assume that the user is moving at constant speed along the road on which the user is currently located. We may use the GPS speed as the constant speed, and we may assume that the user stops when reaching the end of the current road segment. Depending on the lengths of the segments, this representation can be expected to be better or worse than the vector representation. However, for realistic segments, this representation has the potential for outperforming the vector representation.

Next, we may use the route of the user in place of the segments. As argued earlier, most humans who travel do so towards a known destination; and being creatures of habit, and perhaps for efficiency, we tend to follow routes we have previously followed. Therefore, we will frequently be able to predict correctly the route on which a service user travels [BJ07]. Using the correct route in place of a road segment means that the number of updates needed to maintain a user's position with the desired accuracy decreases further. Indeed, updates occur only because of incorrectly predicted speeds—no updates are caused by incorrectly predicted “locations.” It should also be observed that if a route is predicted incorrectly, e.g., because the user makes a turn, this does not lead to a breakdown. Rather, this simply forces an update and a new prediction.

The infrastructures currently available for mobile services support the accumulation of GPS data from vehicles. Based on this data, it is possible to gradually create usage patterns for vehicles that consist of the routes traveled by the vehicles along with usage metadata, which are temporal patterns that describe for each vehicle and route when the route is being used by that vehicle [BJ07]. For example, a pattern may specify that a route is being used in the morning on weekdays. The resulting route and destination data may subsequently be used in services. By also attaching travel speeds to routes, we obtain trajectories, which are routes “lifted” into the time dimension [GS05].

5 Research Challenges

We proceed to present six guidelines for conducting research that we are trying to follow in our ongoing research. These are intended to apply to the area of geo-enabled, mobile services, but are to varying degrees also applicable to other areas of research. (Reference [Je06] relates the guidelines to past research that the author has been involved in.)

For application-oriented research, estimate the time of application and formulate expectations to the reality as of that time; then design for that reality.

Different research activities may aim to be applicable within different time horizons. For research that is intended to be applicable in the near term, we must take care to only make assumptions that are met by current infrastructures. Assumptions concern the available computing and storage capabilities of mobile terminals, the available communication technologies, the available positioning technologies, the available digital road networks, and existing legislation.

Other research may only be applicable in the longer term and perhaps only indirectly. Such research may be more speculative and somewhat more removed from specific applications. Some of the results may not offer the final answers, but may serve as a foundation for further work.

Ensure that at least one appropriate architectural setting exists or may be envisioned for the research contribution.

For research that is expected to have practical application in the short term, the architectural setting is likely to be a concern. For other research, it may be sufficient to ensure that an appropriate architectural setting exists or can be envisioned. And for yet other research, architectural settings may not be an important concern.

Invent solutions for composable functionality.

When research on query processing in relation to moving objects first took off, the efficient processing of many basic types of queries had yet to be explored in the new moving-object settings. Examples included one-time and continuous range queries, nearest-neighbor queries, and reverse nearest-neighbor queries, to name but a few.

As techniques for the processing of these basic types of queries accumulate, it is natural that attention shifts to as yet unexplored or lightly explored types of queries. A potential pitfall is that we start producing highly optimized solutions to very specialized types of queries. This path is not advisable, as the prospects of these solutions finding practical applications are likely to decrease with the degree of specialization of the functionality.

Prioritize versatile and robust solutions over specialized and brittle, although possibly highly performant, solutions.

One lesson to be learned from current, commercial data management technology and existing applications is that versatile and robust solutions have better chances of finding practical use than do very specialized ones, even if these exhibit very high performance in some cases. The objective of a query optimizer is quite modest: it should avoid the clearly inefficient ways of computing a query and identify one good way of executing the query. (Even meeting this objective can be a challenge.)

In the area of data processing for moving objects, the parameter space—the numbers of parameters and parameter settings that characterize a data processing workload—is very large. One consequence is that there is room for many solutions that offer superior performance for certain settings, but may be clearly inefficient for other settings.

The recommendation is that we try to aim for solutions that are versatile in terms of the functionality they offer and that are robust in terms of the settings. A solution—for which there exist other solutions so that for every possible parameter setting, at least one of these other solutions has twice the performance—may still be preferable if it is much more robust than its competitors.

Design query processing techniques that exploit the entire context.

When we design query processing techniques, we should try to use as arguments all the context that we can reasonably expect to have available. So if we can assume to know the likely route of a moving vehicle, we may suggest restaurants or gas stations to the driver that are near to the expected route, rather than merely to the driver's current location, which is the best a service can do if it ignores the route. Moreover, utilizing knowledge of the road network, we can use network distances as opposed to Euclidean distances in our calculations, and we can augment the answers with travels and detour distances and with suggested routes to the points of interest returned.

Pay attention to both query performance and update performance.

Many indexing and query processing techniques for geographical data were originally developed for largely static data. For example, R-trees do not contend well with workloads with frequent updates.

In contrast, mobile-service application scenarios exist that are characterized by frequent position updates. This puts focus on techniques that are capable of supporting workloads consisting of frequent updates as well as queries, and it puts focus on studies of the trade-off between query performance and update performance [Je02, JS02].

Updates of moving-object positions correspond to the sampling of continuous, position-valued variables. Different services may tolerate different position inaccuracies. For example, a localized-weather service may tolerate a relatively high degree of inaccuracy without this affecting the functionality of the service, while a navigation service is dependent on more accurate positions.

One approach to taking advantage of the different accuracy tolerances of different services is to perform updates only when needed to maintain the accuracies needed. Indexing and query processing techniques should be able to exploit this approach to updates. Also, by forming predictions of the future movements of the objects, the numbers of updates can be further reduced.

6 Summary

A new and exciting Internet is emerging—the mobile Internet. While it is difficult to predict the character of this Internet, this paper argues that the mobile Internet will entail services that are intended to aid their users with the varying, primary tasks that they aim to accomplish while being mobile. It argues that services should be able to anticipate the primary tasks of the users and should reduce the interaction with the mobile devices needed in order to benefit from the services. The paper also suggests that mobile services to a larger extent than conventional Internet services will be so-called push services. Context awareness is suggested as a key characteristic of mobile services.

The paper describes a mobile service platform and associated portal that are being developed in a new project that aims to explore data management foundations for mobile

services. The idea is to learn from and apply some of the many concepts underlying novel services already available on the conventional Internet for the creation of novel mobile services. Such services are then prototyped and exposed to users via the portal. Based on lessons learned from this activity, novel data management services are being added to the platform and portal that the mobile services may exploit to realize their functionality while ensuring that the platform is scalable. The project will focus on geographical data management services, e.g., services that capture the geographical aspect of data and the users' geo-contexts. In an attempt to invent technology that is as relevant for practice as possible, six self-imposed guidelines have been identified.

The platform and portal exist in prototype versions and are intended to evolve throughout the duration of the project that is scheduled to last for 4 years. We are currently integrating tracking and route-recording services into the platform (as described in Section 4.3), and we are also adding new service creation templates to the system.

Acknowledgments

I would like to thank the other three current key participants of the *streamspin* project: Kenneth H. Pedersen, Kristian Torp, and Rico Wind. I would also like to acknowledge those other local colleagues with whom I currently conduct research on the topics covered in this paper: Agnė Brilingaitė, Alminas Čivilis, Xuegang Huang, Harry Lahrman, Hua Lu, Stardas Pakalnis, Torben Bach Pedersen, Simonas Šaltenis, Albrecht Schmidt, Laurynas Speičys, Dalia Tiešytė, Nerius Tradišauskas, and Man Lung Yiu.

My interactions with these dedicated and capable individuals have shaped the thoughts and views expressed in this paper, although only I should be held accountable for any inadequacies of these.

References

- [B06] BBC News. Q&A Europe's Galileo Project. news.bbc.co.uk/1/hi/sci/tech/4555276.stm, 2006.
- [BJ07] Agnė Brilingaitė and Christian S. Jensen. Enabling Routes of Road Network Constrained Movements as Mobile Service Context. *GeoInformatica*, 55 pages, to appear.
- [C06] Cisco. GPRS White Paper. Technical report, 19 pages, 2006.
- [CJP05] Alminas Čivilis, Christian S. Jensen, and Stardas Pakalnis. Techniques for Efficient Tracking of Road-Network-Based Moving Objects. *IEEE Transaction on Knowledge and Data Engineering*, 17(5): 698–712, 2005.
- [FZ98] M. Franklin and S. Zdonik. "Data in Your Face:" Push Technology in Perspective. In *Proceedings of the 1998 ACM SIGMOD International Conference on Management of Data*, pp. 516–519, 1998.
- [G07] Groundspeak. Geocaching. www.geocaching.com, 2007.

- [GS05] Ralf Hartmut Güting and Markus Schneider. *Moving Objects Databases*. Morgan Kaufmann, 2005.
- [Je02] Christian S. Jensen. Special Issue on Indexing of Moving Objects. *IEEE Data Engineering Bulletin* 25(2), 2002.
- [Je04] Christian S. Jensen. Database Aspects of Location-Based Services. Chapter 5 of [VS04], pp. 115–147.
- [Je06] Christian S. Jensen. Geo-Enabled, Mobile Services—A Tale of Routes, Detours, and Dead Ends. In *Proceedings of the 11th International Conference on Database Systems for Advanced Applications*, pp. 6–19, 2006.
- [JS02] Christian S. Jensen and Simonas Šaltenis. Towards Increasingly Update Efficient Moving-Object Indexing. In [Je02], pp. 35–40.
- [N03] Nokia. A History of Third Generation Mobile. Technical report, 16 pages, 2003.
- [N06] Nortel. HSDPA and Beyond. Technical report, 28 pages, 2006.
- [S07] The *streamspin* team. The Evolving *streamspin* Internet Portal. streamspin.com, 2007.
- [T07] TMC Forum. TMC Forum Homepage. www.tmcforum.com, 2007.
- [VS04] Agnes Voisard and Jochen Schiller. *Location-Based Services*. Morgan Kaufmann, 2004.
- [W07] Wikipedia. Galileo Positioning System. en.wikipedia.org/wiki/Galileo_positioning_system, 2006.
- [W06] Wimax Forum. IEEE 802.16a Standard and Wimax Igniting Broadband Wireless Access. Technical report, 7 pages, 2006.

References containing URLs are valid as of January 1, 2007.