CHAPTER 39

Location-Based Services Using WiFi Positioning on a Wireless Campus

Barend Köbben

39.1 The Wireless Campus at the University of Twente

In June 2003 the "Wireless Campus" was inaugurated at the University of Twente (UT), allowing cable–free internet access to staff and students anywhere on campus. University of Twente is a young university in the Eastern part of The Netherlands. It employs 2,500 people and has over 6,000 students. On its campus, the university has 2,000 student rooms. The university campus is situated between the cities of Enschede and Hengelo, near the Dutch-German border. Spread over the 140-hectare campus 650 individual wireless network access points have been installed, making it Europe's largest uniform wireless hotspot. Anyone with a PC, laptop, PDA or other WiFi (wireless fidelity)-enabled device can access the university's network and the internet from any building, the campus park and other facilities without cabling.

University of Twente's Wireless Campus aims at a broad range of research and applications of wireless and mobile telecommunication. The UT wants to use the WLAN (Wireless Local-Area Network) in cooperation with the adjacent Business and Science Park and is busy covering this by access points as well. Furthermore, a project has just started in cooperation with Enschede Municipality to install further access points to cover the downtown area of Enschede.

The wireless network facility was made possible with financial support of the Dutch Ministry of Economic Affairs and has been built in cooperation with IBM Netherlands and Cisco Systems. It consisted in the first instance mainly of access points that use the 802.11b wireless networking standard, offering a data transfer speed of 11 megabits per second for most users. There is an ongoing effort to upgrade the entire network to the 802.11g standard, providing data at speeds up to 56 megabits per second.

Research projects investigate the technology and the applications of wireless and mobile communication in several ways, mostly in cooperation with industrial and other knowledge partners. The Wireless Campus has become a 'test bed' for wireless and mobile applications. The major part of this research takes place at the Centre for Telematics and Information Technology (CTIT) and the research institute MESA+. Both are key research institutes of the University of Twente. MESA+ is an institute that conducts research in the fields of nanotechnology, microsystems, materials science and microelectronics. CTIT is an institute that conducts research into the design of advanced ICT (Information, Communications and Technology) systems and their application in a variety of application domains. Its Computer Architecture Design and Test for Embedded Systems Group became interested in using the WiFi technology in the wider framework of the SmartSurroundings research program. This program is

> *investigating a new paradigm for bringing the flexibility of information technology to bear in every aspect of daily life. It foresees that people will be surrounded by deeply embedded and flexibly networked systems …. This presents a paradigm shift from personal computing to ubiquitous computing, …. Relevant knowledge areas include embedded systems, computer architecture, wireless communication, distributed computing, data and knowledge modelling, application platforms, human-computer interaction, industrial design, as well as application research in different settings and sectors"* (Havinga et al. 2004, page 64).

An important part of such systems is establishing the position of persons, services and devices, and one of the possible strategies to achieve that is to use the WiFi network.

39.2 Positioning Using WiFi Technology

Using WiFi technology is just one of the many ways available of using wireless networks for positioning of (mobile) users. An overview of wireless location papers, a website maintained by Youssef (2005), currently distinguishes some 10 categories, among them mobile phone network– and GPS–based techniques. For an overview of many of these techniques, see Hightower and Borriello (2001).

There are various reasons to choose WiFi-based localization over other methods. The most prominent possibly is the fact that it is an economical solution. Because the wireless network infrastructure already exists, localization can be done by software-only methods without adding any additional hardware. This is also true of cell-phone-network-based solutions, but these currently offer only poor accuracy. Secondly, compared to other indoor techniques such as InfraRed, Bluetooth or RFID, the range covered by WiFi is significantly larger. And although coverage is not as ubiquitous as mobile phone networks or GPS, WiFi-based WLANs are being installed at an ever-increasing rate all over the world in public places like airports, conference centers and shopping malls, as well as in universities, hotels, offices and such.

39.2.1 Positioning Methodologies

There are different basic methods of using WiFi signals for determining the location of users (Muthukrishnan, Meratnia et al. 2005), shown in figure 39-1. Hardware-based systems use additional hardware on top of the existing infrastructure to determine characteristics such as Time of Arrival or Angle of Arrival of the signal received from known fixed locations of the WLAN Access Points (APs). Triangulation and other geodetic techniques can then be used to calculate positions. Apart from the disadvantage of the additional hardware need, there is also the problem that the signals are reflected from various objects, especially indoors. Because of this multipath environment, these techniques are complex and reliable results are difficult to obtain.

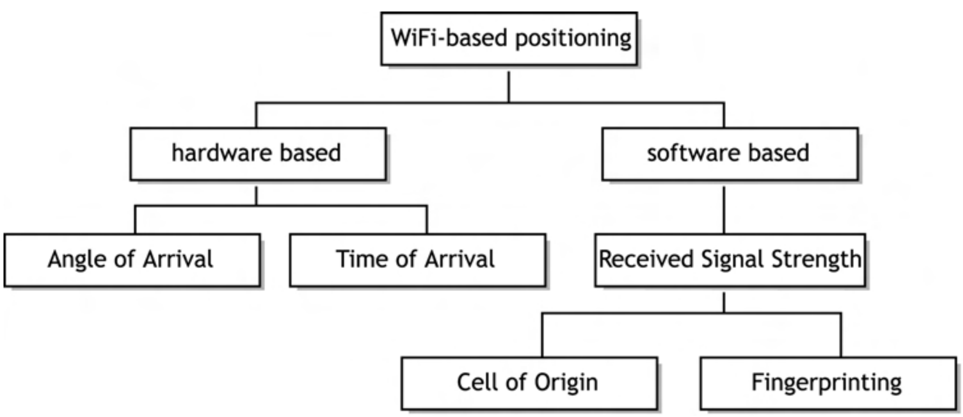

Figure 39-1 Simplified taxonomy of WiFi-based positioning systems, after Muthukrishnan, Meratnia et al. (2005).

The *software-based* systems mostly use the *Received Signal Strength* (RSS). The big advantage of RSS-based techniques is that we can use the existing infrastructure to deploy a positioning system without additional devices, other then the standard WLAN network card in the computer or PDA.

A very simple method is to determine which AP's signal is the strongest one received and then assume the location to be in the area that is covered by this AP's signal. This *Cell of Origin* method, that is similarly used in mobile phone networks, will generally result in very coarse-grained location information.

More-accurate methods use so-called *location fingerprinting* schemes, where selected characteristics of the signal that are location dependent are stored in a database. These fingerprints are then matched to the characteristics actually measured at the current location of a receiver. Unfortunately, the RSS is a highly variable parameter and issues related to positioning systems based on RSS fingerprinting are not yet understood well. The fingerprint information can be obtained in two ways. Firstly there are *Radio Map*-based methods. They involve location fingerprinting done at as many locations as possible, building up a fine-grained RSS map. When a device requests a location, it sends the signal strengths from all access points it can detect to the database, that finds the closest match and returns that as the probable location. The main drawback is the necessary calibration of signal strength as a function of a particular location, and even particular time (since the radio wave properties in an indoor environment can vary greatly depending on the number of people inside the building). There is a trade-off between the amount of effort put into the calibration (it requires lots of time and work and should be performed repeatedly) and the accuracy obtained. Little research as of yet has addressed the issue of optimizing the calibration effort. Secondly, *Model Based* approaches are based on automatically generated fingerprint information based on models of the signal propagation and detailed information about the geometry and topology of the environment, including the materials used in walls, ceiling, etcetera, as these greatly influence the multipath characteristics of the signal.

In Muthukrishnan, Meratria et al. (2005), the authors defined evaluation criteria for WiFi-based location systems that can be used as guidelines to compare and evaluate several indoor location/positioning systems.

1. *Accuracy and Precision* of estimated location are the key metrics for evaluating a localization technique. Accuracy is defined as the deviation of the estimated position from the true position and is denoted by an accuracy value and precision value (e.g. 15 cm accuracy over 95% of the time). The precision indicates how often we expect to get at least the given accuracy. The accuracy of a positioning system is often used to determine whether the chosen system is applicable for a specific use.
2. *Calibration* is also very important. The uncalibrated readings are highly erroneous and device calibration (the process of forcing a device to conform to a given input/output mapping) is needed. Often there is a trade-off between the accuracy and the calibration effort.
3. *Responsiveness* is defined as how quickly the system delivers the location information. It is an important parameter, especially when dealing with mobility. However, this parameter is mostly ignored in the description of the existing systems.
4. *Scalability* is a significant parameter, as the proposed design should be scalable for large networks. If an approach is calibration intensive then eventually it is not a scalable solution.
5. *Self-organization* is of great importance, as it is infeasible to manually configure the location determination processes for a large number of devices in random configurations with random environmental characteristics.
6. *Cost* is also a crucial issue. It includes the cost of installation, deployment, infrastructure and maintenance.
7. *Power Consumption* is of great concern when running the system in a real environment, especially for mobile devices.
8. *Privacy* includes major concerns: Using localization, it is very easy to create a Big Brother infrastructure that can track users' movements and facilitate the deduction of

patterns of behavior. However, this issue is being generally overlooked in the design of systems and considered as an after thought only. Centralized systems are particularly weak with regard to privacy.

39.2.2 Localization Algorithm in the Prototype

The positioning component currently employed in the project described here is based upon an earlier pilot project called "FriendFinder" (Bockting et al. 2004), done in 2004 for two specific buildings on the University campus. In this pilot a prototype client-server architecture was built, where the client program on the mobile device determined its location by detecting the Access Points (APs) in range and comparing them with data about the APs in a server-side database. This database stored the location in XYZ of the APs, their BSSIDs (Basic Service Set Identifier, the unique identifier of an AP), and the strength of their antenna output (in mW). The application first buffered the RSS measurements because not all APs are detected in any single scan. Then it detected probable faulty measurements and deleted them. The accepted measurements were then put through a filter that calculated their centroid. Now the client had a first estimate of its position. Further filtering took place, and the final estimate was determined by so-called "iterative multilateration." In this technique a client's position, with its estimated inaccuracy, is used by other clients as a reference frame. In that way all nodes use each other's information to jointly improve the accuracy of the positioning. An important part is played by further filters that implement a learning effect from the stored positioning history of the application to further improve the accuracy.

Tests have shown that the FriendFinder pilot achieved an average positioning accuracy just under 5 meters, for non-moving devices. In the current project, the positioning component is part of a wider PhD research (Muthukrishnan, Lijding et al. 2005) that aims at significantly increasing this accuracy. To achieve this, the research investigates the possible approaches, the filters and methods used and the positioning algorithms themselves. An area of further research will be the self-learning abilities of the system, which should make the positioning more accurate over time. The Model Based approach, mentioned in 39.2.1, will be employed to achieve calibration-free localization, preserving quality and accuracy. This is ongoing research and the positioning system used for the current prototype "Flavour" (see 39.4.1) will be updated with the emerging knowledge and methods.

At present, the localization works as follows: The WiFi device inside the laptop or PDA periodically scans its environment to discover WLAN Access Points (APs) in the vicinity. The location of the APs in a 3D coordinate system is maintained in a geodatabase (see 39.3). During the access point scanning phase, the BSSID address of the access points and their Recorded Signal Strengths (RSS) are determined and stored. At any unknown location 'n' in the conference venue, the variation of the signal strength will be:

$$\text{at each } n \text{ RSS varies as } 0 >= \text{RSS} <= \text{MAX}$$

However, the signal strengths are usually contaminated by noise. In order to have a better estimation of the actual location, an *exponential moving average filter* is employed to smooth the signal strength. The formula used in doing so, in which $\alpha=0.125$ and RSS denotes the observed signal strength is:

$$\text{Current RSS} = \alpha * (1 - \text{Current RSS}) + \alpha \text{ (previous RSS)}$$

The accuracy of this method can be improved by introducing previous traces (history) of the users. The knowledge of users' earlier determined locations can be used to predict their movement in a certain area. For example, assume that there are already 'p' predetermined positions prior to the user's new location 'q', then there is a high probability that the new location 'q' is near the vicinity of the predetermined positions 'p'. Due to fluctuations in the signal strength, caused for example by reflections off walls, or by obscuring of the antenna by the device user, the system might return a false location. The algorithm computes the difference in the distance between the last determined locations and the currently calculated one. If the difference is larger than a certain threshold then it is can be assumed the location will be in error. This system can be greatly improved by introducing the internal geometry of the building into the equation. If the system would be able to have a notion of the distances as well as topological relations within the environment, major improvements in the final localization accuracy can be expected.

39.3 Mapping the Access Points

The WiFi-based positioning methods described above are highly dependent on an initial mapping stage, in which the coordinates of all the access points in a 3D coordinate system had to be recorded in a database. For the FriendFinder project mentioned above, only a limited number of the Access Points (APs) had been used. As no geo-information experts were involved at that stage, their positioning was done in a rather improvised way. The height of the APs especially was a problem: it was determined only by estimate and with respect to the building's ground floor height. In this limited project that was not a big problem, as only one building was involved, but for the larger project the elevation differences between the buildings (more than 5 meters, which is a lot for The Netherlands!) had to be taken into account.

The 650 individual wireless network APs that have been installed were only indicated on paper maps, one map per floor, of the individual buildings of the University. The base maps are print-outs from CAD–drawings (blueprints) maintained by the Facility Management Services; they have a high level of detail, but they are not georeferenced and thus have a local, arbitrary, coordinate system that's basically just 'paper coordinates.. Furthermore, the location of the APs had been indicated on these maps haphazardly by hand-drawn symbols at the time of installation of the devices.

Therefore the first task has been the digital mapping of the AP locations in a geodatabase. In order to do this, it was decided to digitize all locations using GIS software and digitally georeferenced versions of the CAD drawings. The georeferencing was achieved by transformation of the CAD drawings, using control points from an overview map of the whole campus that is available in the Dutch national coordinate system "Rijks Driehoeksstelsel." It was possible, when using simple first order transformations, to achieve Root Mean Square Errors of less than 0.1 m.

For all buildings a base elevation was determined in meters above NAP (the Dutch vertical datum) by combining the campus map with the Actual Height Model of the Netherlands, a detailed elevation model of the whole country made by airborne laser altimetry, which has a point density minimum of 1 per 16 m2 and a systematic error of 5 cm maximum (Rijkswaterstaat Undated). In order to get precise location measurements, it was necessary to physically visit all APs and use a laser measurement device to determine the relative location of the AP antenna with respect to the elements of the building present in the CAD drawings (walls, floors, windows). By combining these relative measurements with the georeferenced maps a precise 3D location has been determined and put into a geodatabase. The added bonus is that all APs have been checked and additional attributes were gathered, such as antenna type, antenna connection length (for estimating signal loss), etc.

To efficiently store and use spatial data in a database, it has to support geographic data types and methods, preferably using the standards of the OGC (Open Geospatial Consortium 2007). Among other things, the OGC has set the Simple Features SQL Specification that provides for publishing, storage, access, and simple operations on spatial features (point, line, polygon, multi-point, etcetera) through an SQL interface. There are several commercial database systems with OGC-compliant spatial extensions, of which Oracle is probably the most prominent, but there are also open source alternatives: PostgreSQL, a database system with PostGIS as a spatial extension and MySQL. MySQL's recent versions include OGC-compliant spatial extensions, although not implementing the full set of OGC specifications. The reason for the choice of MySQL in this project was the simple 'light-weight' character of the software, as compared to the complicated though more fully-featured PostGIS. By adhering strictly to the OGC standards it should be straightforward to change or even mix database platforms in the future.

39.4 Wireless Campus Location Based Services

There has recently been a lot of industry and research activity in the realm of "Location Based Services" (LBS), which has been defined in Urquhart et al. (2004, page 70) as "*wireless services that use the location of a (portable) device to deliver applications which exploit pertinent geospatial information about a user's surrounding environment, their proximity to other entities in space (eg. people, places) and/or distant entities (eg. destinations)*". The purpose of the project described here is not the development of 'the' or even 'a' Wireless Campus LBS, but rather to investigate and set up the infrastructure necessary to be the basis for LBS's. It combines input from several research projects with the practical application of new as well as established techniques to provide useful services for the UT campus population. The research mentioned has a wider scope then just this project: the Wireless Campus LBS is intended *to serve as a testbed* for the research as well as *to benefit from* the outcomes of the research.

39.4.1 First Tests: FLAVOUR

The first use case test of the Wireless Campus LBS was to provide the participants of a conference held at the UT grounds in the summer of 2005 with an LBS to help them navigate the conference locations and locate fellow attendees. This conference, SVG Open 2005, the 4th Annual Conference on Scalable Vector Graphics (SVGopen.org), was deemed to be a good testbed as it drew a crowd of some 180 people from 20 countries all over the world, from a very wide field of applications: electronic arts & media, geospatial sciences, information technologies, computer sciences, software developers, Web application designers, etc. They share an interest in Scalable Vector Graphics (SVG), the W3C open standard enabling high-quality, dynamic, interactive, styleable and scalable graphics to be delivered over the Web using XML. Most of them are technology-oriented and there is a high degree of interest in, and ownership of, mobile devices.

The application built for testing by the participants has been called FLAVOUR (Friendly Location-aware conference Assistant with priVacy Observant architectURe). Services offered by FLAVOUR can be categorized into:

- *Pull services*, in which location of attendees plays an important role as the attendee's request will be replied by the system on the basis of their whereabouts. Examples of pull services offered are:
  – Finding fellow attendees;
  – Locating resources available in the infrastructure such as printers, copiers, coffee machines etc.
- *Push services*, in which individual and bulk messages are sent to the attendees. This enables the attendees to:
  – Be notified about important events by conference organizers;
  – Communicate with their contacts, i.e., colleagues, friends, etc.

The system architecture, described in more detail in (Muthukrishnan, Meratnia et al. 2005), is based on a Location Manager, which provides services using the Jini platform (Sun Microsystems 2007). Jini is a Java-based open architecture that enables developers to create network-centric services. Each Location Manager registers with the Jini Lookup Service to offer the location of the user it represents. Interested users can look up the service and subscribe to the location of a given conference participant. This is done using a publish–subscribe mechanism. The Location Manager uses a privacy policy to decide if a client is allowed to subscribe to the location of its owner (publisher). It also publishes to all the subscribers relevant changes in the location of its owner.

The Jini architecture also provides other kinds of services, such as a message board to which every conference participant can subscribe. The message board can be used by the conference organization to publish changes in the schedule, information related to the social events, etc. Participants can also use the message board to make announcements to the other participants, as for example asking about lost objects, or to chat.

The graphical depiction of the maps and the location of the users is done in SVG, providing vector graphics in high graphical quality with a small memory– and file–footprint. The system also provides the user with an estimation of the current positioning accuracy. A screen dump of the user interface can be seen in Figure 39-2.

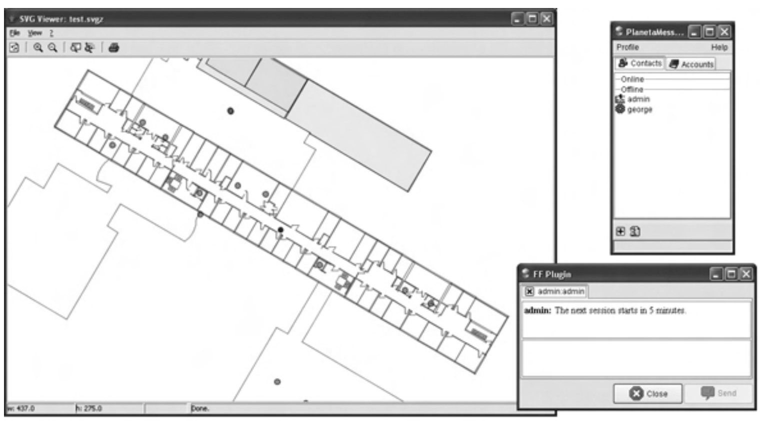

Figure 39-2 Screen dump of the FLAVOUR user interface.

The tests at SVG Open 2005 were relatively successful. Most conference participants experimented with the localization features of the system. The messaging and friend-finder functions were used to a lesser extent. Various extensive interviews have been held with test persons and also written feedback was collected. The localization functionality worked quite reliably, although the accuracy varied quite a bit over the different conference locations. In the computer science building the results were clearly better then in the main conference halls. The tests still have to be analyzed further, but the most obvious reasons are the non-optimal configuration of access points and the fact that the database of these access points still was incomplete at the time of testing.

39.5 Outlook

The implementation of the Wireless Campus LBS described here has only just started. But as it builds on the solid foundations of the well-established infrastructure of the Campus-wide WLAN at the University of Twente, and has had a successful pilot in the FLAVOUR tests at SVG Open 2005, we expect that it will be put into use and expanding relatively quickly in the coming years. On the client-side of the system, ongoing research at ITC on data dissemination for LBS and mobile applications (Köbben 2004) will be concentrating on the Wireless Campus LBS as a testbed for adaptive, task-oriented delivery of mapping information to mobile users using *cartographically aware database objects*. Cartography and GIS more and more involve the use of spatial database technology. In the database world, there is a growing focus on *context awareness* of database objects. One of the important context parameters is location awareness, which is important for all spatial applications. Our goal is to extend context awareness with the idea of database objects that are *cartographically aware*. As a simple example one might think of a location map in the Wireless Campus LBS where the spatial distance of the objects to the focus of the map (the user's position) influences their representation: Interior walls only show themselves in buildings that the user is in, room numbers only if the user is within reading distance. But more complicated systems could be thought up, which might be especially useful for generalization techniques that are hard to achieve in traditional layer-based systems. It is our intention to use the Wireless Campus LBS to test these concepts in practice and provide a proof-of-concept application.

Probably the most exciting aspect of the project is the fact that it provides the opportunity for a very diverse group of people from quite different disciplines to contribute to a technical infrastructure that can serve as a testbed for their respective research projects, and at the same time has the potential to become a useful everyday feature for mobile users at the University Campus.

Acknowledgements

The Wireless Campus LBS project started in early 2005. It's an informal co-operation between people at the University of Twente (UT) Computer Architecture Design and Test for Embedded Systems group (Arthur van Bunningen, Kavitha Muthukrishnan, Nirvana Meratnia, Georgi Koprinkov), the UT department of Information Technology, Library & Education (Sander Smit, Jeroen van Ingen Schenau) and the International Institute for Geo-Information Science and Earth Observation (Barend Köbben).

References

Bockting, S., N. Hoogma, V. Y. Nguyen, M. J. Ooms and G. D. Oostra. 2004. FriendFinder Project 2003/2004, Localisatie met Wireless LAN. MSc project. Enschede, The Netherlands: University of Twente, 54 pp. <http://www.sanderbockting.nl/documents/FriendFinder.pdf> Accessed 23 October 2007.

Havinga, P., P. Jansen, M. Lijding and H. Scholten. 2004. Smart Surroundings. Pages 64–69 in *Proceedings of the 5th PROGRESS Symposium on Embedded Systems*, held in Nieuwegein, The Netherlands, October 2004. ISBN 90-73461-41-3.

Hightower, J. and G. Borriello. 2001. A Survey and Taxonomy of Location Systems for Ubiquitous Computing. *Computer* 34(8):57–66.

Köbben, B. 2004. RIMapper - a test bed for online Risk Indicator Maps using data-driven SVG visualisation. Pages 189–195 in *Proceedings of 2nd Symposium on Location Based Services and TeleCartography*, held in Vienna, Austria 28 – 29 January 2004. Edited by G. Gartner. Vienna: Institute of Cartography and Geo-Media Techniques. <http://kartoweb.itc.nl/kobben/publications/RIMapper_paper_Vienna2004.pdf> Accessed 23 October 2007.

Muthukrishnan, K., M. E. Lijding and P. Havinga. 2005. Towards Smart Surroundings: Enabling Techniques and Technologies for Localization. Proceedings of LOCA2005 – co-located with the 3rd International Conference on Pervasive Computing., held in Munich Germany May 2005. Munich: Springer Verlag, 11 pp.

Muthukrishnan, K., N. Meratnia and M. Lijding. 2005. FLAVOUR- Friendly Location-aware conference Aid with priVacy Observant architectURe. CTIT Technical report TR-CTIT-05-28. Enschede, The Netherlands: University of Twente CTIT, 16 pp.

Open Geospatial Consortium. 2007. Welcome to the OGC Website. <http://www.opengeospatial.org/> Accessed 23 October 2007.

Rijkswaterstaat. Undated. AHN: Actual Height model of the Netherlands. (Actueel Hoogtebestand Nederland.) <http://www.ahn.nl/english.php> Accessed 23 October 2007.

Sun Microsystems. 2007. Jini Network Technology. <http://www.sun.com/software/jini/> Accessed 23 October 2007.

SVGopen.org. Undated. SVG Open 2005 Conference and Exhibition - Home. <http://www.svgopen.org/2005/> Accessed 23 October 2007.

Urquhart, K., S. Miller and W. Cartwright. 2004. A user-centred research approach to designing useful geospatial representations for LBS. Pages 69–78 in *Proceedings of 2nd Symposium on Location Based Services and TeleCartography*, held in Vienna, Austria 28 – 29 January 2004. Edited by G. Gartner. Vienna: Institute of Cartography and Geo-Media Techniques.

Youssef, M. A. 2005. Location Determination Papers. <http://www.cs.umd.edu/~moustafa/location_papers.htm> Accessed 23 October 2007.