Cooperative Relative Positioning

Hans Gellersen, *Lancaster University*; Paul Lukowicz, *University of Passau*; Michael Beigl, *TecO, KIT*;
Till Riedel, *University of Karlsruhe*

Many pervasive applications deal with relative positions between interacting entities rather than global coordinates. The RELATE project has developed different sensing methods and a modular system architecture for peer-to-peer relative positioning, and studied these in application case studies on mobile spatial interaction, firefighter navigation, and wearable activity recognition.

Computing systems that leverage location information for a wide range of applications have become mainstream. Positioning of mobile entities in such systems is commonly based on a dedicated reference infrastructure. In general, the position is determined in a global coordinate system related to the specific infrastructure (e.g., a building in which the infrastructure is installed).

The work presented in this paper is based on the observation that many pervasive applications deal with *relative positions* of interacting entities rather than global coordinates. For instance:

- Services that help users locate and access devices relative to where they are – the nearest printer, the display to the left, the lights that are a few metres away.
- Applications that analyse and support human social interactions – whether they stand together and face each other.
- Navigation systems that guide users relative to their environment – from one exhibit to another in a museum, or relative to a command post in emergency response.
- Sensor networks that track spatial configurations for context analysis – whether hazardous goods are stored to close together, or where a user's hands are in relation to their body.

We propose that such applications can be supported with relative positioning systems in which the interacting devices cooperate to determine their relative positions, without need for an external reference infrastructure. In the RELATE project, a European collaboration on relative positioning, we have designed and evaluated relative position sensing and fusion methods, developed a system architecture that supports modular integration of different sensor modalities, and explored application of relative positioning in case studies on mobile spatial interaction, navigation of firefighters and activity recognition with body-worn sensors. In this article, we provide an overview of the project's key results and contribute a summative analysis of cooperative relative positioning: the challenges involved, the common principles underlying different technical approaches, the diverse application requirements, and the lessons learned from building systems for case studies.

1. Challenges in Cooperative Relative Positioning

In a cooperative relative positioning system two or more devices exchange messages and sensor signals to determine their relative position. The general principle of message and signal exchange can also be found in most infrastructure-based systems. However, unlike in most infrastructure-based systems, all devices are equal peers. Thus, there is no distinction between fixed components (for which the position is known) and mobile devices (for which the position needs to be determined). This has a number of consequences:

Location Sensing Methods. A common way to position a device in 3D is trilateration using at least three fixed reference points. In infrastructure-based systems these are the fixed beacons or receivers (in general more than three) for which the location is known. They are often carefully placed to optimize the localisation performance in a given environment. This method is not directly applicable in a relative positioning system. For one, we may often want two know the relative position of just two or three entities (e.g., two people interacting), which means that there are not enough points for trilateration. A possible solution is to use distance and angle measurements, which however increases device complexity. Secondly, we have to deal with a wide range of possible spatial distributions of the devices, many of them not well suited for optimal location performance. This can be addressed with optimization techniques that leverage redundant measurements.

Related work

Localisation of nodes based on peer-to-peer interactions is widely investigated in wireless sensor networks (WSN). Typically, RF signal strength is used as a range indicator (with well-known limitations) but researchers have also investigated more accurate techniques, such as interferometry. However, generally distributed localisation in WSN is dependent on a high degree of connectivity and a priori survey of anchor nodes [1].

For support of human interactions, capture of spatial relationships is often more relevant than absolute localisation. Location models provide relevant support, for instance, with range and nearest neighbour queries [2]. Some location systems have been specifically designed to measure user proximity. NearMe uses WiFi signals to model device and user proximity by comparing their lists of detected base stations and their signal strengths [3], and PeopleTones uses a similar approach but for GSM readings [4]. This provides coarse-grained proximity information, without need for additional sensors. The sociometric badge, in contrast, is a purpose-built sensing device for detection and support of social interactions [5]. It includes a 2.4 GHz radio for detection of co-located users, and an IR sensor for identification of face-to-face alignment.

The relative position of user actions with respect to entities in their environment can also be used as an interface for triggering actions, for example snapping of fingers to operate controls [6]. Such applications require fine-grained accuracy in order to be usable, in the referred work achieved with audio as sensing modality.

1. K. Langendoen and N. Reijers (2003). Distributed localization in wireless sensor networks: A quantitative comparison. *Computer Networks* 43(4): 500-518.
2. C. Becker and F. Dürr (2005). On location models for ubiquitous computing. *Personal Ubiquitous Comput.* 9(1): 20-31.
3. J. Krumm and K. Hinckley (2004). The NearMe Wireless Proximity Server, *Proc. UbiComp 2004*: 283-300.
4. K. Li, T. Sohn, S. Huang, and W. Griswold (2008). PeopleTones: a system for the detection and notification of buddy proximity on mobile phones. *Proc. MobiSys '08:* 160-173.
5. D. Olguín Olguín, B.N. Waber, T. Kim, A. Mohan, K. Ara and A. Pentland (2009). Sensible Organizations: Technology and Methodology for Automatically Measuring Organizational Behavior. *IEEE Trans. On Systems, Man, and Cybernetics B, 39(1):* 43-55.
6. J. Scott, B. Dragovic. Audio Location: Accurate Low-Cost Location Sensing. *Proc. Pervasive 2005:* 1-18

Coordination and Synchronisation. Cooperative relative positioning takes place within a group of devices that, in general, come together in an ad hoc, dynamic fashion. Thus, appropriate ad hoc protocols are needed to coordinate the transmission of signals and messages within a group of devices. The protocols must scale with respect of the number of devices and allow for dynamic joining and leaving of the group.

Resource Limitations. In an infrastructure-based system complexity can be outsourced to the fixed nodes and their associated resources. Thus, the mobile nodes can be kept simple. In an infrastructure-less relative positioning system the entire functionality must be part of each device. Since the relative positioning capability is, in general, just an additional feature of a mobile device, only limited resources can be devoted to it.

From the application point of view, relative positioning is mostly required over comparatively short distances (1-3 meters), since interacting entities tend to be within close proximity. On the other hand many applications can benefit from orientation information ("which device is the user facing?"), which is not commonly available in fixed-infrastructure positioning systems.

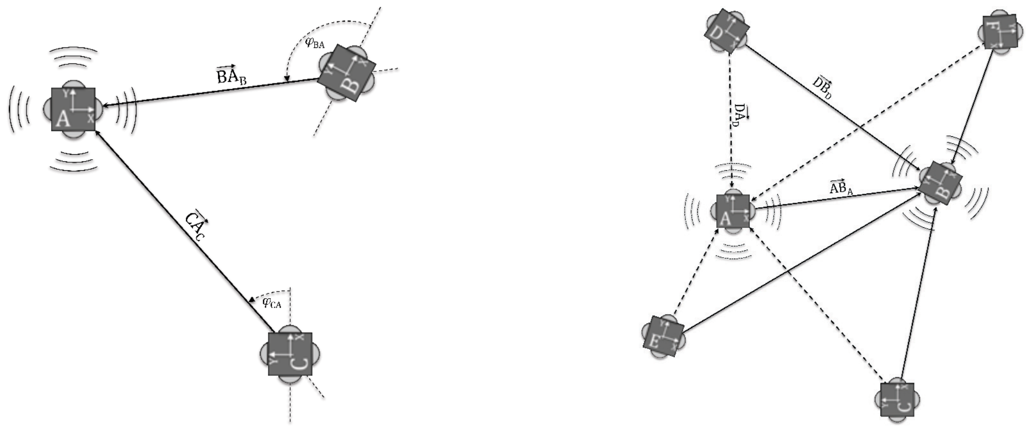

a. Devices individually measure the position of a transmitting node, relative to their local coordinate system.

b. Multiple devices can simultaneously perform a measurement of the same spatial relationship.

Figure 1. Relative Positioning Principles

2. Relative Location Sensing Methods

2.1. RELATE Concept for Relative Positioning

The RELATE concept for relative positioning is illustrated in Figure 1. Each device is equipped with multiple transceivers for sending and receiving of sensor signals. In this example, devices have four transceivers for omni-directional sensing in two dimensions. However, other geometries may be chosen depending on sensor modality and application requirements (e.g., for positioning in 3D). One device at a time can emit sensor signals while all other devices are in receiving mode. The receiving devices use the incoming signal to derive estimates for the range and the angle-of-arrival of the emitted signal. The range and angle-of-arrival together provide a relative position estimate of the transmitting node.

As indicated in Figure 1a, each device defines its own local coordinate system, referenced to its sensor hardware. In the shown example, device A is sending and B and C are receiving. B and C individually compute a relative position estimate of A, expressed as vector. These relative position estimates each refer to the local coordinate system of the receiver (as opposed to a global reference coordinate system), annotated as index in the vector notation (e.g., BA_B expressing the position of A as vector from B to A, in the local coordinate system of B).

The devices take turns as sender, and over a measurement cycle each determine the relative positions of all other devices within their sensor range. The individual relative position estimates can be shared among the devices, to cooperatively compute more accurate estimates. For example, the six nodes in Figure 1b collectively compute 30 relative positions in a measurement cycle and can combine these to cooperatively estimate their relative spatial arrangement including orientations.

2.2. RELATE Ultrasonic System

Location-based services that help mobile users identify and access devices in their environment have been studied widely. They require fine-grained (10-20cm) positioning accuracy to disambiguate location-based interactions, and information on the relative orientation of users. For indoor environments, ultrasonic sensing has proven to be a relatively simple and effective solution. Existing ultrasonic positioning systems can locate mobile devices with high accuracy but require pre-installation of an accurately surveyed transceiver infrastructure [1]. In contrast, we designed a system in which ultrasound is used for peer-to-peer bidirectional sensing, to support mobile spatial interaction in the absence of any external reference system. Others have employed ultrasound in similar ways for distance estimation between peers [2], but in our system it is also used for angle-of-arrival estimation in order to obtain relative orientation information.

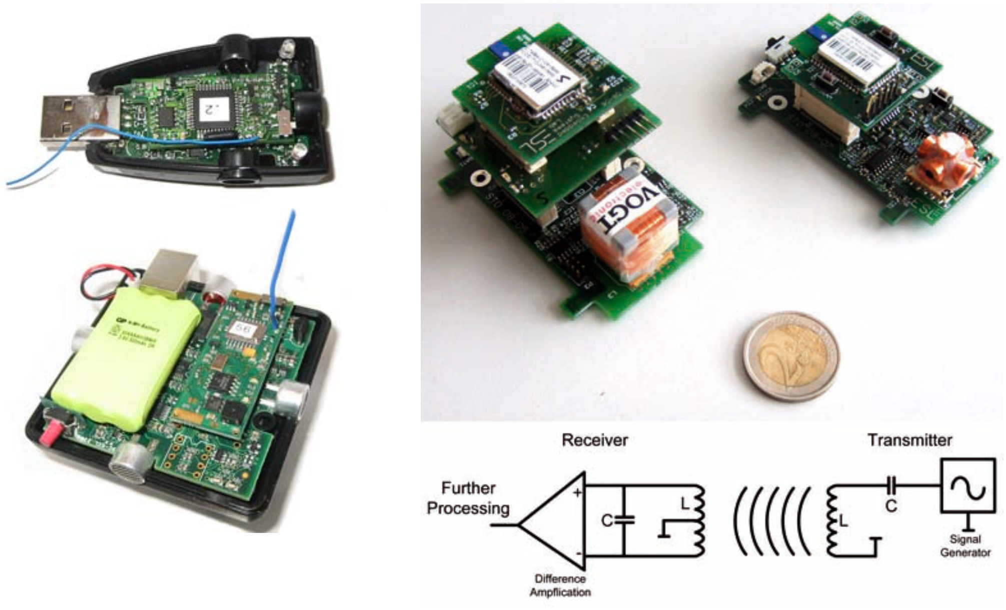

a. Ultrasonic “Dongle” and “Brick” **b.** Relative positioning using magnetic resonant coupling

Figure 2. RELATE Devices and Architecture for Relative Positioning

Figure 2a shows the ultrasonic sensor hardware we developed for relative positioning. The device at the top is a RELATE “Dongle”, designed for ready attachment to mobile host devices such as notebooks and handhelds. The dongle measures 5.5 x 3.5 x 1.5 cm and contains a Particle wireless sensor node, an ultrasonic sensor board with three transducers (to cover front, left and right with 2.5m range), and a USB connector. The device below (named “Brick” to distinguish from the dongle, although of similar small size) has evolved from the dongle to support stand-alone operation, for example as landmark nodes for navigation of mobile users. It employs the same wireless sensor module, but with an on-board battery, a fourth transducer and modified circuitry for larger and omni-directional range (4-5 m).

In sending mode, the devices emit an ultrasonic pulse. The receiving nodes run a peak detection algorithm on all transducers. The earliest peak detected is used to derive a distance measurement from the time-of-flight. The orientation of the transducer on which the peak is detected gives an indication of the angle-of-arrival. If the signal is also received on a neighbouring transducer, then the angle-of arrival estimate is derived from the relative signal intensities. Figure 3a shows experimental results on raw measurement accuracy, obtained with five devices (notebooks extended with dongle) in random configurations, partly with limited line-of-sight between sending and receiving devices [3]. As shown, with good line-of-sight, the measurements are accurate to within 9 cm and 33° in 90% of the cases.

2.3. RELATE Resonant Magnetic Coupling System

For some applications, ultrasound has the disadvantage of requiring free line-of-sight between the transmitter and the receiver. This affects for instance localisation of nodes worn on the user’s body, as the body and user movements will frequently obstruct direct line of sight. However, localisation on the body has interesting potential as a novel approach for analysis of human activity. For instance, the relative position of the user’s hands holds significant information on what they do. More generally, posture and motion can be tracked in terms of the relative position between body parts. This, however, requires very fine-grained accuracy (1 cm) over short range (1 m).

Motivated by the potential for wearable activity recognition, we have developed a new relative position sensing method that does not require line-of-sight and targets short-range applications [4,5]. The method is based on resonant magnetic coupling, known from stationary motion tracking systems, wireless energy transmission and on-body communication, and illustrated in Figure 3b. A transmitter generates a magnetic field that oscillates with certain well-defined, narrow frequency. The receiver contains an LC oscillator circuit (in essence a coil and capacitor) tuned to precisely the same frequency. Only magnetic fields at exactly this frequency can influence the receiver, to the effect that the system is robust against interference from external magnetic fields.

While the general principle is well known, implementing it in a compact, low-cost, low-power way for relative positioning involves a number of challenges [4]. The key issues are:

- Manufacturing tolerances make it difficult to tune the receiver and sender circuits exactly to the same frequency, but even a few percent deviation from the exact resonance can lead to significantly lower efficiency and measurement range. To solve this problem we use a transmitter that “wobbles” around the designated resonance frequency of the receiver. This guarantees that the exact resonant frequency is reached during a measurement cycle.
- The $1/r^6$ distance-dependence of the transmitted energy implies that the system has to work over a large dynamic range (output voltages range from a few V to mV). To deal with this problem our system uses dynamically adaptable amplification.
- For a given distance, the field strength at the receiver depends on the orientation of the receiver with respect to the field and on the location of the receiver around the sender. A given measurement value can thus correspond to many different combinations of distance and orientation. However, it can be shown that by using a perpendicular arrangement of 3 coils in both the sender and the receiver, the relative location of the receiver can be derived.
- The dependence of the angles and distance on the measured field strength cannot be easily expressed in a concise formula. Instead we approximate it with a combination of piecewise approximation and lookup tables.

Figure 2b shows a sender-receiver pair that we have implemented. The devices are tuned to a frequency of 21kHz, achieve a range of around 1m, require around 200mA (receiver and transmitter combined) at 3.3 V and cost around 90$ to produce. The accuracy of the system has been evaluated in systematic experiments using a robotic manipulator to generate positions. As shown in Figure 3b, the distance estimates are accurate to within about 5%. The accuracy of angle estimates is in the range of 10 to 15 degrees. We have also studied the sensitivity of the system to occlusions by objects being inserted between the transmitter and the receiver. As can be seen from Figure 3c, the human body hardly influences the system and even massive metallic objects cause only a small measurement distortion of less than 10%.

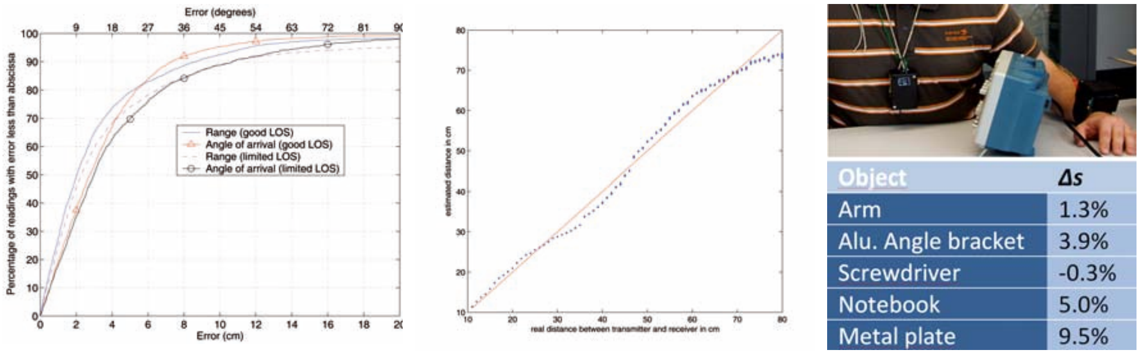

Object	Δs
Arm	1.3%
Alu. Angle bracket	3.9%
Screwdriver	-0.3%
Notebook	5.0%
Metal plate	9.5%

a. Peer-to-peer measurement accuracy of RELATE ultrasonic devices (LOS = line of sight)

b. Distance estimation accuracy of RELATE magnetic resonant coupling devices

c. Distortion of distance estimates by objects between sender and receiver.

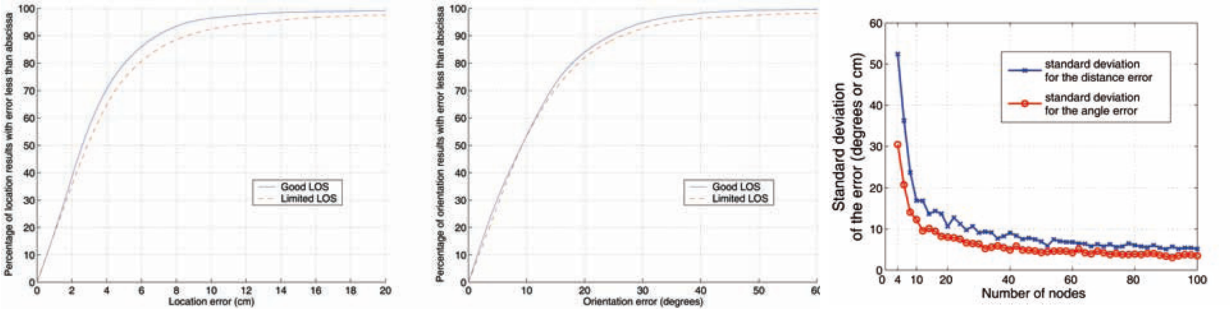

d. Location accuracy (left) and orientation accuracy (right) of cooperative relative positioning with five ultrasonic devices

e. Removal of systematic error with redundant measurements

Figure 3. Performance evaluation of RELATE relative position sensing methods.

2.4. Cooperative Position Estimation

The accuracy of relative position estimates can by improved by fusing individual measurements. Figure 3d shows results with the same data set of ultrasonic measurements used for evaluation of raw measurement accuracy. Here, the measurements of all devices taken over a send-receive cycle are fused for cooperative estimation of the relative device locations, using non-linear regression and studentized residual analysis to remove outlying measurements [3]. Estimates are better than 7cm and 25° for 90% of the results returned in good line-of-sight, and the negative impact of line-of-sight limitations is much reduced.

In further work on cooperative position sensing, we have developed a novel method aimed to remove systematic error by combining redundant measurements [6]. Systematic errors can arise due to sensor decalibration (e.g., misalignment of a transducer) and are common in wireless sensor nodes. Because systematic error is largely dependent on hardware and position of a device at the time of measurement, it is different for each device in the system. Therefore, fusion of simultaneous measurements has the effect that individual systematic errors cancel each other out to a degree, resulting in overall reduced systematic error.

Figure 1b illustrates the method. If devices A and B one after another emit ranging signals, the other devices present can individually and simultaneously compute an estimate for the distance and relative bearing between A and B. For example, D would first estimate the relation position of A and then of B (i.e. vectors DA_D and DB_D), and by subtraction derive an estimate for the vector from A to B (AB_D). With knowledge of the angle-of-emission of the signal from A, D can then perform a rotation of its local coordinate system to that of A. The other devices (C, E, and F) can do the same, resulting simultaneously in four estimates of the relative position of B as seen from A. We have simulated the effect of measurement fusion using error models derived from our ultrasonic sensor data [6]. Figure 3e shows that the standard deviation of measurement error decreases with the number of nodes participating in the scheme. Note the impact, for instance, of scaling up from 4 to 10 nodes for simultaneous measurement.

3. The RELATE Architecture

We have developed an architecture that supports different sensing modalities, to provide a common platform for relative positioning. This includes protocols for cooperative sensing, a device architecture for sensor nodes, and a software platform that mediates between sensors and applications.

3.1. Cooperation and Communication Protocols

The protocols we have developed for cooperative relative positioning are based on AwareCon, a slotted TDMA collision avoidance protocol designed for resource-limited devices. The protocol supports ad hoc discovery and decentralised maintenance of network state. It also provides fine-grained time synchronisation for sensor measurements (e.g., time-of-flight) [7]. For relative positioning, devices access the network on a rota, which they individually derive from the network state. When a device has network access, it transmits network state information, relative position measurements it has taken of other devices while it was in receiving mode, and sensor signals that the other devices use for relative positioning of the sender. Over a send-receive cycle, each device thus takes its own measurements of all other devices, and in addition receives all measurements that the other devices have taken. This enables each device independently to fuse network-wide observations to improve local estimates.

The above protocol is flexible for use with different sensor modalities, and adaptable for context-aware behaviours, for example scheduling network access depending on device mobility. It is limited however in scalability, and insufficient for sharing of redundant measurements at larger scale. To facilitate cooperative schemes for systematic error removal as described above, we have developed a protocol based on synchronous distributed jam signalling (SDJS) for data fusion on the physical layer [6]. In this protocol, a sender announces slots that correspond to measurement ranges. The other devices then reply with a jam signal in the slot that corresponds with their local measurement. This information is then fused based on an estimation theory that takes possible collisions into account [8].

3.2. Hardware and Software System Architecture

We have developed an architecture for RELATE devices that provides a clear separation between the sensing subsystem and the base platform. This enables use of different sensing modalities on the same device core. Multiple sensing modalities can also be integrated and used in parallel, or "hot swapped" dynamically. The base platform can be configured with different wireless communication subsystems. As a reference implementation, we have designed a RELATE "Cube" as a modular stack of boards. The Cube includes a power board as base module, a communication board (with implementations for IEEE 802.15.4 and Particle computer), a DSP board (with a low power dsPIC, basic DSP engine, and 12bit ADC) and sensor boards (with implementations for the ultrasonic and magnetic sensing modalities described above).

The RELATE software system is designed as a modular framework for relative positioning systems. It provides common data abstractions for sensor data (e.g., distance, angle-of-arrival) and data quality (e.g., measurement variance). The framework is implemented in a blackboard architecture that serves as middleware between the sensors and components in the application layer, and the underlying wireless network platforms. Data provided by a sensor is written into blackboard storage for further processing by other modules. This enables integration of heterogeneous sensor types and transparent upgrade or exchange of sensors. A management component on the blackboard supports scheduling and triggering of modules (e.g., to coordinate sensor signalling). The blackboard also provides transparent access to remote data for distributed algorithms [9].

The architecture was validated in the development of three application case studies that each involved different sensing approaches, namely peer-to-peer ultrasound, ultrasound in combination with dead reckoning, and resonant magnetic coupling.

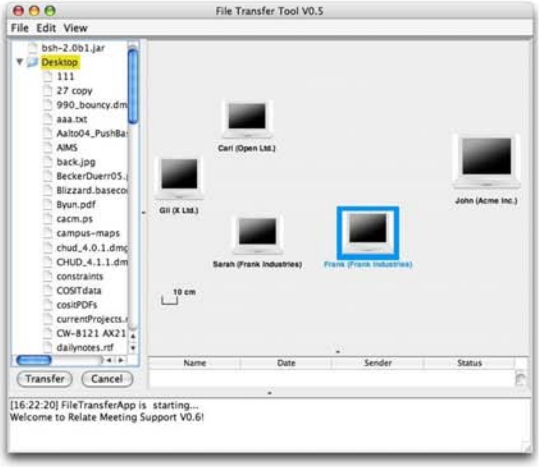

```
[16:22:20] FileTransferApp is  starting...
Welcome to Relate Meeting Support V0.6!
```

a. Supporting users in a meeting scenario with a spatial user interface for direct interaction across their devices.

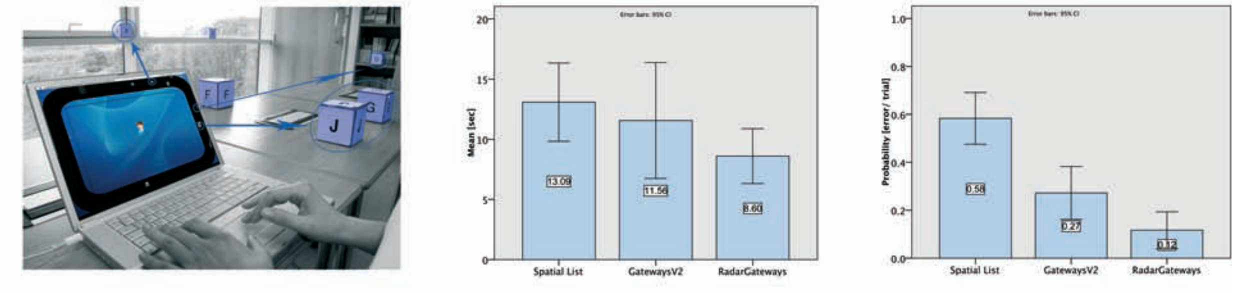

b. Comparative evaluation of user interfaces for selection of objects based on their relative position: experimental setup (left), and results on completion time (middle) and error rate (right).

Figure 4. Supporting mobile users with relative position information for discovery, identification and selection of devices in their environment.

Some content may change prior to final publication.

4. Application Case Studies

4.1. Mobile Spatial Interaction

An application that originally motivated our research is the support of mobile users with spatial interaction mechanisms for discovery, identification and selection of devices they encounter in their environment. Location-based support has previously focussed on providing absolute location information (e.g., room number of a printer location), or coarse-grained proximity (e.g., listing discovered devices in order of their proximity). In contrast, we aimed to provide users with *spatial references* that indicate the relative position of devices from the user's point of view [10].

Figure 4a shows a scenario that we developed for demonstration of spatially supported interaction [11]. Here, mobile users co-located in a meeting can work across their notebooks, supported by a spatial user interface that provides each user with an egocentric view of the meeting situation. In this view, the devices are represented by graphical controls arranged in a spatial layout that reflects their relative spatial arrangement in the meeting. The controls serve as a spatial reference for identification of devices and enable users to access devices without having to know their name or address on the network. For example, if a user in the given scenario wishes to transfer a document to another meeting participant, then they can identify that user's device by their relative position, and drag-and-drop the document onto the corresponding icon in their user interface.

In the above scenario, spatial references were provided in the form of a map view, which is inefficient in terms of screen real estate (esp. on mobile devices). To address this problem we have developed novel spatial user interface concepts. In the RELATE "Gateways" user interface, spatial references are arranged around the perimeter of the screen as gateways to devices located in the corresponding direction [12]. RadarGateways is an extension that in addition indicate the relative distance of the devices. We have conducted extensive user studies to evaluate these user interfaces, in comparison with Spatial Lists as state of the art interface for device discovery. Fig. 4b shows an experimental setup in which users performed a selection task, and results that compare completion times and error rates for Spatial List, Gateways, and RadarGateways. As shown, users are twice as fast with RadarGateways compared with Spatial List, with even more pronounced improvement in error rate.

4.2. Emergency Services "Virtual Lifeline"

A different application, indoor navigation for emergency services, emerged from a collaboration with the WearIt@Work EU Project [13]. In such applications, users (e.g., firefighters) often move in an indoor environment with little or no visibility (due to smoke or dust), no GPS signals, and, in most cases, no other infrastructure that can be readily used for positioning. Attempts have been made to address this problem with pedestrian dead reckoning (PDR), but the irregular movements that are typical for search and rescue (e.g. sharp turns during inspection of rooms, or around obstacles) cause heading errors that such systems cannot compensate [14]. The WearIt@Work project conducted requirements studies with the Paris Firebrigade [15] which led us to develop the concept of a *virtual lifeline,* derived from the firefighters' practice of using physical lifelines as navigation support (see Fig. 5a). The virtual lifeline is conceived as a line of sensor nodes that firefighters deploy as they proceed through an intervention area. By relative positioning from node to node, such a lifeline can track and guide firefighters. Figure 5b shows a demonstrator of the concept, built with Relate Bricks as ad hoc deployed sensor nodes, and sensors fitted into a pair of boots to track the user along the trail of sensors.

A deployment of sensors in trail formation provides only sparse measurements, with little scope for error compensation. We have therefore extended the concept by combining it with PDR. The two technologies are complementary in multiple ways. First, by fusing measurement of the PDR system and RELATE devices, relative position estimates during ad hoc deployment can be improved. Secondly, the user remains tracked when they stray out of range of a sensor trail, or if sensors along the part of the trail fail. Thirdly, RELATE devices can be placed further apart (further than direct measurement range), using PDR to track users in between, and RELATE devices as relative position beacons for re-alignment of PDR estimates. Figure 5c illustrates this with simulation results on user navigation with PDR only (left) and PDR with relative position beacons (right). The path to follow is shown in green, the PDR estimate of the user's trajectory in red, and the user's actual trajectory in blue. A local heading error results in the estimated and actual trajectory to drift apart in the PDR-only condition, but is compensated when relative positioning is used for intermittent re-alignment [16].

a. Members of the Paris Firebrigade using lifelines during intervention training.

b. Firefighter boots augmented with RELATE sensors, and ‘virtual lifeline’ trail of sensor nodes laid for experiments.

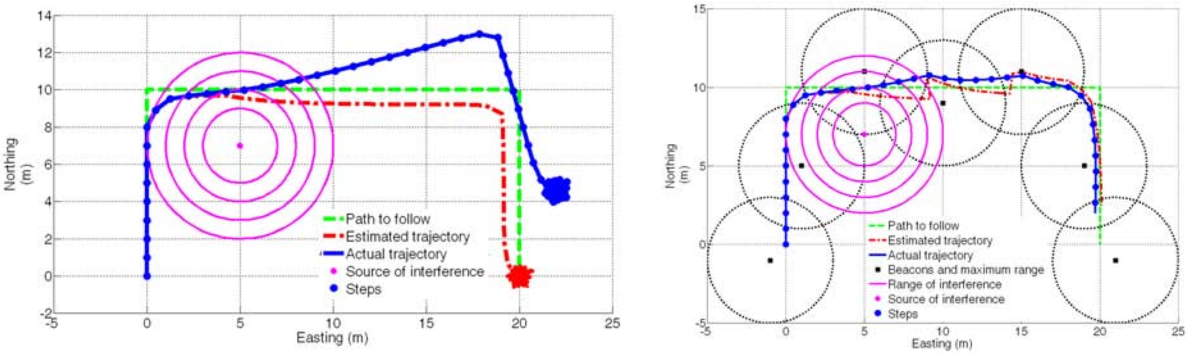

c. Simulation of user navigation with PDR only (left) and PDR with relative positioning beacons (right).

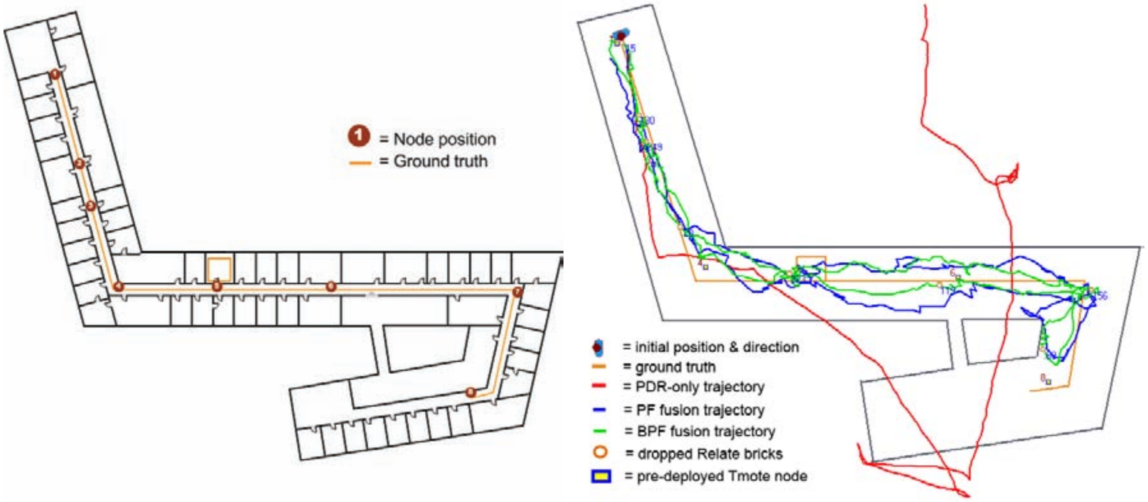

d. Experiments performed in cooperation with the WearIt@Work project: Floor plan of the experimental environment and path used for data collection (left), and prediction of user trajectory (right)

Figure 5. Relative positioning to support navigation of firefighters.

We have conducted extensive experiments to evaluate the use of PDR and relative positioning in real environments. This involved deployment in three different buildings, including a standard office environment and a factory floor as a hostile environment for sensing and communication [17]. A total of five different paths (between 150 and 350m long) were recorded, with up to five different users. We used the Xsens system for

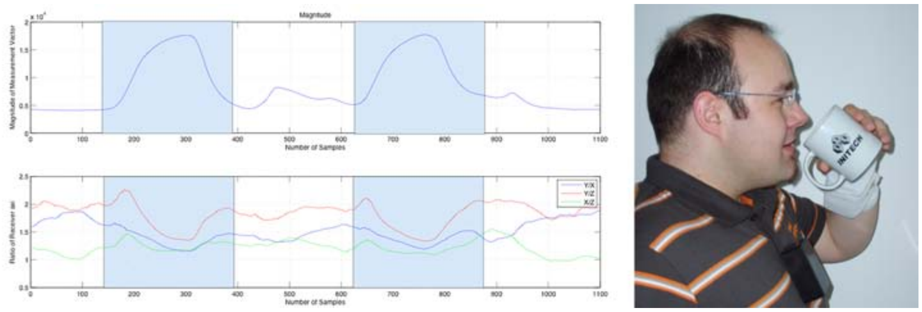

Figure 6. Distance and orientation signals from the Relate resonant magnetic coupling relative positioning node when drinking. Dark areas indicate the drinking motion.

PDR, and RELATE Brick devices as relative positioning beacons placed at distances of 10-20 m. The measurements of the two systems were fused using a fusion engine developed at the Cork Institute of Technology [18]. An example of the results is shown in figure 5d. It can be seen that the PDR-only trace tends to drift off and becomes unusable. By fusion with RELATE measurements, the drift effect is limited. Over the entire data set, we observed that fusion with RELATE reduced the average positioning error of the PDR system by over 50%.

4.3. Wearable Activity Recognition

Most current approaches to wearable activity recognition rely on motion sensors (accelerometers, gyroscopes, magnetic field and combinations thereof), mounted on different body parts. We have developed a new approach using relative positioning based on the RELATE resonant magnetic field technology. It is motivated by the observation that for many recognition problems not the motion as such, but the relative position of body parts (and changes thereof) is the key information. Thus, for example, eating, drinking, and brushing one's teeth all imply proximity between the hand and the mouth with subtle differences in the position and motion between them.

We have performed an initial experiment involving the classification of six gestures – two concerned with nutrition (eating, drinking) and four similar but unrelated gestures. Figure 6 illustrates the signals observed during a drinking motion. All gestures involved a similar hand posture and differed in the wrist placement by less than 10cm. Thus the recognition problem was not trivial. With a sensor setup of receiver on the chest, and emitter on the forearm, a performance of 99.3% correct classification was achieved [4]. The approach was also tried for recognition of movements in martial arts. An experiment on a set of 8 Tai Chi moves was conducted with a combination of magnetic sensing and accelerometers. This work demonstrated a significant improvement in recognition rate in comparison with systems that use accelerometers only. Specifically, with accelerometers only we observed 86% correct recognition, with magnetic sensing 94%, and with the combination of both 100%.

5. Conclusion

There are two broad lessons to be learned from the work described in this paper. The first concerns technology. Although relative positions can be derived from absolute coordinates, there are many situations where dedicated relative positioning systems are a more appropriate choice. They work in environments where global positioning infrastructure is not available, and can be optimised to meet specific requirements of relative positioning applications such as high accuracy over short range, and availability of orientation information. We have shown that cooperation and coordination between the interacting entities are key aspects of relative positioning systems are and described an architecture that implements these aspects in a modular, technology-independent way.

The second lesson concerns the applications. While the work was originally motivated by obvious domains such as human interactions in pervasive computing environments, new and unexpected applications have emerged, with the emergency lifeline and wearable activity recognition as examples. Overall we argue that cooperative relative positioning is an important fundamental concept underlying a broad range of pervasive computing applications.

6. Acknowledgements

This research was supported by the Commission of the EU (FP6 contract 013790, "RELATE"). We gratefully acknowledge the contribution of our colleagues and project partners, in particular Mike Hazas, Albert Krohn, Gerald Pirkl and Carl Fischer who led some of the reported research.

References

1. A. Ward, A. Jones and A. Hopper. A new location technique for the active office. *IEEE Personal Communications 4, 5* (Oct. 1997), 42–47.
2. M. Minami, Y. Fukuju, K. Hirasawa, S. Yokoyama, M. Mizumachi, H. Morikawa, T. Aoyama. DOLPHIN: a practical approach for implementing a fully distributed indoor ultrasonic positioning system. *Proc. Ubicomp 2004*, 347–365.
3. M. Hazas, C. Kray, H. Gellersen, H. Agbota, G. Kortuem, and A. Krohn. A Relative Positioning System for Co-Located Devices. *Proc MobiSys 2005,* 177-190
4. G. Pirkl, K. Stockinger, K. Kunze and P. Lukowicz. Adapting Magnetic Resonant Coupling Based Relative Positioning Technology for Wearable Activitiy Recogniton, *Proc. ISWC 2008,* 47-54.
5. M. Barry, A. Grünerbl and P. Lukowicz. Wearable joint-angle measurement with modulated magnetic field from l/c oscillators. *Proc. 4th Int. Works. on Wearable and Implantable Body Sensor Networks (BSN 2007)*
6. A. Krohn, M. Hazas and M. Beigl. Removing Systematic Error in Node Localization Using Scalable Data Fusion. *Proc. EWSN 2007: 4th European Workshop on Wireless Sensor Networks*, 341-356.
7. A. Krohn, M. Beigl, C. Decker, and T. Riedel. Syncob: Collaborative time synchronization in wireless sensor networks. *Proc. (INSS'07: 4th Intl. Conf. on Networked Sensing Systems*, 283–290.
8. A. Krohn, M. Beigl, and S. Wendhack. SDJS: Efficient Statistics in Wireless Networks. *Proc. ICNP 2004: 12th IEEE international Conference on Network Protocols*, 262-270.
9. M. Beigl, M. Beuster, D. Rohr, T. Riedel, C. Decker and A. Krohn. S2B2: Blackboard for Transparent Data and Control Access in Heterogeneous Sensing Systems. *Proc. INSS'07,* 126–129.
10. H. Gellersen, C. Fischer, D. Guinard, R. Gostner, G. Kortuem, C. Kray, E. Rukzio and S. Streng. Supporting Service Discovery and Spontaneous Interaction with Spatial References. *Personal Ubiquitous Comput.* 13, 4 (May. 2009), 255-264.
11. G. Kortuem, C. Kray, C. and H. Gellersen. Sensing and visualizing spatial relations of mobile devices. *Proc. UIST 2005: ACM Symp. on User interface Software and Technology*, 93-102.
12. D. Guinard, H. Gellersen and S. Streng (2007) Extending mobile devices with spatially arranged gateways to pervasive services. *Works. on Pervasive Mobile Interaction Devices (PERMID), Adj. Proc. Pervasive 2007.*
13. P. Lukowicz, A. Timm-Giel, M. Lawo and O. Herzog. WearIT@work: Toward Real-World Industrial Wearable Computing. *IEEE Pervasive Computing 6, 4 (Oct. 2007), 8-13.*
14. C. Fischer and H. Gellersen. Survey of Location and Navigation Support for Emergency Responders. *IEEE Pervasive Computing, Special Issue on Location-based Services.*
15. M. Klann. Playing with Fire: User-Centered Design of Wearable Computing for Emergency Response. *Proc. MobileResponse 2007: First Intl. Works. on Mobile Information Techn. for Emergency Response,* 116-125.
16. C. Fischer, K. Muthukrishnan, M. Hazas and H. Gellersen. Ultrasound-aided pedestrian dead reckoning for indoor navigation. *Proc. MELT'08: First ACM Int. Works. on Mobile Entity Localization and Tracking in GPS-less Environments,* 31-36.
17. Widyawan, G. Pirkl, D. Munaretto, C. Fischer, C. An, P. Lukowicz, M. Klepal, A. Timm-Giel, J. Widmer, D. Pesch and H. Gellersen. Virtual Lifeline: Multimodal Sensor Data Fusion for Robust Navigation in Unknown Environments. (in submission to *Pervasive and Mobile Computing*).