

## **BIOGRAPHICAL INFORMATION**

Peter Batty

Co-founder and Chief Technology Officer

Ten Sails

### **Specific Responsibilities**

Peter co-founded Ten Sails in 2002, a company focused on building businesses in the area of spatial and location related technologies. As Chief Technology Officer, he is involved in various areas including evaluation of new technologies and development of technology strategy.

### **Past Experience**

Peter has 19 years of experience in the geospatial information technology industry.

He was Vice President of Technology at GE Network Solutions / Smallworld, and was the first person to move to the US when Smallworld started its operations there. He has successfully led a number of major spatial software development projects.

Before Smallworld he worked in the GIS field with IBM.

He has been a member of the GeoWorld magazine Editorial Advisory Board since 1996, has spoken at many conferences around the world, and has received a record eight speaker awards from GITA. He has been involved with industry standards efforts including the Open GIS Consortium and IEC TC57 Working Group 14, and is a member of the GITA Board.

### **Educational information**

B.A. in Mathematics (First Class), Oxford University

M.Sc. in Computing, Oxford University

### **Professional memberships**

GITA

# **SENTIENT COMPUTING: FROM GRAPHICAL USER INTERFACE TO GEOGRAPHICAL USER INTERFACE**

Peter M. Batty  
Ten Sails  
8400 E Crescent Pkwy, #600  
Greenwood Village, CO 80111  
peter.batty@tensails.com

## **ABSTRACT**

A new generation of location tracking technologies with accuracy levels as good as a few inches, indoors or outdoors, enables a variety of entirely new spatial applications which track people and assets in real time, in many different application areas. These enable new ways for technology to interact with people, often called pervasive or sentient computing. This presentation will examine technologies and applications in this exciting area which is poised for rapid growth.

## **INTRODUCTION**

Imagine a world where intelligent environments, or “smart spaces”, know where everyone and everything is, to a high degree of accuracy, in real time. You could be in a shopping mall and walk up to one of the many screens located around the mall. The display instantly changes to show a small map which indicates where your children are in the mall, and you also see a live video of them – the system knows which of the many digital cameras covering the mall they are in front of, and automatically switches to that digital video feed. You didn’t even need to press anything to make all this happen – you just walked up to the screen and the system knew you were there and took the appropriate action. This is just one simple example of “sentient” or “pervasive” computing in action.

Hospitals have enormous potential for this type of technology. Electronic record keeping and workflow management does not currently have as high a penetration in hospitals as one might expect. A key reason for this is that staff are so busy caring for patients that anything that takes away time from this is not well received – if it is quicker to initial a piece of paper on a clipboard at the foot of a patient’s bed to record that they have been visited by a staff member, this may be seen as preferable to entering the information on a tablet PC or PDA. However, a sentient computing system which is monitoring the location of staff and patients in real time could record this interaction without any effort on behalf of the staff. This has many potential applications. For example, from the location of staff and patient in an operating theater the system could determine that an operation was taking longer than planned and automatically reschedule subsequent work. This sort of streamlined workflow can provide huge benefits, particularly in a busy emergency room environment. Another example is tracking elderly patients in homes. Not only can the system raise an alarm if they go where they shouldn’t, enabling a more friendly and open environment, but it can also make sure that they are exercising enough by monitoring their movements, and can alert care workers that they may not be taking

their medication if they do not go to the appropriate cupboard in their kitchen during the day.

There are applications for sentient computing in many other areas, including security, workplace, retail, military and emergency service training, theme parks, and observing animal behavior.

This paper examines the technologies that are being developed to enable this type of vision, and some of the challenges that need to be overcome. Managing large volumes of spatial data is central to these applications, but the applications are very different from traditional GIS. The paper considers the similarities and differences of these systems compared to traditional geospatial software.

## **CHALLENGES IN INDOOR LOCATION TRACKING**

The fundamental enabling technology for sentient computing systems or smart spaces is indoor location tracking. A system to do this is often known as a Local Positioning System, or LPS. The well established Global Positioning System (GPS) uses signals from satellites to locate objects on the earth's surface, to an accuracy of a few meters, or even a meter or less with appropriate correction. However, GPS does not work effectively indoors. Local Positioning Systems address this issue by using sensors that can locate objects within a more localized region.

There are several difficult challenges to be overcome in accurately locating objects in an indoor environment. RF (Radio Frequency) or other types of signal that may be used to determine an object's location may be reflected off walls, floors or ceilings, giving false readings that can result in inaccurate location determination. This is known as the multipath problem. A second problem is the line of sight problem: often in an indoor environment there may be obstacles such as furniture, partitions or people between a sensor and the device being tracked (which will be referred to as a tag in the remainder of this paper). Different technologies have different characteristics in this regard, which will be discussed below.

## **CHARACTERISTICS FOR EVALUATING LPS**

There are a variety of characteristics for evaluating location tracking technologies for a given application. The various technologies each have their strengths and weaknesses: no single technology is suitable for all applications. While there are many different applications for LPS with a wide range of characteristics, one broad distinction is that applications involving the tracking of people generally have distinctly different characteristics from those involving the tracking of assets.

The following are some of the key characteristics to consider in evaluating a location tracking technology:

**Accuracy:** some technologies provide accuracy as good as a few inches, others may be accurate to a few meters. When looking at accuracy figures, it is important to understand what these really mean. Sometimes you see things like "accuracy as good as 1m", or

“average accuracy of 1m” – neither of these are especially useful measures. A much more useful measure is the 95% confidence level measurement, which is the value which 95% of readings will be accurate to. For example, a system with an average error of 1m may well have a 90% confidence level of 2.5m. In other words, 1 in 10 readings will be worse than 2.5m (8 feet), which could easily place a person in the wrong room or wrong supermarket aisle, which could cause major problems in many applications. So accuracy needs some careful analysis – in many cases, accuracy may need to be higher than one might initially think.

**Cost:** this is an obvious factor but an important one. Important considerations with regard to cost are infrastructure cost (sensors, networking, etc) and tag cost. When tracking very large numbers of objects, such as boxes in a warehouse, or individual items on supermarket shelves, obviously a very low tag cost is essential. On the other hand, when tracking people, typically a higher tag cost can be justified in exchange for greater accuracy and faster update rates.

**Update rates:** some applications, especially those involving people, may require very dynamic real time update rates (multiple times a second in some cases). Other applications, especially those involving storage of assets, for example in a warehouse, may not need frequent updates at all.

**Tag size and form factor:** some applications will require very small and lightweight tags, others will be able to accommodate bulkier tags. Tag size may be related to other factors – for example using a watch battery rather than a AA battery will reduce tag size but also reduce battery life.

**Tag battery life:** this is an important consideration in real world applications. Asset tracking applications typically will require a very long battery life, or passive sensors that require no battery at all. Some applications may be able to tolerate more frequent battery replacement or recharging, though even when this is acceptable, the length of the battery life can have a significant impact on the ongoing maintenance costs of the system. Again, battery life is related to other factors – in addition to size, battery life will be affected by update rate, with higher update rates reducing battery life. Some systems have sophisticated power management, varying update rates intelligently to maximize battery life.

## LOCATION TRACKING TECHNOLOGIES

This section considers the major location tracking technologies that are now available, and briefly compares their characteristics. For completeness, global as well as local positioning technologies are included.

### GPS

GPS is now very much a mainstream technology. A search on amazon.com for GPS (in November 2004) returned over 1700 items, from fish finders to personal training devices. There are plenty of sub-\$200 devices which can provide 3m accuracy with WAAS. Even higher (1m) accuracy can be obtained using appropriate differential techniques.

## **Phone locating technologies**

Most new cell phones are now location enabled, either using GPS or other mechanisms such as E-OTD (Enhanced Observed Time Difference), which does not require special hardware in the handset, but uses a system based on triangulation from cell phone towers. In the United States, the move toward location-aware mobile phones has been driven by the E911 government mandate, which obliges phone companies to be able to locate people who make emergency calls. The original compliance deadlines have been moved back; the current mandate states (in overview) that most new mobile phones activated by the end of 2003 must be location aware, and by the end of 2005 the phone companies have to ensure that 95 percent of handsets in operation are location aware (full details are available at <http://www.fcc.gov/911/enhanced/>).

Location accuracy achieved by cell phones is generally less than that achieved by dedicated GPS units – the E911 mandate requires location accurate to 50 to 100 meters (150 to 300 feet) in most cases. This is not sufficient for accurately locating objects for many GIS applications, but it is good enough for many applications, like tracking work crews, for example. Location enabled phones also open up many opportunities for more consumer oriented applications, or “Location Based Services” (LBS).

## **Microwave**

Microwave signals can be used for local positioning system and work very well outdoors, or in large open buildings such as sports arenas, with accuracy as good as a few inches. Microwave is commonly used for sports tracking. However, microwave does not work well indoors.

## **Ultrasound and Laser**

These are entirely different technologies, but have two key features in common: both can provide excellent accuracy (to less than an inch), but both require a very dense network of receivers and so are very expensive and impractical for most applications. Laser systems are used for applications like controlling robots in high end manufacturing facilities, and ultrasound has been used for some successful research systems.

## **UWB**

UltraWideBand radio, or UWB, provides probably the best combination of accuracy and price for most indoor applications which require relatively high accuracy. Accuracy of six inches or less is typical, with relatively low infrastructure (compared to ultrasound and laser in particular). The nature of the UWB signal enables good handling of multipath.

## **WiFi**

WiFi location tracking has received a reasonable amount of press recently. It provides relatively coarse accuracy, using signal strength rather than triangulation to determine position. While vendors claim accuracy “as good as” 1m, typically the 90% confidence level is several meters. Since signal strength is used (rather than calculation based on triangulation or other means, which is the approach used by all the preceding methods), the whole site needs to be surveyed, and re-surveyed whenever furniture is moved around or the WiFi access points are changed. While there is a superficial attraction in being able

to use WiFi access points, which an organization may have anyway, it is actually necessary to have a signal from at least 3 access points and preferably as many as 6 access points at each location – so it is necessary to buy many times more infrastructure than would be required for standard WiFi network access. WiFi is attractive for applications involving tracking PCs or PDAs where only a very coarse degree of accuracy is required, but it is not accurate enough for many types of application, and suffers from large and expensive tags with poor battery life (for applications which do not require a PC or PDA).

## **RFID**

RFID is somewhat different to the other technologies mentioned above. RFID tags are very cheap, currently less than 50c in high volumes. The id on a tag can be read when it is within a short distance (typically a few feet) of a reader. So RFID does not continuously track the location of a tag like the other approaches; rather it records discrete locations whenever a tag comes within range of a reader. Therefore RFID is very suitable for applications like logistics or supply chain management, where many items pass through a relatively small number of locations like warehouses, depots or stores. Because of the low cost of tags, it is very suitable for certain types of asset tracking applications.

RFID can be combined with other tracking technologies. For example, a UWB tracking system might record the location in a warehouse where a fork lift truck left a pallet stacked with boxes, and at the same time an RFID reader on the truck records the id of all the boxes that were left at that location. Obviously this approach is not active – it assumes that nobody will subsequently move those boxes (without recording that move) – but it means that more expensive active tags are not required for each box.

## **TYPES OF LOCAL POSITIONING SYSTEM APPLICATION**

### **Discrete tracking**

As described above, RFID essentially enables discrete tracking of objects. In most supply chain applications, the readers are relatively sparse spatially, and the number of objects being tracked relatively large. While it could be useful to see locations of readers (and hence last known locations of objects) on a map, this is not necessarily a major requirement.

In some applications, a denser network of readers might be used, in which case the type of application that is enabled becomes more like a coarse grained continuous application, as discussed in the next section. For example, one could install RFID readers on each side of every door in a building, which would enable you to tell which room someone was in (a reader on each side of each door would be necessary to know whether someone had entered or left a room). In many cases the infrastructure cost for an approach like this might be more expensive than using a continuous tracking technology, though if very large numbers of objects were being tracked in this setting then it could be cost effective.

For further discussion of issues relating to this type of application, see the next section on continuous tracking.

## **Continuous tracking**

Continuous tracking applications track an object or person at all times (when within the area covered by the LPS). Depending on the nature of the application, the location of the tag may be transmitted wirelessly at periodic intervals, it might just be transmitted on request, or positions might even just be stored periodically in a local device of some sort, for analysis or reporting later.

An example of this type of application is analysis of space usage in an office. Many companies are interested in a “hot desk” approach to offices, in which people do not have a fixed desk, but just use any free desk whenever they are in the office. In this environment it is hard to tell how much office space is really being used – if analysis can be done to determine that less space is needed, then this can result in significant savings. Using a continuous tracking system to record locations of everyone in the office at a very fine-grained level can provide exactly this functionality.

Other applications in this situation include provision of basic location information – if people no longer have fixed desks, it is useful to be able to look at a map of the office to determine where a given person is, at any point in time. Of course this raises questions about privacy, which are discussed below.

Certain aspects of a system for these types of application are similar to traditional GIS. For example, the requirement to display multiple types of object on a map, with the ability to define flexible symbology, which may be dependent on properties of the object; the ability to display annotation, and to click on an object and display further information about it; the ability to search for an object based on a query, and go to it on the map (where is Peter?).

One key difference is that the data for this type of application is much more dynamic than in a traditional GIS application, potentially updating multiple times a second. In this environment the display should change automatically, without user intervention, as things change in the real world. Traditional GIS is designed to handle largely static data, and typically the map displayed on the screen will only be refreshed when the user requests it. Obviously current systems can be customized to provide some degree of dynamic display, but for high update rates a different architecture will be needed. Some existing software systems are designed to handle this type of dynamic environment, such as SCADA (Supervisory Control and Data Acquisition) systems in a utility, or air traffic control systems, but these are typically highly specialized and with limited spatial capabilities. Similar capabilities will be needed in more general purpose geospatial systems.

Another requirement is that systems will need to handle a much wider range of graphics than before, including increasing use of 3D visualization and multimedia content such as digital photos or video.

Both these trends will lead to the spatial industry taking much greater advantage of the capabilities of modern graphics cards than has previously been the case. Though we are starting to see a few exceptions, the graphical interaction a user currently gets with expensive GIS packages is far less dynamic than with a typical mass market video game.

These types of system will generate enormous quantities of spatio-temporal data, which requires a whole new set of tools for reporting and analysis. In the example mentioned above, we could have data about the location of each person in an office building, once a second for the past year. How do you can sense be made from this mass of data, to determine which parts of the building are over or under used, where there are bottlenecks in the usage of the building, etc? Spatio-temporal analysis has been a topic of research for some time and a few tools exist in this area, but the availability of these types of incredibly rich datasets will provide great impetus in this area. Display of historical data may also be a requirement in some types of real time operational systems, since it may be important to know where someone or something was previously, as well as where they are now.

There will also be a market for specialized systems for storing and querying these huge amounts of spatio-temporal data. These will be analogous to existing systems that archive other types of real time sensor data, but there are additional complications in being able to effectively do spatio-temporal queries against the data. The general principle of such systems is that instead of storing the location of an object every second, they will reduce the amount of data stored by only storing a new location and time when the location of the object changes by more than a specified tolerance.

### **Sentient computing applications**

A continuous location tracking application becomes a sentient computing application when it can initiate actions based on real time spatial analysis, without needing any specific initiation from the user. For example, in a hospital, when a doctor is stationary and within a specified distance of a patient for a given number of seconds, an event could be generated to indicate that the doctor visited the patient at that time – and similarly the system would automatically record when the visit ended. This is a simple example, but it adds an order of magnitude of complexity to the software that is required to manage the system, since many thousands of spatial interactions need to be analyzed every second.

Some types of interaction are easier to handle than others. For example, checking whether a visitor has entered an area of a building which only employees are authorized to be in is relatively easy – whenever a visitor location changes it just needs to be compared to a relatively small number of fixed areas. Interactions involving multiple mobile objects are more complex – like the example above, where we decide that a doctor is visiting a patient if they are within a certain distance of each other for a certain period of time. Other complex interactions to process include working out that a person is carrying an object. To take another hospital example, an alarm could be raised if a baby was carried out of a specified area by anyone other than a family member or an authorized member of staff.

Other examples of sentient computing may not be so obvious. In an office, if a person scans a document on a shared scanner, the file can be automatically saved in their home directory, since the system knows who is standing next to the scanner. Portable electronic devices like cameras or voice recorders can be easily shared when they are also tagged. Multiple people can take pictures or record audio notes, and again the system can tell which files should be sent to which people when they are uploaded, by cross-referencing

the time each file was created with the person who was holding the recording device at that time. There are many uses for cross-referencing the location of people with a large network of video cameras which covers a whole building or site. In a military training exercise, each soldier can receive a personalized video showing their movement throughout an exercise, even if hundreds of people were involved. A theme park could offer a family a personalized video or set of still photos showing each family member at interesting locations in the park.

Another less obvious example of how sentient computing can be used is to draw a box a few feet square in the corner of a hospital ward. If any piece of equipment (which has a tag) needs repairing, a staff member just places it in this square. This immediately triggers the system to record that this equipment needs repair and schedule someone to pick it up. Smart posters on a wall are another interesting idea – just by pressing a button on a tag while standing in front of a poster, a person can request more information or sign up for a service.

These examples just scratch the surface of what is possible with sentient computing – it not only enables applications which have an obvious requirement for tracking the location of people or objects, but it enables entirely new ways for people to interact with computers in a very intuitive way.

## **HOW SENTIENT COMPUTING WILL IMPACT GIS AS WE KNOW IT TODAY**

The types of sentient computing applications described in the previous section require high location accuracy and update rates. Initially these technical requirements can be best met by LPS, so applications will be focused on indoor locations or specific local areas. However, as both local and global location tracking technologies evolve, wireless networking matures, and tag prices drop, the same concepts will be able to be applied anywhere to any objects that are tagged.

This will radically change the way that we update spatial databases – in many cases they will update themselves automatically in real time. For example, when a utility worker mounts a transformer on a pole, its location (and all other relevant information) will just be immediately passed back to the master database, without the field worker having to take any action. There are various intermediate steps that could move us towards this goal. Current technologies would enable the field worker to record field work to a good degree of accuracy with simple actions like pressing a button, or saying “transformer installed” to a Bluetooth-enabled microphone mounted in their helmet, which would initiate a GPS reading to record their current location. This could also be combined with a portable RFID reader to record the identity of the equipment being installed. An LPS mounted on the truck could provide a more accurate location than could be provided by GPS alone (the truck could get a more accurate GPS reading than the worker, through being stationary for a longer time and having a better antenna).

## **PRIVACY**

The ability to track the location of people to a high degree of accuracy in real time has many potential benefits, as outlined in this paper. This obviously raises potential concerns over people's privacy, and it is important that appropriate policies are put into place to address these issues. In some ways, the issues raised by location-aware cell phones are potentially greatest, since they track people all the time, whereas LPS just tracks people within a defined area such as a hospital or office.

There are a number of cases where the privacy issue can be addressed in a straightforward way:

- The application can run with anonymized data – for example, a basic space utilization analysis in an office just needs to know when people were in an area, it does not need to know the identities of the people. Individuals will still need to be reassured that their identity genuinely is not included in the recorded data, but technically this is straightforward
- In some cases, the benefits are sufficiently compelling that people are not likely to have significant objections – for example, if the technology can save lives in a hospital environment, it is likely to be accepted there. The technology can provide significant safety benefits in other cases, for example for prison guards.
- In other cases, location tracking can be made voluntary and individuals can be allowed to turn it on and off as they please

The harder cases are those where individuals are identified, and where it is essential that everyone is tracked in order for the application to work effectively. In these cases, policies need to be developed to make sure that location data is not misused.

It is also worth noting that our privacy is being eroded by technology in general, not just by location tracking. More and more data is recorded about us and our actions – what we buy, where we buy it, when and where we take money from an ATM, when our car goes through a toll plaza, etc. While it is important to have policies to manage this, it is hard to see how the trend will be reversed. Scott McNealy of Sun has said “You have no privacy anyway, get over it!”

## **CONCLUSION**

This paper has described new types of location tracking technologies with a variety of characteristics – some can provide real time accuracy as good as a few inches, others enable assets to be tracked with a coarser degree of accuracy very cheaply indeed. These technologies enable a variety of entirely new spatial applications which track people and assets in real time, in many different areas. Some of these applications have obvious location tracking elements; others are less obvious but provide new ways for technology to interact with people, often called pervasive or sentient computing. This moves us from a graphical user interface to a geographical user interface, in which the location of people and objects drives the behavior of a computer system. Ultimately we will have the ability to know where everyone and everything is all the time, which has enormous potential, but raises some difficult issues with regard to privacy.