

Sensor Aggregation and Integration in Healthcare Location Based Services

Lorcan Coyle, Steve Neely, Paddy Nixon and Aaron Quigley
Systems Research Group
School of Computer Science & Informatics
UCD Dublin
Belfield, Dublin 4
Ireland
Email: first.name@ucd.ie

Abstract—Complex and dynamic working environments such as health care facilities consist of staff, patients and equipment constantly moving in response to changing medical requirements. Knowing the current location of people and equipment is essential for the smooth running of a facility, yet creating a global view through tracking is a challenging task. It is clear that many common hospital situations can be improved with real-time access to the various actors' location information. One of the main problems with implementing such services is that current location based applications tend to be proprietary and the data generated closed. The realisation of ubiquitous location based services demands the exploration of hybrid models and methods that can utilise existing and subsequent infrastructures in novel and complimentary ways. We describe a number of hospital scenarios that use location-based services and make available all the location data gathered. We propose that by aggregating location data by a range of acquisition methods it is possible to improve the performance of location applications and readily adapt to the introduction of new location detection technologies.

I. INTRODUCTION

Hospital environments are complex and dynamic with multiple staff, patients and equipment constantly moving in response to medical requirements. The continual change in the environment state is a result of a plethora of disciplines (doctors, nurses, administrative staff, porters, support staff etc.) all working together toward the common goal of providing medical support for patients. From the moment a patient is first admitted to a facility they are moved through wards and departments to the appropriate professional treating their case. At the same time, numerous staff are on the move around the hospital during their shift.

Improving efficiency is a commonly cited operational goal of modern hospitals. Efficiency can be measured by throughput of patients, response time, patient to staff ratio or time taken to complete individual tasks. Technology can serve to improve efficiency in many cases such as automating administrative tasks and calculating streamlined work flow procedures.

Efficiency is also improved by being able to locate the most suitable staff members along with the required equipment for a given patient in a timely manner. Improved communications and administrative support afford this. Automated people- and asset-tracking can further reduce response times. Although key

personnel can be located through ad-hoc communication with pagers or mobile phones, locating patients and equipment can be more difficult. Location tracking systems have a natural place in such hospital environments. Emerging technologies in real-time tracking make location-based applications in this area an affordable, practical reality [1].

We outline three workplace use cases where location based services (LBS) can improve the efficiency of hospital tasks. Commercial LBS solutions are typically offered as part of a specific end-to-end package where only application layer data are exposed. We propose an alternate approach in which we describe an open architecture with an accessible sensor layer and a separate, loosely coupled, application layer.

Since the application layers of each of these use cases all consume location data, the ideal case is to pool the sensor layer data and make it available to each of the applications either in its raw form or in a more processed, aggregated form. This will make it possible for developers to improve performance, easily allow the introduction or upgrade of location-detection systems and assist in the development of new location-based services.

II. BACKGROUND ON TECHNOLOGY

Location detection systems facilitate the discovery of a user or object's location within a given physical space. The systems manifest themselves in a number of varieties. Some require users to carry identifiable tags detected by fixed sensors that are installed and calibrated; there exist systems that simply track movement of objects in an anonymous fashion; whilst others employ existing fixed infrastructure, such as WiFi access points or Bluetooth devices, to determine location. These location detection systems are commonly used in context-aware applications, where the roles or current activity of the user influence the behaviour of the system. For example a doctor performing rounds in a hospital may be assisted by automatically sending the latest patient reports to his hand-held computer as he arrives at their room.

Reusing existing infrastructure is a cost-effective method for determining location. It does not require the installation and calibration of new sensor systems. Instead we utilise equipment in well-known locations to act as beacons. The

determination of a GSM mast's ID and the trilateration of signals from GSM mobile telephone masts can do this. The accuracy of GSM trackers depends strongly on how many masts are in range of the receiver. Accuracy can range from 50m in densely populated areas, such as cities, to kilometre accuracy in the countryside. A current example of this is the London ambulance service which uses mobile phone signals to locate callers when responding to emergencies.

Place Lab [2], Ekahau [3], Wherenet [4], Microsoft Live [5] and Aer Scout [6] are all examples of location detection systems that use wireless access points (APs) as beacons. Discoverable APs afford such a service as they provide a unique address to devices when queried. WiFi based positioning systems need to undergo a bootstrapping phase in which the locations of the APs are mapped out, stored in a database and shared through sites such as WiGLE [7]. The relative signal strengths from APs can later be interpreted and analysed with this database to calculate the location of a device. The granularity of these systems again depends on how densely populated an area is with APs. It is difficult therefore to make absolute statements about accuracies of these systems. Place Lab, for example, has been reported to provide a median accuracy of 15-20m [2].

To provide planetary-scale location determination, the Global Positioning System (GPS) or the Galileo positioning system (which is currently under development) are most effective. These systems employ a constellation of satellites orbiting the planet tracked by specialised receivers to determine location. The cost of placing and maintaining satellite-based systems is extremely high but the cost for receivers is relatively low. These receivers, augmented with extra features such as Dead Reckoning technology systems typically provide 1-3m accuracy. Basic GPS systems require a direct line-of-sight with the sky so suffer indoors or in cities where shadows of tall buildings create urban canyons in which satellites cannot be tracked or where severe multipath effects are felt.

On a smaller scale systems may utilise Bluetooth [8], RFID, Infra-red [9] and ultrasound technologies, such as Crickets [10] and Bats [11]. These would typically be deployed at room-level and provide sub-metre accuracy.

Recent advances have seen the emergence of location detection technologies based on ultrawideband [12] [13]. These require the installation and precise calibration of fixed sensor infrastructure to detect tags within a building. The benefits of using ultrawideband are that the signals are less affected by attenuation and multipath distortion as compared with GSM, WiFi, or Bluetooth. Ultrawideband technologies typically provide 15-50cm location information at a high cost from thousands to tens of thousands of dollars for the equipment to cover $1000m^3$ of room space.

Computer vision has become a standard for tracking objects ranging from coarse grained crowd control and head count determination to fine grained full body motion capture [14]. Using simple off-the-shelf cameras with specially patterned or coloured markers the system can track objects at sub-centimetre accuracies. Such system can be an attractive choice

for a range of specialised tasks but require line-of-sight to work.

An alternative approach to determining location is to ask software that has no connection to physical sensor infrastructure but operates using hearsay. The term hearsay is used as the data provided are not derived directly, but are told to the system indirectly by a human. One example of this is by looking at the location information contained within calendar entries, e.g. meeting locations attached to times. While these data are unlikely to have the accuracy of dedicated location based systems, they may be used as a last resort or complimentary source.

Location detection systems vary in accuracy, range, area coverage and cost. These can be from sub-centimetre to kilometre accuracy; room, building, campus, city or planetary coverage; and by reusing existing infrastructure can range from free to multi-million dollar costs to implement.

III. USE CASES FOR LOCATION-BASED SYSTEMS

To demonstrate the utility of location-based systems in a hospital setting, we describe a number of use cases where a LBS can make the running of the facility more efficient. The first of these scenarios uses passive observations (proximation [8]) of doctor's and nurses location as they visit patients on their rounds [15]. As doctors or nurses visit patients, their own LBS monitors their position, and a journaling application appends meta-data to this location information to drive documentation of patient updates. This scenario does not require real-time updates nor for the LBS to expose any data beyond the doctors or nurses own system. Further, the collocation of doctor and patient coupled with a procedural trigger is used to select the appropriate logging mechanism.

The second scenario is the tracking of assets. Tag-based systems have proven successful for asset-tracking in many areas, e.g., retailing [16]. In a hospital setting, tags may be attached to high value equipment, e.g. IV pumps, ventilators, patient transportation equipment, medical supplies, and even blood [17]. Tag readers can be fitted to portable computers or embedded in strategic locations, e.g. door frames [18]. The difficulty with these systems is that the installation of readers comes with a significant cost and a maintenance cycle.

The final scenario is based on the current work practice of doctors being paged using a dedicated infrastructure, which requires each doctor to wear their pager at all times. Patients or visitors are not part of this system and are often contacted by public address. Public address systems typically broadcast across a larger area than is necessary, contributing to the overall noise pollution. By segregating the public address system and tailoring the broadcast to speakers in the areas where the targeted people are expected to be, such intrusion can be reduced.

IV. LOCATION-BASED SYSTEM GRANULARITY AND COST

Location detection can be viewed as a boundary detection problem [18] rather than one of accuracy. To find patients, staff, and assets, the level of accuracy need not necessarily be

high; it may be enough to know which room they are in. This means that RFID, Bluetooth, UWB and WiFi based solutions would all be suitable.

The cost of deploying a location detection system can vary widely. UWB offerings require specialised sensors installed and calibrated in every room and corridor. Deploying UWB tracking in a large building can become prohibitively expensive due to this. Bluetooth and RFID installation is less complex and devices are significantly cheaper. WiFi based solutions may reuse existing communications infrastructure which negates the requirement for purchasing new beacons. In areas where WiFi is not available a switch to a different detection system or combination of them would be ideal.

Given the diversity of technologies already in place in hospital environments coupled with restrictions placed by financial and technical considerations, a combination of best fit location detection systems is desirable. Being able to combine inputs from WiFi, UWB, RFID and Bluetooth trackers to give more accurate results — or even any results, when other technologies are not available — would afford the location based system to run more effectively.

V. AGGREGATING HETEROGENEOUS SENSOR DATA

A typical LBS offers an end-to-end solution. Sensors give readings, the application interprets these readings and behaves appropriately. The data generated by location sensors in location detection applications are often tied to a proprietary format or otherwise are not exposed to other applications easily.

In an organisation with many applications there is usually no way of translating and delivering data from the sensors of one application to the outputs of the other. Middleware can be used to overcome this problem. By decoupling sensors from applications and using a middleware to mediate between the two, it is possible to free the data for use in other application areas [19]. By aggregating sensors from different location based services it is possible to improve the overall accuracy of the system. However, care needs to be taken to ensure that the individual accuracies of the systems are taken into account.

A. Construct

We have designed the Construct platform [20] as a system for integrating noisy data sources in a clean, dynamic, flexible and semantically well-founded manner. Construct provides applications with a uniform view of information regardless of how that information is derived. Construct supports extensive sensor fusion and aggregation within the platform to be shared between applications. All data are modelled using the Resource Description Framework (RDF) [21], which provides a standardised way in which to model contextual information and properties. The use of the RDF standard for modelling context provides a common level of abstraction with which to represent information generated within a smart-space [22]. This allows all data to be managed, manipulated, and disseminated independently of the technologies used to

acquire it. This makes it easier for developers to handle heterogeneous location sensors and devices.

We are complementing this work with the development of an uncertainty framework that performs aggregation of data and deals with inconsistencies that may arise when the same type of data is produced from different sensors (e.g. location sensors providing data that says a person is in two places at once) [23].

B. Aggregating Readings from Multiple Sensors

Representing and operating on uncertainty is recognised as key in any system using many sensors [25]. *The Location Stack* [24], [25] uses Bayesian filtering techniques to manage such data. It uses raw sensor data, models of the physical characteristics of the sensor and the environment, as well as using previous estimates of location to estimate the location of an agent at the current time. Our approach is similar to this; we model the inaccuracy of location data using three parameters: *precision*, *decay*, and *confidence*.

Precision relates to the physical accuracy of a sensor reading. It must be pre-determined by the sensor itself — e.g. Bluetooth sensors tend to be far less accurate than UWB (granularity of metres versus centimetres for UWB) — and sensors must provide an estimation of their precision along with each reading. The decay parameter attempts to capture the reduction in belief in a sensor reading as time passes since it was generated. If readings are generated regularly then the most up-to-date reading is likely to be the most accurate. The decay rate is directly related to the frequency of readings — UWB sensor readings have a higher decay rate than RFID as they occur much more frequently.

The confidence measure is used to counter the problems that occur when sensors behave incorrectly. Confidence reflects the middleware's belief that a sensor is accurate. This is most useful when there are redundant sensors; if one sensor is consistently generating readings different from the others, its accuracy should be questioned. This can happen if a sensor is broken, incorrectly calibrated or behaving maliciously. Each sensor's confidence parameter is altered over time as its readings diverge from the consensus view. In effect, the middleware corrects or *learns* the accuracy of each sensor by altering the confidence parameter. The learning of confidence is dependent on redundant sensors; where there is only one source of location data in an area the system must accept that sensor's value as it is the only one available. The location based system may test the value of the confidence measure for each sensor and recommend that a sensor be replaced/tested if its confidence value drops below a critical level.

If a single reading is required by an application, e.g. *what is the location of hospital cart X?*, all the sensor readings for that cart are aggregated using the precision, decay and confidence parameters to generate a single estimation of the cart's current position [23]. This is done by giving each sensor reading a weight depending on the strength of each of the accuracy parameters. In this way, old readings, readings from

imprecise sensors and readings from sensors that the system does not trust/believe will be given low weights and vice versa.

VI. CONCLUSION AND FUTURE WORK

The very nature of healthcare facilities, where there is a lot of movement and activity, results in a complex and dynamic environment. The ability to track people and equipment facilitates the development of location based services which can improve the efficiency of such environments; there are a host of independent technologies that do this. Since all hardware and software location based systems have a degree of *inaccuracy* our goal is to reduce this error via aggregation and fusion.

Our work is ongoing; currently we are collecting a corpus of context and location data from different sources, including Ubisense and hearsay sources. Our goal is to develop and evaluate our aggregation algorithms against these datasets. Next we will generate realistic data to test and demonstrate the problems outlined in our hospital scenarios.

Dedicated LBS are typically targeted at a single application which limits their general purpose utility for new services or to help improve the coverage, accuracy or response of complimentary services. A significant advantage of using a sensor middleware is that patients may have additional sensors attached to add extra value without redesigning the whole system. By adding temperature sensors to patients' identity tags it would be possible to collect the temperature of many patients and assess it centrally. This has been shown to cut down the likelihood of staff infection, as well as ensuring that fever outbreaks are spotted at an early stage [26].

In general, location based services in healthcare, although in their infancy, may have great potential to reduce operating costs and improve medical services and patient safety.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the support of Science Foundation Ireland (under the grant "Secure and Predictable Pervasive Computing" and Enterprise Ireland under the grant "Pervasive Visualisation").

REFERENCES

- [1] A. Quigley and D. West, "Proximation: Location-Awareness Through Sensed Proximity and GSM Estimation," in *Proceedings of the International Workshop on Location- and Context-Awareness, (LoCA 2005)*, vol. 3479, May 2005, pp. 363–376.
- [2] A. LaMarca, Y. Chawathe, S. Consolvo, J. Hightower, I. Smith, J. Scott, T. Sohn, J. Howard, J. Hughes, F. Potter, J. Tabert, P. Powledge, G. Borriello, and B. Schilit, "Place lab: Device positioning using radio beacons in the wild," in *Proceedings of PERVASIVE 2005, Third International Conference on Pervasive Computing*, Munich, Germany, 2005.
- [3] Ekahau, Inc, "Ekahau RTLS," www.ekahau.com.
- [4] WhereNet, Inc, "WhereNet RTLS," www.wherenet.com.
- [5] Microsoft, Corporation, "Microsoft Live Local," local.live.com.
- [6] AeroScout, Inc, "Aeroscout RTLS," www.aeroscout.com.
- [7] WiGLE, "Wireless Geographical Logging Engine," www.wigle.net.
- [8] D. C. C. O. Aaron Quigley, Belinda Ward and B. Kummerfeld, "BlueStar, a Privacy Centric Location Aware System," in *IEEE Position, Location and Navigation Symposium 2004 (PLANS 2004)*, Monterey CA, USA, April 2004.
- [9] R. Want, A. Hopper, V. Falcao, and J. Gibbons, "The Active Badge Location System," Olivetti Research Ltd., ORL, 24a Trumpington Street, Cambridge CB2 1QA, Tech. Rep. 92.1, 1992.
- [10] A. Smith, H. Balakrishnan, M. Goraczko, and N. B. Priyantha, "Tracking moving devices with the cricket location system," in *2nd International Conference on Mobile Systems, Applications and Services (Mobisys 2004)*, Boston, MA, June 2004.
- [11] A. Ward, A. Jones, and A. Hopper, "A New Location Technique for the Active Office," *IEEE Personnel Communications*, vol. 4, no. 5, pp. 42–47, October 1997.
- [12] Ubisense, Ltd, "Ubisense RTLS," www.ubisense.net.
- [13] Parco Merged Media Corporation, "Parco PCTS," www.parcowireless.com.
- [14] B. Bodenheimer, C. Rose, S. Rosenthal, and J. Pella., "The process of motion capture: Dealing with the data.," in *Eurographics Workshop on Computer Animation and Simulation*, 2006.
- [15] A. Holzinger, K. Schwabberger, and M. Weitlaner, "Ubiquitous computing for hospital applications: Rfid-applications to enable research in real-life environments." in *COMPSAC (2)*. IEEE Computer Society, 2005, pp. 19–20.
- [16] M. Roberti, 2003, on CIO Insight web site, Analysis: RFID - Wal-Mart's Network Effect, <http://www.cioinsight.com/article2/0,1397,1455103,00.asp>, September 15, 2003.
- [17] intel, "Using rfid technologies to reduce blood transfusion errors," 2005, white Paper by Intel Corporation, Autentica, Cisco Systems, and San Raffaele Hospital.
- [18] M. Crawford, F. Fairmichael, K. O'Riordan, A. Quigley, M. Roche, and S. Tapley, "RFID Enabled Awareness of Participants Context in eMeetings," in *Pervasive 2006 Workshop Proceedings (ISBN 978-3-00-018411-6), Real-World Experiences with RFID and Sensor Networks, May 7 Dublin Ireland, 2006*, pp. 355–360.
- [19] L. Coyle, S. Neely, G. Rey, G. Stevenson, M. Sullivan, S. Dobson, and P. Nixon, "Sensor fusion-based middleware for assisted living," in *Proceedings of ICOST'06*, 2006, to appear.
- [20] G. Stevenson, L. Coyle, S. Neely, S. Dobson, and P. Nixon, "Construct — a decentralised context infrastructure for ubiquitous computing environments," in *IT&T Annual Conference, Cork Institute of Technology, Ireland, 2005*.
- [21] W3C, "Resource Description Framework (RDF)," <http://www.w3.org/RDF/>.
- [22] A. K. Clear, S. Knox, J. Ye, L. Coyle, S. Dobson, and P. Nixon, "Integrating multiple contexts and ontologies in a pervasive computing framework," in *Contexts and Ontologies: Theory, Practice and Applications*, Riva Del Garda, Italy, August 2006, to appear.
- [23] S. Dobson, L. Coyle, and P. Nixon, "Hybridising events and knowledge as a basis for building autonomic systems," in *Journal of Trusted and Autonomic Computing. September 2006*, 2006, to appear.
- [24] J. Hightower, D. Fox, and G. Borriello, "The location stack," University of Washington, Department of Computer Science and Engineering, Seattle, WA, UW CSE 03-07-01, July 2003.
- [25] D. Fox, J. Hightower, L. Liao, D. Schulz, and G. Borriello, "Bayesian filtering for location estimation," *IEEE Pervasive Computing*, vol. 2, no. 3, pp. 24–33, July-September 2003.
- [26] S.-W. Wang, W.-H. Chen, C.-S. Ong, L. Liu, and Y.-W. Chuang, "Rfid application in hospitals: A case study on a demonstration rfid project in a taiwan hospital." in *HICSS*. IEEE Computer Society, 2006.