

Ultrasound-Aided Pedestrian Dead Reckoning for Indoor Navigation

Carl Fischer
Computing Department
Lancaster University
Lancaster, UK.

Kavitha Muthukrishnan
Faculty of Computer Science
University of Twente
Enschede, The Netherlands.

Mike Hazas
Computing Department
Lancaster University
Lancaster, UK.

Hans Gellersen
Computing Department
Lancaster University
Lancaster, UK.

ABSTRACT

Ad hoc solutions for tracking and providing navigation support to emergency response teams is an important and safety-critical challenge. We propose a navigation system based on a combination of foot-mounted inertial sensors and ultrasound beacons. We evaluate experimentally the performance of our dead reckoning system in different environments and for different trail topologies. The inherent drift observed in dead reckoning is addressed by deploying ultrasound beacons as landmarks. We study through simulations the use of the proposed approach in guiding a person along a defined path.

Simulation results show that satisfactory guidance performance is achieved despite noisy ultrasound measurements, magnetic interference and uncertainty in ultrasound node locations. The models used for the simulations are based on experimental data and the authors' experience with actual sensors. The simulation results will be used to inform future development of a full real time system.

Categories and Subject Descriptors: C.3 [Special-Purpose and Application-Based Systems]: Real-time and embedded systems

General Terms: algorithms, experimentation, measurement.

1. INTRODUCTION

Search and rescue is a challenging and dangerous activity. The environment is often unfamiliar and changing, and visibility can be limited. The rescue operations are time-critical and hence quick decision making support and close coordination within teams are required. Ad hoc tracking and navigation support for emergency response is an impor-

tant and safety-critical challenge. A report on the Worcester warehouse fire, in which six firefighters died, highlights the difficulty to keep track of firefighters within the building as one of the major causes for loss of lives [1], and a report on fatalities in structure fires linked 29 casualties between 1990–2000 to firefighters becoming lost inside the structure [3]. The application pull for new technologies to address safety of emergency responders is evident in major initiatives including fire services, fire protection agencies and relevant industries [9, 15, 18] but new research is required to tackle the problem of ad hoc tracking and navigation.

We envisage a system that will aid the search and rescue operation by *tracking* the responders' position and informing the incident commander about their location inside the building, and by *guiding* the responders within the building under poor visibility conditions, thereby helping them reach victims faster and leave the building quickly and safely when necessary.

Inertial navigation or pedestrian dead reckoning (PDR) has been applied to tracking and navigation of first responders with promising results. However the position error in a purely inertial system increases with time and requires correction from external sources. A common practice is to periodically use GPS to correct position estimates [13]. But for most indoor scenarios GPS is unavailable. Embedding sensors or tags into the building fabric to act as landmarks is another solution but this only works in modified buildings and cannot be rapidly deployed in arbitrary locations.

We plan to address the problem of positional drift by having the responders themselves deploy landmarks as they progress into an unknown environment. We will specifically use ultrasound nodes. The "breadcrumb" trail thus created can be used to assist the PDR in guiding the responders back to their starting point, or guiding other responders towards a victim or an alternative exit. The benefit of the deployed landmarks is particularly interesting when locating multiple responders relative to each other and enables better coordination within teams.

In this work we look at what could be achieved using such a system for guiding the user to the required destination. Through simulations we show that PDR alone is not sufficient but that by deploying ultrasound sensors along the path the user can be successfully guided to their destination.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MELT'08, September 19, 2008, San Francisco, California, USA.
Copyright 2008 ACM 978-1-60558-189-7/08/09 ...\$5.00.

We also show that this applies even in the presence of noisy ultrasound measurements, magnetic interference and when the locations of the ultrasound nodes are only known with some uncertainty.

2. RELATED WORK

Emergency response is an area where distributed sensing and localisation not only provide extra services to the users but are intended to save lives. Different sensing technologies have been used in literature to solve localization and tracking problems in search and rescue missions. The Fire project [15] has developed SmokeNet, a wireless network of smoke detectors which doubles as a location system based on radio signal strength. The Flashlight by Peterson and Rus [11] guides a person through a sensor network avoiding danger zones by providing tactile feedback when they are facing the right direction. We believe that RF-based sensors are not suited to indoor navigation because they do not account for walls. Ultrasound propagation on the other hand is inherently limited by walls and doors thus guaranteeing room-scale granularity or better.

The indoor positioning system [6] developed by Thales works similarly to GPS but indoors: firetrucks parked around a building act as “satellites” that use ultrawide band (UWB) RF signals to locate firefighters inside a building by means of time of arrival measurements. Although this system might perform well for lightweight residential buildings, UWB may not penetrate larger structures that extend underground for instance. For this reason we choose to deploy a physical chain of sensors that can create a link to the outside both for navigation and communication purposes.

Dead reckoning has the distinct advantage of providing autonomous positioning capabilities and is thus particularly attractive for indoor search and rescue operations. However positions provided by this method will unavoidably drift over time due to errors in measurements being integrated [4]. The drift can be reduced by using shoe-mounted inertial sensors and resetting the velocity to zero at each footfall [10] and by combining the inertial measurements with data from an electronic compass through a Kalman filter in order to avoid drift in heading [5]. It has been shown that disruptive motion such as side-stepping, back-stepping, tight turns that are typical in search and rescue scenarios produce scaling errors and cause the travelled distance and thus the estimated position to drift even more than during normal walking. Despite these limitations dead reckoning is the only completely self-contained location technique that requires no prior knowledge of the environment. This is why we and others attempt to address these limitations by combining dead reckoning with other complementary technologies.

In most cases it is essential to correct positions and headings with data from external sources. GPS is one possibility but only for outdoor navigation with short periods of GPS outage [13]. Another possibility is to predeploy RFID tags at known locations and use these to correct positions [19]. Indoor location systems such as UbiSense have also been used in combination with PDR [8]. However there is no guarantee that a building will be equipped with any particular location infrastructure.

The navigation system developed by Renaudin et al. [14] combines PDR with map matching in order to prevent drift. Inertial measurement units (IMUs) on the chest and legs are used to measure movement and posture. The first team to

enter the building place an RFID tag on each door frame they pass through. The position computed by the inertial navigation system (INS) can then be corrected according to a database of the coordinates and directions of all doors in the building. The second team are equipped with an RFID reader and can therefore determine their positions as they scan each tag. This is an attractive solution since it is entirely adhoc. Nevertheless it requires floorplans of the building and will fail in areas with few doors such as open plan offices or airport terminals.

In our system we will use ultrasound nodes from the Relate project [7] as landmarks to correct the drift in PDR. Ultrasound has also been used in several other location systems [12, 16].

3. CHARACTERIZING PEDESTRIAN DEAD RECKONING

Dead reckoning is a self-contained navigation technique in which measurements — typically from inertial sensors in the case of *pedestrian* dead reckoning — are used to track the position and orientation of an object given an initial position, orientation and velocity. No infrastructure is required but the position error will increase over time due to noise.

3.1 Our PDR algorithm

XSens’s MTx [17] is an inertial measurement unit (IMU) comprising a tri-axis accelerometer, gyroscope and magnetometer. The on-board processor computes drift-free 3D orientation. Our pedestrian dead reckoning algorithm is similar to other work described in [5, 2] which also use shoe-mounted IMUs and apply periodic zero velocity updates (ZUPTs). In order to convert the MTx measurements into meaningful positions, the raw accelerations are rotated from the sensor coordinate system into the world coordinate system using the rotation matrix computed by the MTx as shown in Figure 1. The accelerations are then double integrated to yield position estimates. In order to reduce the position error which increases quadratically with time we reset the integrated velocities to zero at each step thus making the error linear with distance covered.

Two phases in walking are identified: the stance phase, when the foot is in contact with the ground, and the swing phase. During the stance phase the velocity is reset and kept at zero; during the swing phase the acceleration is double integrated. Our algorithm detects the stance phase of each step by applying a threshold to the product of the norm of the acceleration by the norm of the rate of turn as suggested in [2]. If this product is below an empirically determined threshold for more than 0.2 seconds then a stance phase is detected. When the product rises above the threshold again a swing phase is detected. This is illustrated in Figure 2. If some steps are taken at a faster pace then the stance phase may not always be detected and some opportunities for ZUPTs may be missed.

In all of our experiments we sampled orientation and inertial data at 100 Hz, the maximum speed at which the on-board processor can compute orientation, but our algorithm also performs with similar results at 50 Hz.

3.2 Performance evaluation

In this subsection we report the performance of our PDR algorithm with real data for various trail topologies in different environments.




Figure 1: Transformation from sensor to world coordinates via the direction cosine matrix: $x_w = R_{GS}x_s$.




Figure 2: PDR algorithm: each step has a stance phase (shaded) and a swing phase. Velocity is reset to zero during the stance phase, acceleration is double integrated during the swing phase.

3.2.1 Experimental setup

In all the experiments the IMU was firmly attached under the laces of the user’s shoe. We had an on-line implementation of the algorithm described above, recording the user’s trajectory on a Sony Vaio hand-held computer connected to the IMU. One set of experiments was run in our university building (Infolab). We considered different trail topologies: straight line (88 meters in total), L-shaped (54 meters in total) and rectangular (11.5 meters in total). This was done by two users. Another set was run in a similar building in another institute (TZI). We walked along a long corridor, entering several offices along the way (140 meters in total). A third set was run in a large industrial workshop (BIBA1). A single user walked a complex path of over 200 meters around heavy machinery. A final set was run in the office corridors around the workshop (BIBA2, 220 meters in total). This was done by six different users, three times each. In all the experiments the user returns to the starting point.

3.2.2 Error analysis

There are two major sources of errors in the PDR approach — error in distance and error in heading. For the straight line in Figure 3, the estimated distance drift is +2 percent of the total travelled distance. For the L-shaped path in Figure 4 we get an error of -8 percent of the total travelled distance. For the rectangular path (not shown) we get a closed loop where the starting and ending points are

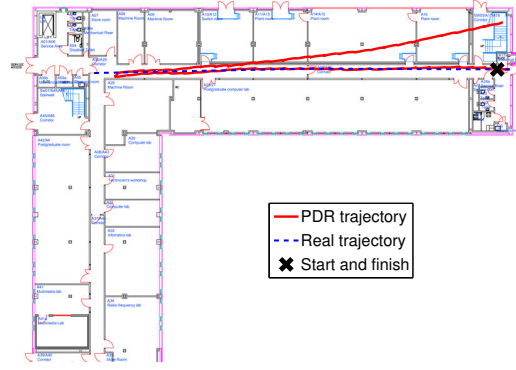


Figure 3: PDR straight line path, Infolab.



Figure 4: PDR L-shaped path, Infolab.

the same, but the error is -7 percent of the total travelled distance. We notice that heading errors tend to occur when the user does a 180 degree turn.

Figures 5 and 6 show that the performance of PDR can be impacted significantly by heading errors. We tested it for cases where the user walks along a corridor and enters several rooms along the way; the path shown in Figure 5 starts well but severely drifts off after 40 meters. The drift happens in one particular place and then again just after the 180 degree turn. For the U-shaped path in Figure 6 the error in heading is extreme due to interference from machinery in the nearby workshop. All experiments in the same corridor at BIBA2 show an almost identical error pattern suggesting that there is some particular magnetic interference in certain locations. It remains puzzling that Figures 5 and 6 exhibit strong interference on the forward path but not on the return path. We know that the MTx internal filter is sensitive to the amplitude of accelerations — hence to the speed of walking — and also attempts to compensate for magnetic interference. This may explain some of the differences between the forward and return paths.

Although some distance drift is inevitable due to the integration of noise and offsets in the raw sensor data, we also believe that most of the distance error is due to the MTx incorrectly estimating its orientation as explained by Foxlin in [4], thus we might interpret some of the forward motion as




Figure 5: PDR path with user entering offices along the corridor, TZI.




Figure 6: PDR path “worst case scenario” with strong magnetic interference due to nearby machinery, BIBA2.

vertical motion or vice-versa. The MTx internal algorithm is a black box meaning we have very little information about how the different sensors are used in computing the orientation and almost no control over any of the internal parameters. We assume that most of the heading errors are due to metallic objects or magnetic fields interfering with the MTx magnetometers since these extreme heading errors occur systematically in the same locations. We also note that when using the system outdoors in an open space the results are much better and the orientation drift is negligible. So it seems that magnetometers help in outdoor situations where they accurately determine magnetic North but that indoors they cause heading errors due to interference.

3.3 Consequences

We consider the consequences of these observations on guiding a user along a path. The biggest problem we observe in the PDR approach is the drift in orientation. Even

if the position is corrected by some other sensor modality, heading error means we cannot guide the user because we do not know which direction they are facing. Drift in the distance estimates are unavoidable but they remain small and consequences for guidance are less important. We believe that most of the drift in both distance and heading is due to the MTx internal Kalman filter incorrectly estimating the sensor orientation and that it will improve with future developments by XSens and others. The type of errors we have observed make it difficult to quantitatively evaluate performance which can vary from “almost perfect” to “unusable” depending on the level of magnetic interference. In all the cases we observed that most individual segments of the recorded paths are very accurate — even a spiral staircase at BIBA1 was correctly recorded — but that strong heading error occurs at particular locations. Manual correction of the position and heading can give good results but the challenge is to make these corrections automatic.

4. SIMULATION OF A GUIDANCE SYSTEM

We plan for search and rescue teams to deploy small ultrasound beacons as ad-hoc landmarks along their path. These beacons can then be used by other teams or by the same team to assist them on their way back. The team members wear boots equipped with ultrasound transmitters that can be located by the beacons, and inertial sensors. We investigate how such a system might perform through simulations.

4.1 Measurement model

We model the ultrasound and inertial measurements based on our observations of data from deployments in realistic environments outside the lab.

4.1.1 Ultrasound measurement model

The ultrasound location estimates are very noisy. We model the range and bearing measurements as Gaussian with standard deviations of 5 centimeters and 30 degrees respectively [7]. A fraction of the range measurements are made outliers by adding 3 meters to the real value. Because the ultrasound location estimates are so noisy we only use them to correct the PDR location estimates if the discrepancy between the ultrasound and PDR estimates is greater than a threshold (on the scale of a meter or more). If the PDR location estimate and the ultrasound location estimate are reasonably consistent then we continue to rely on the PDR since this will give smoother results. If the estimates are not at all consistent then we trust the ultrasound location estimate. The ultrasound location estimate is used as the new location and the heading of the PDR is adjusted using a simple trigonometric formula which returns the angle between the current (wrong) location estimate, the last reliable location estimate, and the new (almost correct) ultrasound location estimate. This formula gives good results in practice but only if the ultrasound measurements are frequent enough.

4.1.2 Pedestrian dead reckoning model

The successive positions of the user are not known in advance and the error in heading is dependent on position so it must be calculated dynamically. We assume that the error in heading is mostly due to magnetic interference however the internal Kalman filter of the inertial measurement unit means the heading is not only affected by the local magnetic

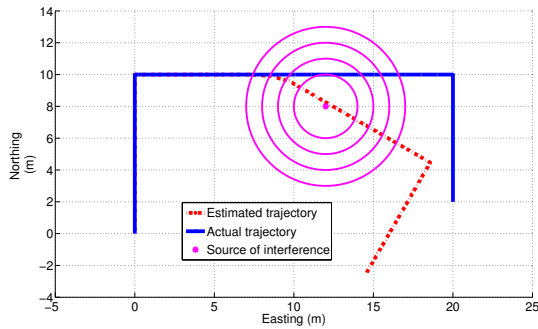


Figure 7: Magnetic interference model affecting PDR estimates.




Figure 8: Guidance algorithm: (1) estimate the person’s position, (2) project onto path, (3) find target position, (4) compute guidance angle.

field but also by the magnetic field at previous locations. In our model we define sources of magnetic interference and for each source a radius and an amplitude. When the user moves closer to the source than the given radius then the heading is modified by the given amplitude. The sign of the modification in heading depends on the direction that the person approaches the source. This empirical model illustrated by Figure 7 replicates the effects that we have observed during our experiments. Note that in our current simulator the only random element is the ultrasound measurements, the PDR error is deterministic.

4.2 Guidance algorithm

One important goal of our work is to guide search and rescue personnel follow a predefined path. The initial scenario we envisage is a wide open area such as a dark underground parking lot or an empty smoke-filled warehouse where a path has already been defined as the team went in and deployed ultrasound nodes along the path. As the team attempt to return to the exit back along the path they are guided by an arrow on a head-mounted display (HMD) showing them which way to walk. The path to follow is defined as a series of segments. Given the estimated position of the user we find the point on the path that is closest to their estimated position by projecting the estimated position onto the successive segments of the path. Then we direct the user to a point that is a few metres ahead along the path as shown in Figure 8.

In order to check the feasibility of this system we assume that the user always follows the direction provided. This shows us how often they reach their destination and how often they stray too far from the path and get lost. The simu-

lator is event based. Ultrasound and inertial measurements are generated periodically (e.g., every 200 milliseconds and 10 milliseconds respectively) and processed by the fusion algorithm to estimate the user’s position. Periodically (e.g., every 2 seconds) the guidance system computes which direction the user should travel and the user takes a step in that direction, effectively creating a feedback loop. A simulation run is considered successful if the person gets within a short distance of the end of the path under a certain delay.

4.3 Simulation results

If we run a simulation with PDR alone, that is without using the ultrasound measurements to correct position and heading, the user will be guided to the wrong location. In the sample simulation shown in Figure 9 the PDR wrongly believes the person is too far South and so they are guided towards the North. The simulation ends without the person reaching the destination because the system wrongly believes that they are already there. However if ultrasound measurements are used to correct the position estimates the person is successfully guided to the end of the path as shown in Figure 10.

Initial results show that if we do not use ultrasound measurements enough, the user will be guided away from the path and out of range of the beacons due to incorrect position and heading estimates. If that happens in the simulation the user is lost unless by chance they stray back into range of the beacons. In reality new nodes could be automatically deployed to create a new branch in the path or some special action could be taken if this occurs, at the very least by warning the user. If we use ultrasound measurements frequently then the user is likely to reach the end of the path safely. Using frequent ultrasound measurements to correct the position makes the position estimates and hence the guidance rather “jumpy”. This is not a problem in the simulation but may be a usability issue in an online implementation. Ideally the displayed arrow should rotate smoothly.

In another batch of simulations we introduce different levels of uncertainty to the ultrasound beacon positions and orientations and see how this affects the success of the guidance system. We discover that even for large errors in the estimated beacon positions the user can still reach their destination. Errors in the estimated beacon orientation are even less important as long as the user’s estimated position does not drift to far from their real position. This is good news for it means that the requirements for locating the beacons should be achievable. However improving the accuracy of beacon locations and orientations does improve the success rate so this is one way of achieving a more reliable system. Simulations also confirm that increasing beacon range and beacon density improve success rates.

5. CONCLUSION AND FURTHER WORK

The initial simulation results are promising and show that using both PDR and deployed ultrasound beacons to estimate a person’s position we would be able to provide sufficient information to guide them along a predefined path even in presence of magnetic interference and with noisy ultrasound measurements. The next phase will be an online implementation of this system and an experimental study.

It is important to realize that given the field of application such a system should be extremely reliable. Success

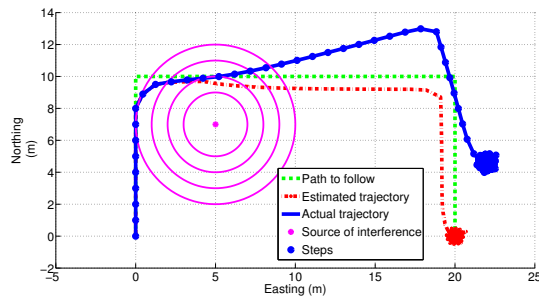


Figure 9: A simulation run showing the user following guidance along a trail but failing to reach the end because of drift in the PDR estimates.

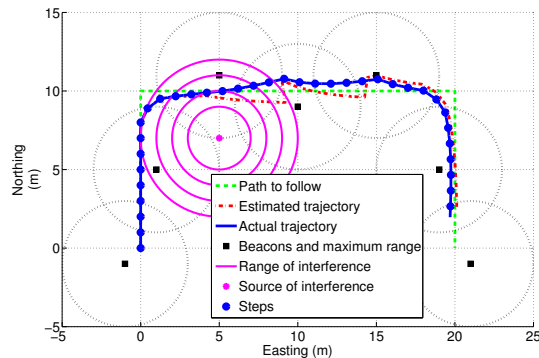


Figure 10: A simulation run showing the user successfully following guidance along a trail, their position is estimated with PDR corrected by ultrasound beacons.

rates of 90 percent are not enough or the system will not be trusted and the user's will continue to rely on other navigation methods. If 100 percent successful guidance proves unrealistic then it will be necessary to investigate other ways of informing the user about the current situation, about what has gone wrong, maybe providing a reliable way to retreat back to a previously known position rather than all the way to the exit. Physically deployed beacons have the advantage of being visible, especially if the casing is carefully designed and they include lights or sirens, and thus provide a fallback navigation method. We believe however that the harsh lighting and noisy environment of a fire scene combined with the high stress levels of the firefighters make our full guidance system far preferable since the cognitive load will be less.

Following these simulations we will now investigate how well a real person is able to follow such guidance as provided by our system during an experimental study. We may require smoother guidance data for real users. This could be provided by more sophisticated fusion algorithms based on Kalman or particle filters. The simulator will also be improved to take into account errors in the user's following of instructions and better models of PDR drift and ultrasound noise. Simulations will be used to examine the effect of other parameters such as beacon position uncertainty and beacon density for various trail topologies. A major challenge is calibrating the positions of the ultrasound beacons. We plan

to address this as a simultaneous localization and mapping problem (SLAM). This is a common topic in robotics but due to the nature of movement in pedestrian navigation and the trail topology of the beacons the solutions will be different. Finally the guidance provided to the user is limited to a simple direction which may only be sufficient for navigating in open spaces. Under low visibility this is already a challenge, but there are other navigation scenarios — for instance in a maze of cubicles — where our guidance system may be inappropriate because of obstacles, and different solutions may be required, at least in terms of how to provide visual guidance to the user. For instance it may be important to help the user distinguish between several doors or avoid walls by a visualisation of available paths. This requires work on the visualisation rather than the underlying location system.

6. REFERENCES

- [1] J. R. Anderson. Abandoned Cold Storage Warehouse Multi-Firefighter Fatality Fire. Technical report, USFA, 1999.
- [2] S. Beauregard. Omnidirectional Pedestrian Navigation for First Responders. In *WPNC'07*, pages 33–36, 2007.
- [3] R. F. Fahy. U.S. Fire Service Fatalities in Structure Fires, 1977–2000. Technical report, NFPA, 2002.
- [4] E. Foxlin. Motion Tracking Requirements and Technologies. In K. Stanney, editor, *Handbook of Virtual Environment: Design, Implementation, and Applications*, pages 163–210, 2002.
- [5] E. Foxlin. Pedestrian Tracking with Shoe-Mounted Inertial Sensors. *Computer Graphics and Applications*, 25(6):38–46, 2005.
- [6] D. Graham-Rowe. Indoor 'sat-nav' could save firefighters. *New Scientist*, 196(2634):24, 12/2007.
- [7] M. Hazas, C. Kray, H. Gellersen, H. Agbota, G. Kortuem, and A. Krohn. A relative positioning system for co-located mobile devices. In *MobiSys'05*, pages 177–190, 2005.
- [8] E. P. Herrera, R. Quiros, and H. Kaufmann. Analysis of a Kalman Approach for a Pedestrian Positioning System in Indoor Environments. In *Euro-Par 2007*, pages 931–940, 2007.
- [9] P. Lukowicz, A. Timm-Giel, M. Lawo, and O. Herzog. WearIT@work: Toward Real-World Industrial Wearable Computing. *IEEE Pervasive Computing*, 6(4):8–13, 2007.
- [10] L. Ojeda and J. Borenstein. Non-GPS Navigation for Emergency Responders. In *2006 International Joint Topical Meeting: Sharing Solutions for Emergencies and Hazardous Environments*, pages 12–15, 2006.
- [11] R. Peterson and D. Rus. Interacting with a Sensor Network. In *ICRA'04*, pages 180–186, 2004.
- [12] N. B. Priyantha, A. Chakraborty, and H. Balakrishnan. The Cricket location-support system. In *MobiCom'00*, pages 32–43, 2000.
- [13] V. Renaudin, O. Yalak, and P. Tomé. Hybridization of MEMS and Assisted GPS for Pedestrian Navigation. *Inside GNSS*, 2(1):34–42, January 2007.
- [14] V. Renaudin, O. Yalak, P. Tomé, and B. Merminod. Indoor Navigation of Emergency Agents. *European Journal of Navigation*, 5(3):36–45, 07/2007.
- [15] D. Steingart, J. Wilson, A. Redfern, P. Wright, R. Romero, and L. Lim. Augmented Cognition for Fire Emergency Response: An Iterative User Study. In *1st International Conference on Augmented Cognition*, 2005.
- [16] A. Ward, A. Jones, and A. Hopper. A New Location Technique for the Active Office. *IEEE Personal Communications*, 4:42–47, 1997.
- [17] XSens. XSens website. <http://www.xsens.com>. Accessed 2008.01.09.
- [18] S.-H. Yang and P. Frederick. SafetyNET - a wireless sensor network for fire protection and emergency responses. *Measurement and Control*, 39(7):218–219, 09/2006.
- [19] S.-Y. Yeh, K.-H. Chang, C.-I. Wu, H.-H. Chu, and J. Y.-J. Hsu. GETA sandals: a footstep location tracking system. *Personal and Ubiquitous Computing*, 11:451–463, 2007.