Opportunistic Radio SLAM for Indoor Navigation using Smartphone Sensors

R. M. Faragher, C. Sarno, and M. Newman.
Sensor Systems, Advanced Technology Centre,
BAE Systems
Chelmsford, UK
Ramsey.faragher@baesystems.com

Abstract — **This paper provides the experimental results of a system utilising only the sensors available on a smartphone to provide an indoor positioning system that does not require any prior knowledge of floor plans, transmitter locations, radio signal strength databases, etc. The system utilises a Distributed Particle Filter Simultaneous Localisation and Mapping (DPSLAM) method to provide constraints on the drift of a simple hip-mounted Inertial Measurement Unit (IMU) integrated into the smartphone and providing the core information on the movement of the user. This system was developed during a project investigating methods of providing relative positioning systems to a team operating for extended periods without GPS. The paper concentrates on the DPSLAM positioning technique suitable for use by an individual with no prior knowledge of the area of operation before deployment. As with all SLAM systems, the user is simply required to revisit locations periodically to enable IMU drifts to be observed and corrected.**

Opportunistic radio positioning, SLAM, indoor navigation

I. Introduction (Heading 1)

Over recent years there has been increasing interest in ubiquitous positioning, or the ability to determine a location in any environment, outdoors and indoors. We have all become used to the availability and performance of Global Navigation Satellite Systems (GNSS) for accurate outdoor radio positioning with a reasonable degree of reliability and availability. However radio indoor positioning is more challenging since GNSS signals do not penetrate buildings well, and indoor positioning therefore relies typically on local infrastructure and other support to aid the user. Indoor radio positioning is available today to the general public in conurbations via WiFi and cellular measurements, by exploiting a database of signal strength fingerprints managed and provided by a third party provider such as Skyhook [1]. The user can access this database via a cellular or WiFi data connection. These systems therefore have two clear constraints: the area must already have been surveyed, and the user must have a data connection available to them.

An ideal system would not rely on these constraints, but would develop its own database during operation. Such a system is described and demonstrated here. The benefits of this system are significant - it can provide situational awareness and asset tracking in new and unknown environments for the military, emergency services, lone workers, security personnel and autonomous vehicles. This method does not require a data link to function, nor any prior surveying of the radio environment, nor any other prior knowledge such as a floor plan or map. The system can also be used however to quickly and easily generate maps of the radio environment or floor plans, which can be beneficial for organisations wishing to provide positioning services to the public using a simpler positioning method - i.e. this method can be used to rapidly survey an area and generate a signal fingerprint database for others users to exploit.

II. Indoor Positioning

A. GNSS challenges

The problems with GNSS availability indoors are well documented. The weak signals cannot easily penetrate building materials, especially not through multiple floors. While high sensitivity receivers exist [2] that can provide indoor signal tracking with degraded positioning performance, they have not provided a viable solution to the indoor positioning problem.

B. ZUPTS benefits and challenges

An existing indoor positioning technique that does not rely on any infrastructure or prior knowledge is the Zero Velocity Updates (ZUPTS) method. In this method, an inertial measurement unit (IMU) is attached to the foot of a user and a *strap-down* IMU solution [3] tracks the movement of the user as they walk. A typical strap-down solution using low-cost and highly-portable IMU sensors would normally suffer rapid degradation in positioning performance with no external assistance from GNSS or other sensors. The ZUPTS method can however exploit a particular feature of pedestrian motion to constrain inertial drift. Each foot is regularly static during normal walking motion during the periods when the users exploit the friction between their foot and the ground to propel their body. Since an IMU mounted to the foot must also be known to be static during this short period, the IMU accelerometer and gyroscope biases are observable during this short period with every step. The regular observations of the IMU biases permit much more accurate inertial navigation than would be possible if these biases were not regularly re-estimated. The accumulation of error associated with the ZUPTS location estimate is therefore reduced, but not

The work presented here was funded by the UK Ministry of Defence CDE (Centre for Defence Enterprise)

removed completely without further constraints. Extended operations indoors or in GNSS-denied locations will still result in a user's positioning estimate degrading with time. Mounting a device specifically to a foot may also not be an attractive solution for some applications, e.g. military personnel, casual consumer users navigating a shopping centre, etc.

C. Deployed beacons

Indoor navigation and tracking can be provided by deploying dedicated tracking beacons in the area of operation [4]. The benefit of this approach is the ability to exploit positioning measurements from a well-known and carefully designed system. The drawbacks are the issues of signal penetration through multiple walls when systems are constrained to operating within legal broadcast limits and the practicalities of rapid beacon deployment with good signal geometry in a new area of operation. Typically any system exploiting the license-free Industrial, Scientific and Medical (ISM) bands are limited by the maximum permitted broadcast power, resulting in a system with a maximum usable distance of a few tens of metres if there are multiple interior walls or other dense objects between the user and the beacons. This problem can be reduced by using low frequency signals [5] or by only operating in fixed environments with permanent indoor positioning needs, such as airports or warehouses, but is a significant problem for a system that must be capable of rapid deployment in new unknown environments (e.g. for military operations, rescue operations conducted by emergency services, etc).

D. Opportunistic radio positioning benefits and challenges

An alternative method of indoor radio navigation exploits pervasive opportunistic radio signals such as television, commercial radio and cellular broadcasts. These signals are typically received at much higher signal strengths than GNSS signals and so are capable of penetrating deeply into buildings. Much work has already been performed by various authors and companies in the field of outdoor opportunistic radio positioning exploiting standard radio positioning methods employing timing measurements to infer range between a given transmitter and the receiver or to assist GNSS signal acquisition [6 - 8]. However, the signal environment is highly complicated indoors, with rapid fading variations and highly-variable multipath interference corrupting these simple, traditional positioning methods [9 - 10] The only feasible method of accurate opportunistic radio positioning in difficult signal environments is signal fingerprinting, where the pattern of signal strength measurements gathered at a particular location is assumed to be repeatable and unique [11]. This method is currently provided by pre-mapped databases of signal fingerprints which users access via a network connection. An obvious and desirable extension to this concept is the automatic generation of this database as a user explores a new, unknown area. This can be achieved by developing a Simultaneous Localisation and Mapping (SLAM) technique, and such a method is described and demonstrated in this paper.

III. Radio Signal Strength Mapping

First it is important to test the hypothesis that radio signal strength maps in indoor environments exhibit high spatial variation, but low temporal variation (i.e. each map is complex, but does not vary significantly over time). To do this we use a Gaussian Processes regression scheme.

A. Gaussian Processes

A thorough discussion of Gaussian Processes (GP) is available in [12] and a discussion of its use for generating radio signal strength heat maps is given by Ferris [13]. The Gaussian Processes technique is a well-known multi-dimensional regression method that takes a set of training data and user-defined Kernels to generate multidimensional Gaussian mixture models for the states of interest.

The BAE Systems Advanced Technology Centre Research Facility provided the indoor environment for this study. This two storey building is roughly 100 metres by 50 metres in dimension, with a dense structure of multiple rooms, computers, servers, laboratories, and other objects. The building is located on the outskirts of a large town, and so enjoys a good coverage of opportunistic radio signals. Signal strength maps generated using training data and Gaussian Processes methods are given below in Figures 1, 2 and 3 for VHF signals (FM public broadcast radio), cellular signals (GSM 900) and WiFi (2.4 GHz) signals. The WiFi maps are shown as both GP mean value maps, and GP variance maps, demonstrating the ability of the Gaussian Processes method to not only provide a prediction of the estimated signal strength at an un-surveyed location, but to also provide an estimate of the error associated with the prediction.

The aim of this set of measurements was to determine the complexity of these signal strength maps, and to determine their variation with time. Due to the range in opportunistic transmitter locations, frequencies and transmit powers there is a stark difference between maps. The most significant contributing factor to this is likely to be the variety of transmitter locations, resulting in signals entering the building from different directions.

A simple experiment provided a useful test of the validity of Gaussian Processes for generating signal strength maps from training data while also testing the temporal variation in the maps. A set of GP signal strength maps was generated on a given day and then a week later the building was surveyed again. The new measurements were made at arbitrary locations within the building within the same regions (i.e. the same rooms and corridors) but no attempt was made to record new measurements at the exact old survey locations. The predictions extracted from the old GP maps at these new survey locations were then compared to the new survey data. The results are shown in Figure 4. It was determined that the new measurements typically agreed with the predicted values to within a few dBm, and within the error estimate of the Gaussian Process prediction. Large changes to the structure of the building or its contents are of course expected to cause more significant variations, and so ideally signal strength maps generated from previous visits to an area should only be

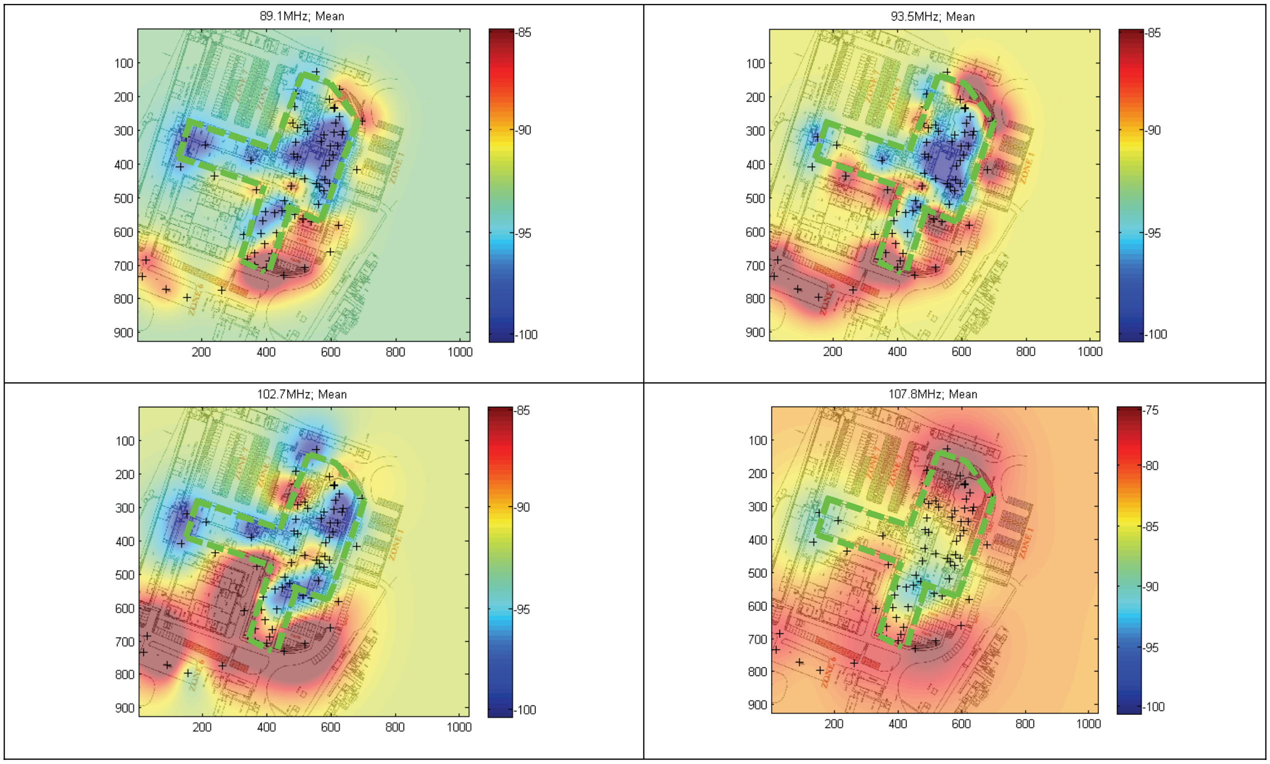

Figure 1: Gaussian Process mean signal strength maps for a number of VHF frequencies. The black crosses show the locations of the training data measurements used to generate the maps. The green dashed line marks the edge of the main building. The units of the image axes are pixel number. The colour bar scale is dBm.

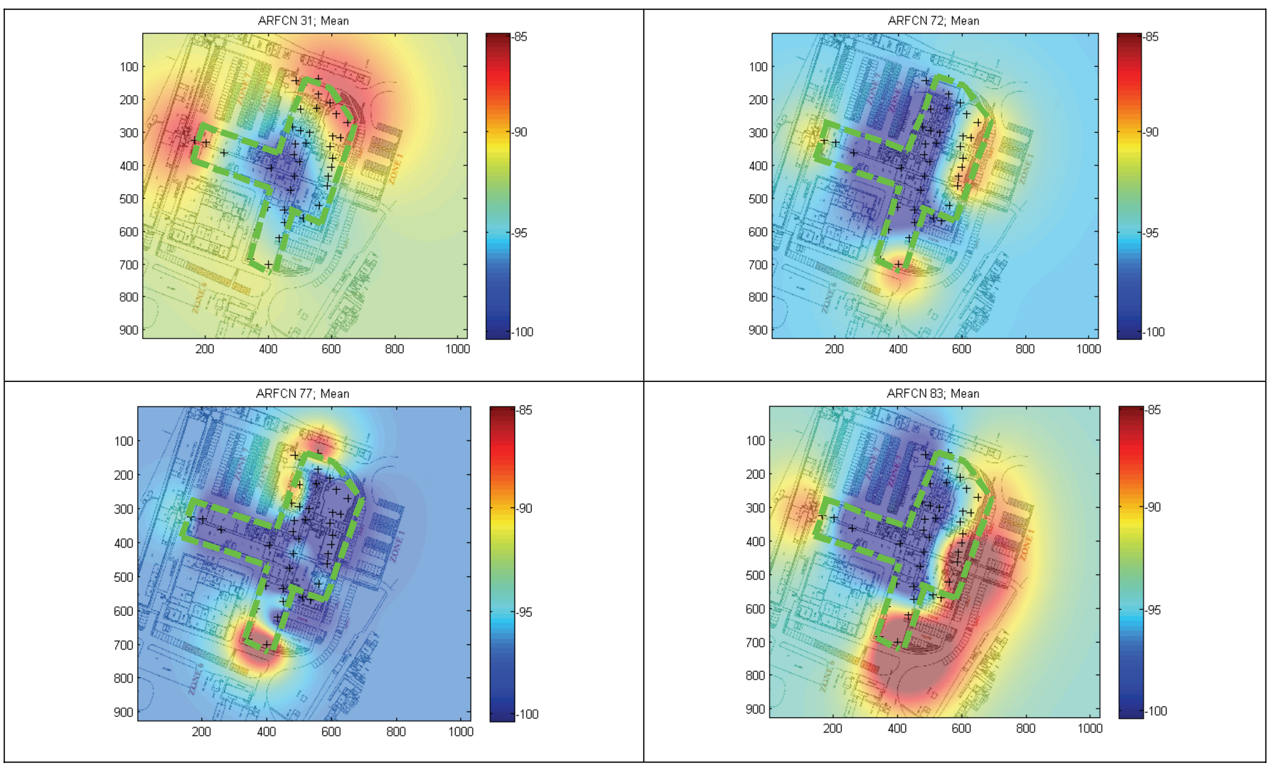

Figure 2: Gaussian Process mean signal strength maps for a number of GSM 900 frequencies, given by Absolute Radio Channel Frequency Number (ARFCN). The black crosses show the locations of the training data measurements used to generate these maps. The units of the image axes are pixel number, and the colour bar scale is dBm.

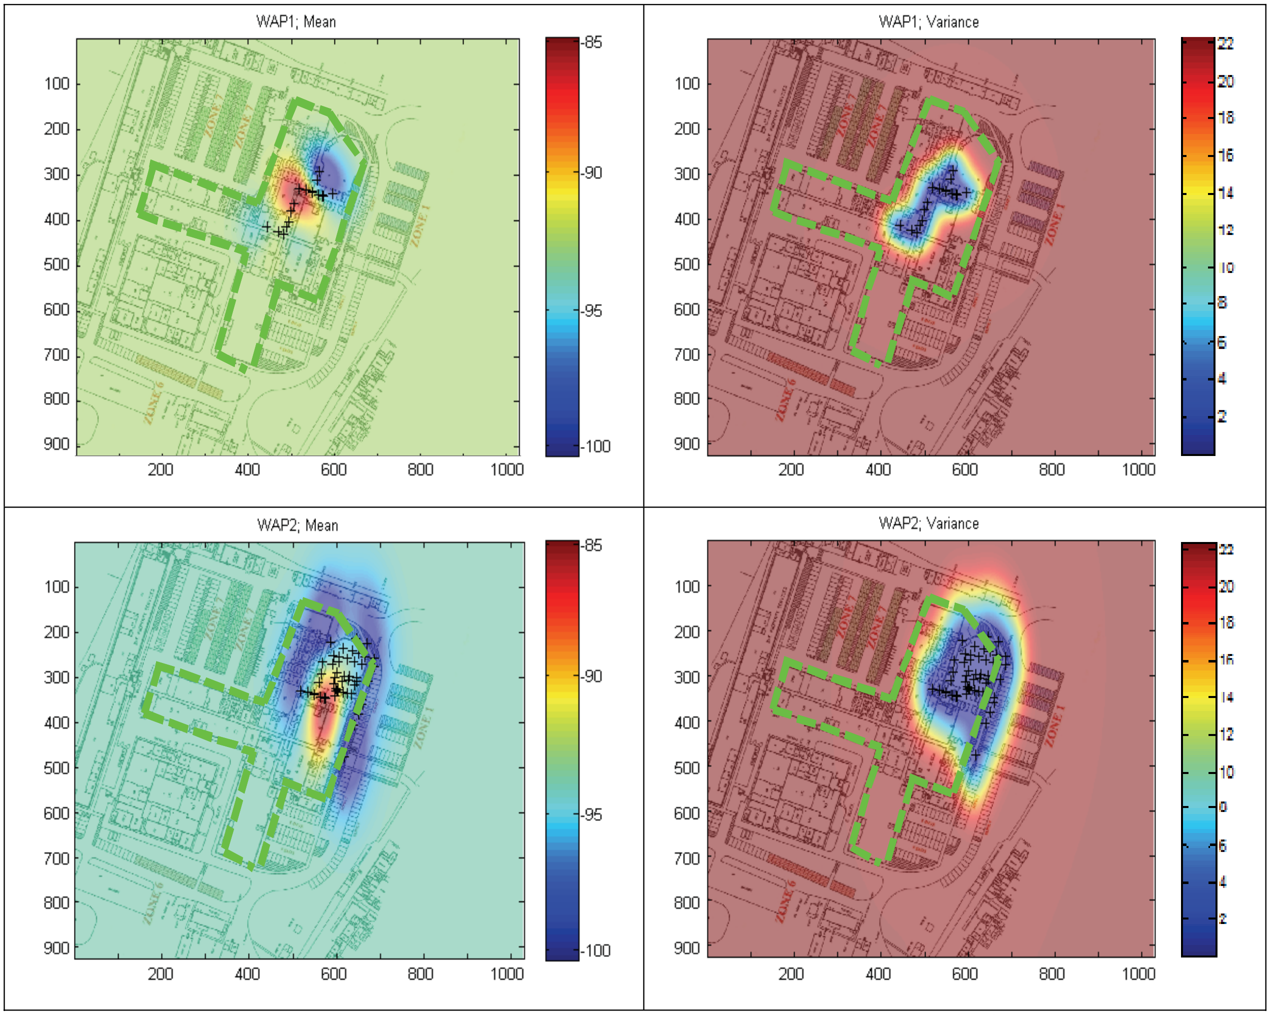

Figure 3: Gaussian Process mean signal strength (left) and variance (right) maps for two WiFI routers in the test area. The black crosses show the locations of the training data measurements used to generate these maps. The green dotted outline shows the main building outer wall. The units of the image axes are pixel number, and the colour bar scales are in dBm.

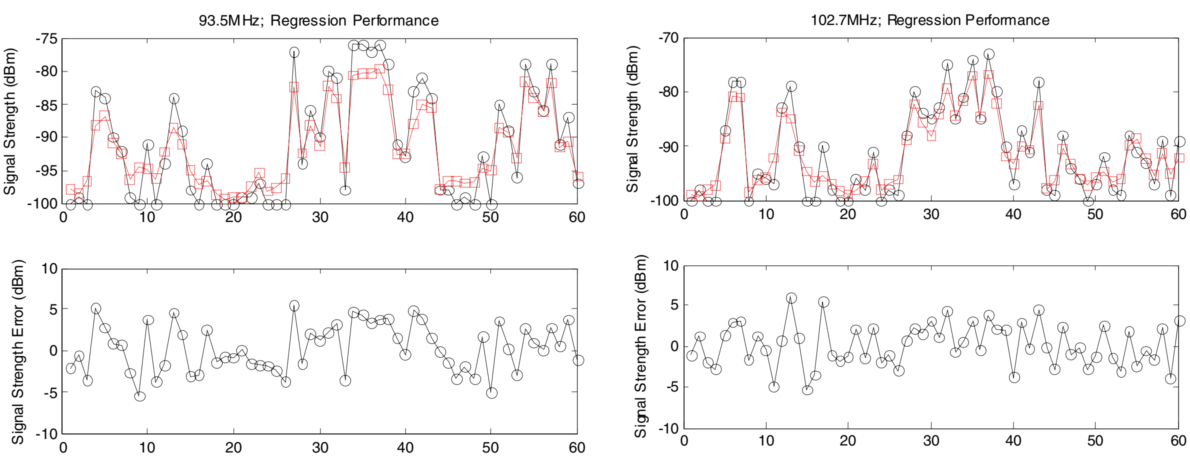

Figure 4: Examination of the predictive capability of Gaussian Processes regression (see main text for discussion). The figure shows the results of two experiments, provided by the two columns. The top image in each column provides a comparison of the predicted values at each location (black circles) and the measured values at those locations (red squares) for two opportunistic VHF public radio broadcast frequencies. The lower plots show the residuals between those measurement sets. The x-axis provides the sample number through the surveys.

used as a coarse guide. The variance of Gaussian Process signal strength maps could grow over time, or "age" to reflect a reduction in confidence regarding the structure or contents of a given building. Ideally a system should not have to rely on any prior surveying of an environment at all, which is the great benefit of the SLAM method described by this paper.

IV. Opportunistic Radio SLAM

Previous authors have considered WiFi SLAM using Gaussian Processes Latent Variables [14] and an efficient extension in the form of GraphSLAM [15]. These methods utilize an iterative process to converge upon a SLAM solution. Here we utilize a method based on Distributed Particle SLAM (DPSLAM) [16]. The advantage of this method is that it provides a continuous "online" SLAM solution, a continuous representation of the current user location probability density function, and also provides a flexible core navigation engine to permit fusion of other sensors and constraints when available (e.g. GNSS data, road-snapping, indoor floorplans, etc).

A. Occupancy Grid

The occupancy grid is a simple concept. As the user moves through the indoor environment and makes radio measurements, the spatial resolution for storing these radio measurements defines a grid of permitted locations. For example, there is no point choosing to store fingerprints on a millimetre scale, or on a kilometre scale, when aiming to provide metre-level positioning indoors. A metre-scale occupancy grid is much more sensible. As a user moves through an environment, the particle cloud providing an estimate of the Probability Density Function (PDF) of position will be spread across a number of cells in the occupancy grid. At each measurement epoch the user will record the signal fingerprint at their current true location and each particle will enter their identity and the current time into the cell they currently occupy in the occupancy grid. When a particle revisits a cell in the occupancy grid, it can "look up" if it has been in that cell before and compare the current and old set of signal strength measurements. This provides the basis for the SLAM loop closure, as discussed below.

B. DPSLAM

Distributed Particle Simultaneous Localisation and Mapping is a particle-filter-based SLAM method [16]. DPSLAM exploits efficient databases via pointers and binary trees to maintain a history of the states of interest for each particle over time. The particle cloud states are initialized by some primary positioning system (e.g. GNSS signals outdoors before a user enters a building, or via signal fingerprinting from a prior map in a region where a prior map database is available, etc). While operating in a new GNSS-denied region, all positioning updates are provided by inertial measurements alone (e.g. ZUPTS or the simple step-and-compass method used here), resulting in a gradual decrease in the certainty of the user's true location (i.e. the particle cloud disperses as the PDF expands). When the user is provided with some positioning constraint (e.g. GNSS availability) the particle cloud is reweighted (particles near the GNSS location estimate are given a high probability of representing the true user location) and resampled in the usual manner. The particle cloud then collapses accordingly, representing improved confidence in the user location. In this way, low probability particles and their history (including their entries into the occupancy grid) are "pruned" from the particle database, and high probability particles are duplicated to maintain a high density of particles around the maximum of the probability density function. The history of the particle cloud can also be updated at this point by reweighting the history of all particles, resulting in the track of the user being updated as well as the current location.

This method extends to a SLAM framework when the positioning update is not via an external aiding mechanism, but is provided by the user revisiting locations and re-observing "landmarks", in this case making signal fingerprint measurements in revisited locations within the occupancy grid. If the signal environment is complex and varied, as is typical indoors, then signal fingerprints can vary measurably on short length scales (e.g. a few metres). When a user revisits a location, this will be reflected in the signal fingerprint measurements and the particles can be reweighted accordingly. A particle that revisits a cell in the occupancy grid and retrieves its old signal fingerprint, only to find that the current measured fingerprint is completely different, will be given a low probability weighting (it is unlikely that the user is really revisiting an old location, else the fingerprints would be similar). However a particle that revisits a cell, retrieves its old fingerprint and discovers that it is similar to the current fingerprint measured by the user will be given a high probability weighting, i.e. it is plausible that this particle represents the current user location because the old and new fingerprints match. As the user continues to move, if this particle continues to make a sequence of similar fingerprint measurements, then its probability weighting will stay high, it will spawn new particles at every epoch and the particle cloud will collapse onto it (low probability particles will be removed and high probability particles will be duplicated).

C. Method

The indoor tracking experiments discussed here were all performed using data gathered on a "smartphone" cellular telephone. The data was then processed offline using Matlab, although the processing time was much faster than the length of the experiments. The computational overhead was dominated by the number of particles (1500) used in the particle filter, which in turn is affected in part by the resolution of the occupancy grid (1m) and the performance of the inertial measurement unit. The availability of ZUPTs would greatly reduce the number of particles required.

1) Smartphone measurements

A simple Android program was written to log all sensor data from the smartphone to a file. The data included accelerometer, compass, GPS, WiFi, and cellular measurements. The WiFi measurements could be polled every second, but cellular measurements could not be polled by the user and were returned by the device at a varying update rate of approximately 0.1 – 0.25 Hz. It seemed that the device only returned cellular measurements when a change in power level was recorded.

2) Pedestrian motion estimation

A foot-mounted IMU method (ZUPTS) was not employed for a variety of reasons, including the lack of a gyroscope within the smartphone used. The inertial measurement process was relatively simple, with accelerometer thresholds used to determine that the user had made a stepping motion, and the compass used to provide an estimate of the direction of this motion. When walking with a hip-mounted IMU, the total accelerometer magnitude varies dramatically, typically registering a sudden but characteristic spike from around 0.8g to around 1.5g as the torso drops towards the ground with the leading leg, then is suddenly arrested as the leading foot hits the ground, as shown in Figure 5. A simple moving window across the streaming accelerometer data can therefore be used to register and count steps. The compass data was low-pass-filtered to remove the high frequency perturbations caused by the user stepping motion. The IMU (smartphone) was mounted on the hip of the user during use. Aspects such as detecting motion types (walking and running, moving up and down stairs, sidestepping, etc) were also investigated. For the experiments described here the user walked around the ground floor of the indoor environment. In future a MEMS barometer may be useful in detecting floor changes to permit 3D indoor tracking.

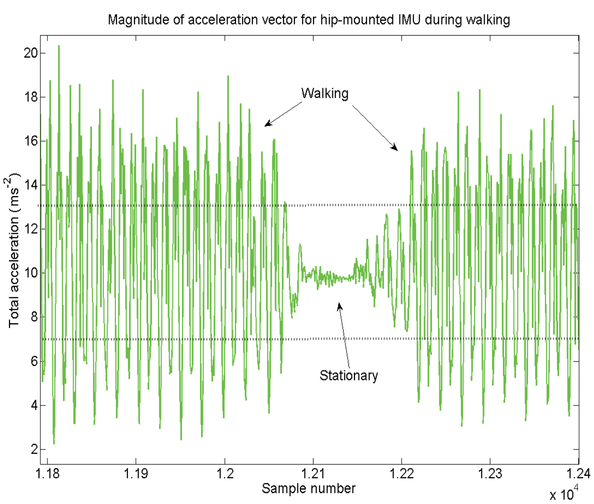

Figure 5: Stepping motion is easily detected from accelerometer data using thresholds (shown by the dotted horizontal lines)

3) Navigation Engine

The navigation engine utilized a particle filter to track the step length, compass bias, latitude and longitude of the user. The engine was initialized outdoors with an initial position estimate and uncertainty provided by GPS. This defined the initial size and position of the particle cloud. The engine was also initialized with an expected range in possible user average step length and compass bias. When the engine was initialised, each particle was allocated a randomly chosen step length and compass bias within this range. As the user moved, the accelerometer thresholds triggered each step event and the low-pass-filtered compass reading was passed to the engine. Each particle propagated their step length in the direction of the compass heading, with each particle's slightly different step length estimate and compass bias setting resulting in the particle cloud growing over time. Each particle was also perturbed by a small randomly-generated amount to represent the small random variations in a user's step length at each pace, and also to represent the measurement noise on the magnetic compass measurement. The particle filter could be constrained via GNSS measurements or DPSLAM processing. The user step length and compass bias became observable during periods of GNSS availability and whenever DPSLAM loop closure occurred. To further account for variations in step length and compass bias during the course of the tracking process, whenever GNSS measurements or a loop closure caused the particle cloud to collapse, the new particles generated by the reweighting and repopulating process were each given a small perturbation in step length and compass bias to the values originally allocated to their parent particle.

D. Experimental results

The results of a typical experiment are provided here in the form of images at various stages in the process (see Figure 6). The first image (marked A) shows the initial position of the user (represented by the particle cloud), provided by an outdoor GPS measurement. The red arrows visible on the map show the true path taken by the user. The floor plan itself is not available to the navigation engine, i.e. there is no knowledge of walls, rooms, outdoor regions, indoor regions etc. The next image (B) demonstrates the short path taken by the user to enter the building. Two paths are now visible. The blue path is the result of using the step detection and heading measurements alone, with no radio information provided. The green path represents the output of the full navigation engine fusing inertial measurements with opportunistic radio and GPS data. In these first few seconds of the journey, the availability of GPS provides some constraint on the step length of the user, and compass bias of the device, such that the navigation engine is partially calibrated when the user enters the building. The distance covered by the user by this stage is approximately 50 metres.

With no prior knowledge of the opportunistic radio signal behaviour inside the building the positioning estimate is provided only by the step detection and magnetic compass estimates, and so the accuracy degrades with time. The particle cloud grows accordingly as the user moves through this new environment (image C). However by image D, the user has revisited a region encountered earlier on in the journey, when the true location of the user was known more accurately. Some of the particles in the cloud (the ones best representing the real path taken by the user) also revisit cells in the occupancy grid that they have populated before. The current radio measurements are compared to the old radio measurements stored in these occupancy grid locations, and the particles are reweighted according to how well their past and current measurements match. The new particle weights

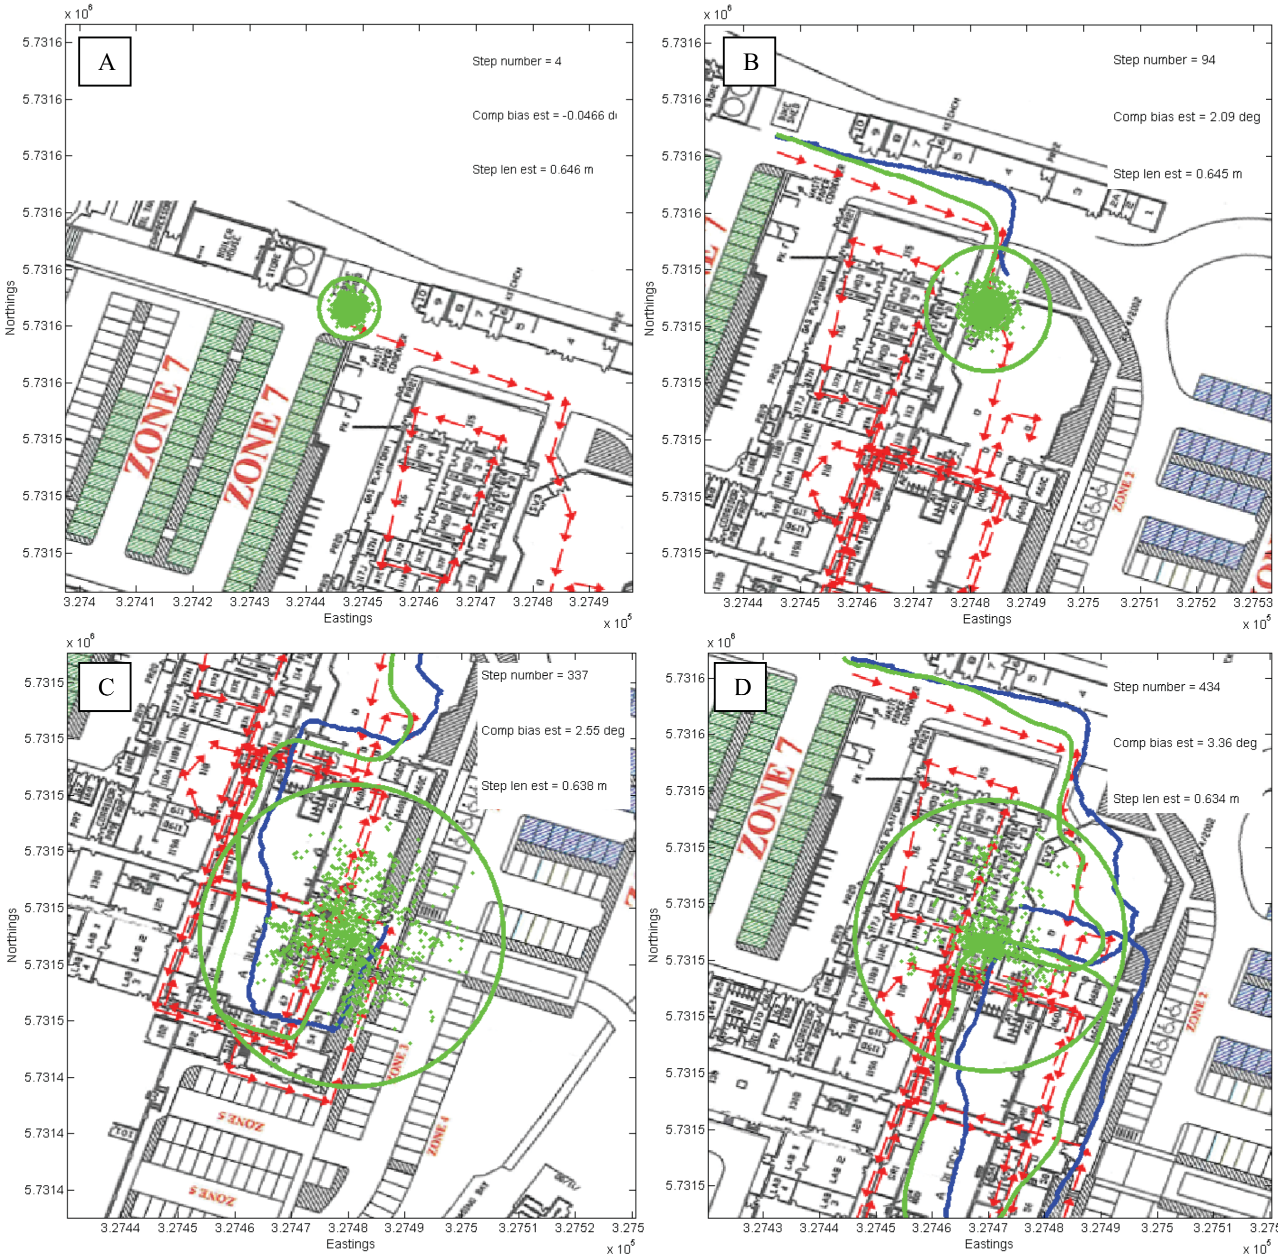

Figure 6: These images show the first loop closure for an indoor positioning experiment using Opportunistic Radio DPSLAM. The red arrows show the true user motion during the whole experiment (see Figure 7). The blue line is the track from only the step detection and compass measurements. The green line represents the DPSLAM track. The information represented by the green circle is discussed in the main text. During images A and B, GNSS measurements allow the step length and compass bias to be observed. Image C demonstrates the growth of the particle cloud when using dead reckoning alone. Image D demonstrates the collapse of the particle cloud caused by the user moving through a previously-visited corridor such that the current radio fingerprints match those stored by the user earlier in the journey. The floor plan is provided as a visual aid here, but was not available to the navigation engine.

are applied to both the current particle cloud, and the history of the track. As the user continues to operate in the environment the system is able to constrain the drift on the user's position while generating a map of the signal strength fingerprints within this building. By the end of the 15 minute walk around the building the user's final position error is only 4 metres, compared to the error on the uncorrected inertial measurement estimate of 86 metres. The final user track is shown in Figure 7.

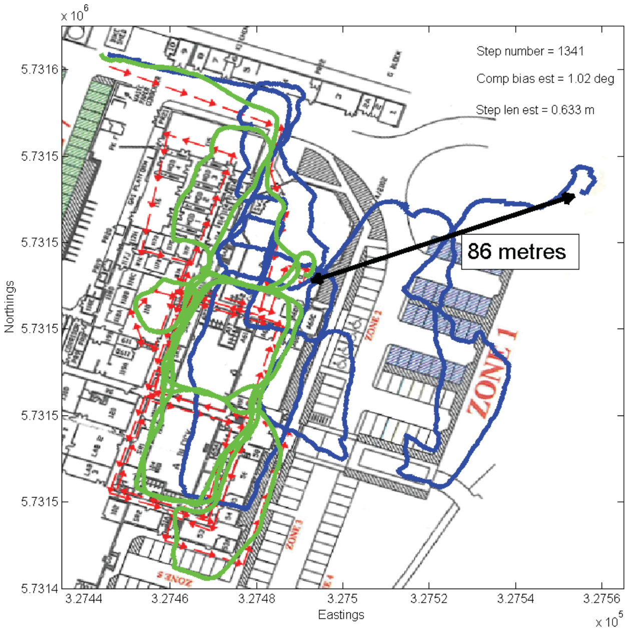

Figure 7: The entire track produced by this fifteen minute indoor walk. The final position error for the unassisted IMU track (blue line) is 86 metres. The final position error for the Opportunistic Radio DPSLAM solution is 4 metres. The largest error at any point along the Opportunistic Radio DPSLAM solution track is 12 metres. It should be noted that during the system operation the instantaneous error can be larger than this, but that the history of the user motion is updated and improved whenever the user revisits locations or when GNSS measurements become available.

The images in Figure 6 also contain a green circle which varies in size. This circle represents the overall global position error at this point. At early stages of the journey, before any particle re-weightings occur, the size of this circle is set by the size of the particle cloud and is stored at each epoch by each particle in the occupancy grid. This is critical for later particle cloud collapses as it maintains the history of the error associated with the dead-reckoned sections of track. The particle cloud can collapse down onto a previous section of track quite well, seemingly suggesting that the user's location is well known. However that section of track may itself carry a large global error, as it may have been created after a significant period of dead reckoning. When a particle cloud collapse occurs due to a DPSLAM loop closure therefore, the particles extract the error circle size associated with that section of track history, and display that error estimate to the user as the error circle. The user may therefore see a tight particle cloud demonstrating that the user location is well known along its own historical path, but a large error circle demonstrating that the global estimate of the user location is not as well known. These historical error estimates could be corrected by the user generating a nearby GNSS estimate at some time in the future or by sharing Gaussian Process signal strength maps or training data amongst other users.

V. CONCLUSIONS AND FURTHER WORK

A new method of indoor radio positioning has been developed that exploits opportunistic radio positioning and the DPSLAM technique. The flexible multi-sensor navigation engine provides indoor and outdoor radio positioning capabilities without any prior surveying or knowledge of the signal environment. The opportunistic signal source locations do not need to be known, nor do the signals need to be demodulated, decoded, or contain any specific structure. The only critical assumption is that the opportunistic transmissions maintain a fixed power output, and the use of signals populating known radio bands such as VHF FM, cellular, television, etc, will allow the user to confidently make that assumption. The user's indoor position estimate is provided primarily by a simple inertial measurement unit, but the associated errors that normally increase with time can be corrected and bounded by opportunistic radio measurements within the SLAM framework.

The indoor global positioning accuracy is dependent on the initial calibration of the system and so factors such as the accuracy of the initial GNSS position estimates while operating outdoors, and the calibration of the user step length estimate and initial compass bias during initial periods of GNSS availability. The indoor system accuracy is also dependent on the user periodically revisiting previous locations, to allow the user's error to become observable and to be corrected.

A powerful capability of this system lies in the ability to generate signal strengths maps for an area quickly and easily using this DPSLAM approach and Gaussian Process regression methods, rather than through a slow manual surveying approach.

Current on-going work includes the addition of a simple wideband radio scanner to the sensor set to allow the incorporation of other radio bands than those recorded by a smartphone. Since this DPSLAM positioning system can exploit any signal that is broadcast with fixed power from a static transmitter, even the broadcasts from GNSS jammers could provide useful underlying indoor signal strength maps for this this system to exploit. This system therefore not only provides a navigation aid when GNSS is denied by difficult signal environments, the indoor positioning performance can actually improve in the presence of GNSS jammers.

Magnetic anomaly measurements can also provide a useful fingerprinting metric to add to the set of opportunistic radio measurements. It will also be desirable to combine the Gaussian Processes regression methods with the occupancy grid to enable particles to be sensitive to measurements stored in nearby occupancy grid cells. This should permit the system to function with a sparse particle cloud, and may permit useful loop closures to occur when the user moves to locations nearby to sections of historical track, rather than having to revisit exact previous locations.

ACKNOWLEDGMENT

The authors wish to acknowledge and thank the UK Centre for Defence Enterprise for funding this work.

REFERENCES

[1] http://www.skyhookwireless.com/

[2] S. Schon, O. Bielenberg, "On the capability of high sensitivity GPS for precise indoor positioning", Inst. fur Erdmessung, Leibniz Univ., Hannover; Proceeding of Positioning, Navigation and Communication, 2008. WPNC 2008B.

[3] L. Ojeda and J. Borenstein, "Personal dead-reckoning system for gps denied environments," in Safety, Security and Rescue Robotics, 2007. IEEE International Workshop on, 2007, pp. 1 –6.

[4] http://www.ubisense.net/en/

[5] http://www.q-track.com

[6] CSR Positioning Center (CPC) for Mobile Operators http://www.csr.com/assets/common/CSR_Positioning_Center.pdf

[7] M. Rabinowitz, J. J. Spilker, "The Rosum Television Positioning Technology," Proceedings of the 59th Annual Meeting of The Institute of Navigation and CIGTF 22nd Guidance Test Symposium, Albuquerque, NM, June 2003, pp. 528-541.

[8] T. Pratt, R. M. Faragher, P. J. Duffett-Smith, "Fine Time Aiding and Pseudo-Synchronisation of GSM Networks," Proceedings of the 2006 National Technical Meeting of The Institute of Navigation, Monterey, CA, January 2006, pp. 167-173.

[9] R. M. Faragher, P. J. Duffett-Smith, "Measurements of the effects of multipath interference on timing accuracy in a cellular radio positioning system" Radar, Sonar and Navigation, December 2010, Vol 4 Issue 6 Page 818 – 824

[10] R. M. Faragher "The Effects of Multipath Interference on Radio Positioning Systems" PhD Thesis, Cambridge University 2007

[11] L. Tsung-Nan, L. Po-Chiang, "Performance comparison of indoor positioning techniques based on location fingerprinting in wireless networks" International Conference on Wireless Networks, Communications and Mobile Computing, 2005 13-16 June 2005

[12] C.E. Rasmussen, C.K.I. Williams, "Gaussian Processes for Machine Learning", MIT Press, 2006, ISBN 026218253X

[13] B. Ferris, D. Hähnel, D. Fox, "Gaussian Processes for Signal Strength-Based Location Estimation", Proceedings of Robotics: Science and Systems, 2006

[14] B. Ferris, D. Fox, N Lawrence "WiFi-SLAM Using Gaussian Process Latent Variable Models" Proceedings of IJCAI 2007

[15] J. Huang, D. Millman, M. Quigley, D. Stavens, S. Thrun, A. Aggarwal, "Efficient, generalized indoor WiFi GraphSLAM", 2011 IEEE International Conference on Robotics and Automation (ICRA), 9-13 May 2011

[16] A. Eliazar, R. Parr, "DP-SLAM: Fast, robust simultaneous localization and mapping without predetermined landmarks", International Joint Conference on Artificial Intelligence 2003