Available online at www.sciencedirect.com
SCIENCE DIRECT®
Computer Communications 28 (2005) 1575–1585
computer communications
www.elsevier.com/locate/comcom

Redundant positioning architecture

Tom Pfeifer*

Telecommunications Software and Systems Group (TSSG), Waterford Institute of Technology (WIT), Cork Road, Waterford, Ireland

Received 31 August 2004; accepted 15 December 2004
Available online 11 May 2005

Abstract

Positioning has been a driving factor in the development of ubiquitous computing applications throughout the past two decades. Numerous devices and techniques have been developed—few of them are actually used commercially.

The precision is limited to specific applications, the availability limited to the provider of specific services. Occasionally, few methods have been combined to recalibrate each other by means of data fusion.

We present a novel architecture for processing the vast amount of data from pervasive devices penetrating everyday objects to the cheapest level. Location information can be inferred from infrastructure deployed for different purposes, only partly designed for positioning in the first place.

The massive redundancy of such nodes and the synergetic heterogeneity of completely different recognition principles allows to tailor the perceived positioning probability to the specific requirements of the target application. A self-learning and self-healing approach to misleading, wrong and outdated pieces of information provides a new quality of interworking position and context aware systems.

Keywords: Positioning; Location awareness; Context awareness; Distributed objects

1. Introduction

Location Awareness in general describes applications in computing and telecommunication, which alter their behaviour in dependence on the location of an entity. The latter might be the user of the application, a person the user of the application wants to communicate with, or an object capable of changing location. Such location represents a major category of *context* [9], and is derived by various methods of *positioning*.

To achieve tradeability of contextual information among different organisational domains, it requires qualitative parameters regarding precision, trustworthiness and timeliness, recently phrased as *Quality of Context* [2].

Authentication is an issue implicitly related to certain methods of positioning, employing the same technology.

With Ubiquitous Computing thriving recently, dozens of applications demand such location information. It has sufficiently been discussed that localisation indoors has different technological and topological requirements than outdoors. Even if we could derive a probable geographical positioning of one meter indoors, this left us uncertain if we were at the one or the other side of the wall between two rooms. Thus, a different approach is required in local environments.

Our novel architecture massively *combines* numerous individual positioning technologies to obtain more precise and more reliable results according to the various needs of the whole range of location based services.

1.1. Categories of localisation technology

In the relation of localisation and communication [31], we can distinguish two major categories, seen from the perspective of the mobile device:

The first category is *receptive* localisation. The position information is distributed ubiquitously, and the mobile device can derive its own location from this information. Satellite distributed GPS is a typical example. The mobile device can either relate the derived location to a map independently and provide a local service (without revealing its own position to any third party), or purposely transmit this location in order to obtain a value-added service.

---

* Tel.: +353 51 30 2927
E-mail address: t.pfeifer@computer.org.

0140-3664/$ - see front matter © 2005 Elsevier B.V. All rights reserved.
doi:10.1016/j.comcom.2004.12.042

The second category is *transmissive* localisation. The position is derived by a fixed station which either sees the mobile device or receives a beacon from it. The station can then either transmit the derived location information to the mobile device, or use it to generate the value-added service for the mobile user. Sub-cell GSM positioning is the example, where the beacon of the mobile communication channel is also used for positioning. At cell level, receptive location is possible in this example also, when the mobile device can identify the very cell it is in.

Hybrid techniques combining the principles above are possible.

However, despite years of experimentation in the labs, very few positioning technologies have currently significant economical impact—most prominently GPS. Most other technologies, for indoor as well as outdoor, have niche markets only, if they are commercialised at all.

Reasons include high costs compared to the added value achieved, or immature precision and reliability.

1.2. Interworking localisation

Progress in the latter issues is expected in most recent proposals to combine heterogeneous positioning data from different technological sources, in order to obtain a higher probability for a certain position scan by principles of data fusion [7]. The number of different parameters overlapping into a final position assessment is still small.

However, an increasing number of technologies suitable for positioning are becoming available. We are at the advent of the penetration of everyday objects with pervasive devices to the cheapest level. Reliable, camera-based visual tracking and shape/face recognition becomes very feasible with falling costs of the imaging components. Biometric devices might be available in every office. Wireless and wired sensor networks of a variety of categories can detect presence and proximity of people and objects. Position information can be derived from sources not previously designed for this purpose—such diverse as triangulations in wireless communication networks (Wifi, GSM), sightings of campus cards at cash registers, usage of IP or MAC addresses at certain wired network patches, etc.

The interworking of all these systems will provide a *massively redundant approach* of positioning and a new quality of context awareness. An *architecture* harnessing the parallel data streams from all these sources for positioning is presented in this paper, leading to a unprecedented level of precision and well-structured confidence in the data presented to the respective applications.

The heterogeneity of the recognition principles being included, used in a synergetic manner, allows to tailor the perceived positioning probability to the specific requirements of the target application. This means in an example, that for setting some room or music preferences in a smart environment, a recognition probability in the upper 80% would be sufficient, while paying money from an ATM would require much higher confidence in the authentication.

The massive redundancy of such nodes implies that not every piece of information might be correct, it might be wrong in the first place, temporarily misleading or be out of date. This aspect has been considered in our architecture, and self-learning, self-adapting and self-healing approaches keep it alive. Support for manageability is provided.

Within this paper, we briefly discuss in Section 2 some existing technology and current approaches of fusion, with regards to the above goal in mind. Section 3 introduces our redundant positioning architecture. In Section 4, the required research and implementation is outlined and Section 5 summarises the paper.

2. Related work

2.1. Impact of individual technologies

Research into positioning systems has focused around leveraging location information from various media such as sound, infrared, radio and vision. Each system has addressed the aggregation of sensor data into location estimations via their own specific methods.

The following collection of related localisation technology is not meant to provide an exhaustive analysis, but merely a collection of examples in order to discuss the pros and cons of each approach, and in particular to categorise them in regards of their current economical impact.

2.1.1. Widespread commercial systems

The list of widely deployed positioning services with everyday visibility is very short.

Radar for locating ships and aircrafts is the longest used positioning system, however, with no impact expectable for ubiquitous computing.

The Global Positioning System (GPS) [12], as receptive radio localisation, is currently the most reputable location system, widely used in navigation and open-area logistics. GPS receivers can compute their location anywhere on the planet to a high degree of accuracy. Well defined protocols for transferring location information from receivers to other devices have been established, thus making it easy to integrate with other systems.

Contactless inductive transponder based positioning and visual on-the-fly bar-code recognition are widely used in factory hall logistics processes in industry (e.g. tracking items on, or delivered to, an assembly belt), delivery services (parcel carrier, etc.), and road toll collection.

Cell based GSM location is explicitly used for its primary purpose of enabling communication.

2.1.2. Location inference

Another class of possible positioning are technologies not primarily developed for this purpose, but bearing implicit information which can be derived by inference of proximity.

Examples are readers (at fixed positions) for magnetic swipe cards or chip cards, used for access control and time registration, for money-related functions (credit cards), or combined functionality in a certain area (campus cards).

BlueTooth devices, can be recognised in proximity to each other.

Further, network addresses (IP, MAC) can be mapped to positions. On a global scale, the mapping of blocks of addresses assigned to countries and institutions is already used for countrywise or citywise localisation of web-services. In terms of a company, the internal patch-database provides office-level granularity.

2.1.3. Laboratories and niche markets

Lots of positioning technologies have been developed in the labs, and to some extent commercialised for a number of niche markets.

Passive *RFID transponders* becoming cheaper. E.g. wrist-worn devices are available for child-location and payment functions in theme parks [27]. Extremely cheap RFID tags (few cents per item) are expected in particular from the AutoID initiative of the Automatic Identification and Data Collection industry [29], significantly driven by global consumer product manufacturers and retailers targeting to individually serial-number any item in the supply chain with the Electronic Product Code (EPC) [30].

Sub-cell GSM positioning [32]—as an example for transmissive radio localisation and as demanded in the 'E-911 mandate' [33] and similar European requirements—is currently too immature, requires still relatively expensive mobile devices and/or hardware investment on the provider side, and cannot distinguish room boundaries. The diversity of network providers—binding services to this specific provider—hinders, e.g. the company-wide usage of location services where employees have different mobile contracts.

Indoor room-level positioning, requiring a different kind of granularity than geographical coordinates out-door, lead to a number of approaches exploiting Infrared, Ultrasonic sound, and local-range Radio Frequency (RF).

RF *localisation* in general, for a long-time, could not provide the dynamic granularity to know on which side of the wall the object is (10 cm), or the proximity (e.g. 1 m) to a particular Point of Interest (PoI). Bahl et al. (Microsoft Research) [10] describe the 'RADAR' approach based on existing WLAN infrastructure, analysing the signal strength already measured in any WLAN card (RSSI: Receiver Signal Strength Indication). Within the coverage of multiple base stations for triangulation, they achieve a resolution of 2–3 m, without requiring extra hardware beyond the WLAN, derived from either a mathematical or an empirical radio-mapping of the building. Weaknesses of this method derive from the propagation characteristics of its radio frequency. A privacy bonus is the fact that the location is known to the receiving mobile only, leaving the dissemination to the discretion of the user.

AeroScout is a commercial WLAN based location system from Bluesoft. It employs additional custom hardware to remotely detect and locate WLAN enabled devices [28]. The system uses a central server to collect, process and expose location data.

Ultra-wideband (UWB) radio technology [13] currently becomes a promising technology for RF positioning, as the extremely short pulses can be easily distinguished from multi-path receptions. Most recently, *UWB* has been demonstrated in positioning approaches with centimeter-wide precision by UbiSense [14].

Modulated Infrared (IR) light based technology provides advantages such as the restriction of signals within rooms (IR does not pass walls) and absence of electromagnetic interference. The signal power can easily be adjusted to cover small areas only.

Olivetti Research (ORL) and Cambridge University, UK, pioneered with the nowadays 'classical' approach of the (transmissive) Active Badge [15,16] infrared sensor system. Designed, prototyped and trademarked by ORL between 1989 and 1992, a small device worn by personnel, transmits a unique IR signal every 10 s.

Commercial follow-ups, like the EIRIS Infrared Localisation System from ELPAS Electro-optic Systems Ltd, Raanana, Israel [17], provide more reliable sightings due to advanced use of diffuse infrared transmissions. The sensitivity of the IR receiver had been pushed to the limit by carefully shielding the diode and amplifier with meshed copper and complex signal regeneration—thereby also increasing the price, significantly.

The general problem consists of battery power being the scarce resource in the badge, allowing limited power of the IR emission only. Sensitivity of the wall-mounted receiver is crucial—but expensive.

An approach to overcome these limitations has been reported from Fraunhofer FOKUS [18], converting from transmissive to receptive badges. Wire-powered, arbitrarily strong, stationary infrared beacons blast their ID in the room, which is received by the badge and, more battery-efficiently, communicated via RF signals to a small number of base stations in the building.

All these badges are currently used in niche markets and labs only, also suffering (to a different degree) from their intentional or accidental concealability, sometimes considered as a privacy advantage, but unsuitable for security applications. Environmental conditions, such as direct sunlight, often interfere with the signals.

The transmission of *ultrasonic sound* within rooms led to a number of interesting prototypes. The basic idea is, compared to GPS, to move to a slower carrier (speed of sound instead of speed of light), to make the received timing differences more measurable in a room environment. Strong dampening effects of walls again limit the signal to the desired room-coverage. Again both, transmissive and receptive approaches are found.

Randell [11] compares the timing of four ultrasonic signals received from four transmitters within a room, while the timing of the signals has been transmitted via RF, and calculates the relative position. This approach is two-way receptive, in the ultrasonic band as well as in the RF. The localisation is very precise (10–25 cm). However, the cost of setting up the ultrasonic transmitters, which must be networked for synchronisation, allows the deployment in selected rooms only.

The Cricket Compass [21] also receives ultrasonic signals accompanied with RF location IDs. Five receivers in the mobile device form a dedicatedly shaped layout. It then can triangulate from ceiling beacons and derive location as well as orientation (‘compass’) information.

In conversion of the principle, the Active Bat (AT&T labs, the former ORL) [20] transmits ultrasonic pulses from the mobile device, received from a matrix of ceiling microphones. It obtains a resolution of 8–10 cm.

From the commercial aspect, the price of such ultrasonic installations would limit application to a small number of mobiles attached with the technology.

Weight or force measuring on the floor is another laboratory approach. It starts with the ORL Active Floor [19], analysing the movement of people and carried objects by Markov Chains, and Georgia Tech’s Smart Floor [24]—still requiring expensive installations.

Currently, semiconductor manufacturer Infineon prototypes, in cooperation with the Vorwerk Teppichwerke carpet plant in Hameln, Germany, a ‘smart carpet’, within their ‘intelligent textile materials’ programme [25]. A self-organising network of robust chips allows to embed sensor functions, such as monitors for pressure, temperature or vibration, directly into woven fabric materials. Linked by a wired grid it allows to cut the material arbitrarily.

Wireless sensor networks—similarly to the wired carpet above—are a hot research topic currently, and might be used to sense the proximity of people and objects (e.g. by temperature or chemical impacts), as their main purpose or as a side effect. While from the current perspective they might not be deployed as primary means of positioning, they could contribute to the redundancy approach discussed later in this paper.

The *visual* channel is the most developed one for human orientation—and most challenging for technology.

Visual localisation and tracking based on advanced stereoscopic image processing is a very promising approach [22,26]. Falling costs for small cameras and increasing processing power make it feasible that such systems might take the place in the ceiling corners of the room, where nowadays the motion sensors of the burglar alarm reside. Privacy issues are of course most debatable within this area.

Concluding this section, none of the developments alone as discussed above would fulfil requirements such as affordability and precision targeted to specific applications. To provide a solution for a situation where the number of localisation areas (office/meeting rooms, multiple PoIs in exhibitions, campus) has the same magnitude than the number of persons or objects to be localised, further approaches have to be considered.

2.2. Location fusion approaches

Many researches in the area of location sensing have come to the conclusion that location information needs to be centrally managed. Some of these researchers have come to the added conclusion that one location system is not enough for reliable positioning and that the aggregation of data from multiple location systems is necessary for robust person and device positioning [34].

A common approach to this issue is the use of a layered model to transform specific data from individual location systems into generic location information. Part of this transformation is the use of Sensor Fusion techniques to combine information form multiple heterogeneous sensor networks.

Sensor Fusion is defined by Hightower/Borriello [4] as ‘the use of multiple location systems simultaneously to form hierarchical and overlapping levels of sensing … to increase accuracy beyond what is possible using any individual system’.

This group at Intel Research has proposed the standardisation of location systems for ubiquitous computing through a six layered model that incorporates location sensor fusion as well as context information fusion [5,6,8]. The system is proposed as a replacement to ‘monolithic’, specific location systems.

Similar work by Leonhardt/Magee from the Imperial College, London [7] results into a three layer stack that can then be layered on to itself. This recursive approach presents a way of mapping the stack to all aspects of a location information system from sensors to clients.

Enhancement of location estimations by evaluating previous, historical data have rarely been discussed. In RADAR, Bahl [23] reports improvements by using short-term, history-based user-tracking to follow a user by evaluating the likelihood of a new position depending on the immediate previous ones.

A completely different perspective is described by Ferscha [1] in terms of the Digital Aura of objects, allowing for spontaneous interaction, based on detected experience of the same context—and certainly people can develop such aura as well. Inherently, such aura considers historical data for its self-awareness.

Many approaches discuss the individual error, precision or confidence level of the respective technology, but wrong and misleading information is not being considered. Typically in the fusion approaches, an overall confidence level is calculated—abstracting from the underlying input technology.

3. Redundant positioning architecture

We assume that the whole avalanche of ubiquitous computing devices of most different kinds that is set to arrive in our fives can be exploited for positioning purposes—even if not originally built for this—and supported by a number of dedicated positioning systems. It is the combination of all of them, delivering small pieces of information into a mosaic representation of the real life—Fig. 1 gives some illustration.

Naturally, not all of this information will be correct, as different classes of errors occur, leading to different degrees of imperfection:

- Uncertainties in the actual position are mostly discussed in the literature, they are caused by the limited precision of the respective measurement, and can vary within the same technology, depending on the actual Dilution of Precision (DoP), e.g. depending on satellite constellations, RF multi-path reflections, etc.
- Expectations of how long an object is still in proximity after a position scan might be based on typical assumptions. On a footpath, the typically expected speed would be pedestrian; or a location scan at the cash register of the cafeteria would assume a typical eating time of 20–30 min. But as exceptions, objects can be faster, and the lunch could have been interrupted.
- Relations between objects change, incidentally or intentionally. Labels fall off, the RFID tagged sweatshirt—so far a typical characterisation of a certain person—has been handed to the sister, or the cash card has been borrowed from a colleague to buy some coffee in the canteen.

If we compare these issues with ordinary human life, then we have learnt to recognise people by a large number of individual characteristics. Facial properties, size, typical glasses worn, the style of clothing, voice, and many more. Even if one of these properties does not match, we recognise due to the redundancy of the other, compensating misleading information, and we adjust in a learning process some properties to new values, e.g. the new hair cut or other clothing.

When we meet new people, this learning process is obvious. First, we lay out the business cards in front of us in the seating order, but soon we hear the people speaking, and individualise them by behavioural characteristics—sooner for strong characters, after some reinforcing cycles of further meetings for the other.

3.1. Characteristics of redundant positioning

Applying this behaviour to a technical system leads to the proposed approach of *redundant positioning* with a vast amount of sensing nodes contributing to the whole image, including the consideration of imperfect, misleading and wrong information from individual nodes, and the accumulation of history data from previous sightings for the learning process.

3.1.1. Redundancy

As discussed in Section 2.1 there is a large number of positioning systems and methods available. Each system has established characteristics including the overhead involved in installing its infrastructure support (hardware and/or software) and expense to procure the system. Extensive physical installation or high costs may limit large-scale deployment to key areas. Other systems, such as those based on Wifi, may be easily deployed on large-scale, but are limited in accuracy, and may be able to locate a sub-set of devices only. When every possible source of location data is pooled, greater coverage of an area can be achieved. As well as coverage issues, multiple sightings from different location systems can be used to reinforce or discount a user's estimation.

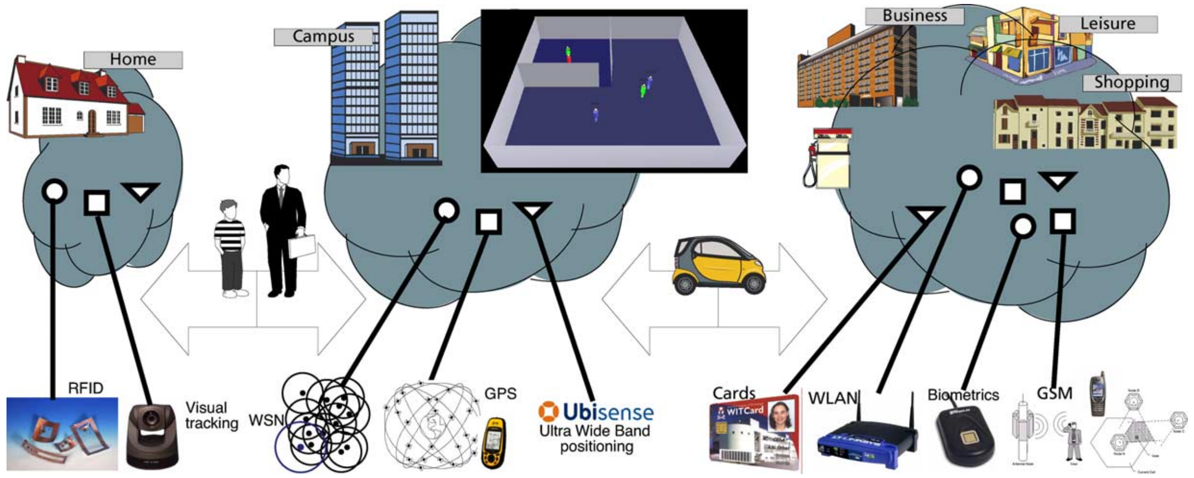

Fig. 1. Positioning across domains.

3.1.2. Temporal validity

It can be argued that every sighting has a temporal validity attached to it. For example, a device sighted by a WiFi location system will be valid for a small period of time (1 or 2 s). A user sighted using his debit card to buy a cup of coffee from a till in a coffee shop could be valid for a number of minutes (the time taken to drink the coffee). By attaching temporal validity data to a sighting the system can maintain location estimations between sparse sightings.

3.1.3. Confidence

The confidence is a measure of the potential error in a location estimation. It is not just a distance measuring error—it could also be an error in the perceived carrier of a located device. In other words, it is geometrical precision vs. detecting the wrong person. Confidence is important as applications potentially using the location information may reject it if the confidence is not high enough. As there are different classes of confidence (or errors, as discussed above), giving an application just a numeric value of confidence, e.g. in percent, is not sufficient when the type of possible error is relevant.

3.2. Input channels

No possible input channel should be excluded from the system approach, even if implementation will be limited by cost and pragmatic reasons. Naturally, different combinations of such channels will be exploited at different locations. Modularity is the key issue in this part of the architecture. Listing a number of possible channels below, therefore, provides examples in the context of redundancy, not completeness.

3.2.1. GPS

The satellite based positioning is embedded in an increasing number of devices such as PDAs and phones. It provides an interesting relevance for indoor positioning when the sudden absence of a GPS signal coincides with the last reception at the entrance of a building—this a valuable piece of information triggering the transition from the outdoor model into indoors.

3.2.2. RFID

Readers for RFID in gate situations like doors detect objects and persons passing, RFID readers in container objects (desks, briefcases, bags) detect their content, RFID readers at furniture and desktop inventory detect proximity of people and associated objects.

3.2.3. Visual tracking

Cameras in corners of the room track recognised contours (which are then personalised, if not by the visual technology itself, by mapping to read RFIDs at specific points).

3.2.4. Floors

Weight-measuring floors, or the more economical Infineon carpet, qualify the trajectory observed by other positioning.

3.2.5. Wireless sensor networks (WSN)

Measuring proximity of people (e.g. temperature) by wireless sensor networks might be used alternatively to the floors. Alternatively, RFID readers might become so small and inexpensive that they can be deployed as self-organising WSN.

3.2.6. Active badges

The badges are an option for places where the obtrusiveness of these devices counts as positive argument (e.g. factory visitors), and the maintenance cycles can be fulfilled.

3.2.7. Biometric sensors

Providing a high (certainly not absolute) level of confidence, biometrics can be used punctually, when positioning requires that level of authentication.

3.2.8. Ultrasonic and ultra-wideband positioning

As long as the installation cost per room can be kept reasonable, the high precision obtained from this principles can be helpful in calibrating other channels in the redundancy approach, as discussed in Section 3.4.

3.2.9. Wifi positioning

Using additional hardware provides more precision, but already the software-based approach alone is of major relevance for the redundant positioning approach, given the rapidly increasing installation base and coverage of wireless LAN.

Already the cellular nature of an installation (supported by some more sophisticated approaches discussed in Section 2.1) provides good coarse positioning, well suitable to be complemented by other technologies.

3.2.10. GSM and beyond

Cell and sub-cell positioning is an inherent feature of mobile communication, however, the large business models behind the competing telecoms demand for exclusive exploitation of value-added location based services. This makes a cooperative approach difficult and requires a change of thinking.

3.2.11. Location inference

All software approaches based on already existing systems might be considered as location inference. Typically, these methods depend on the cooperation of the owner/operator of the specific system. Typical examples are: recognising devices at wired ethernet patches, the proximity of Bluetooth enabled devices, the Wifi positioning already discussed, and the usage of swipe or smart cards

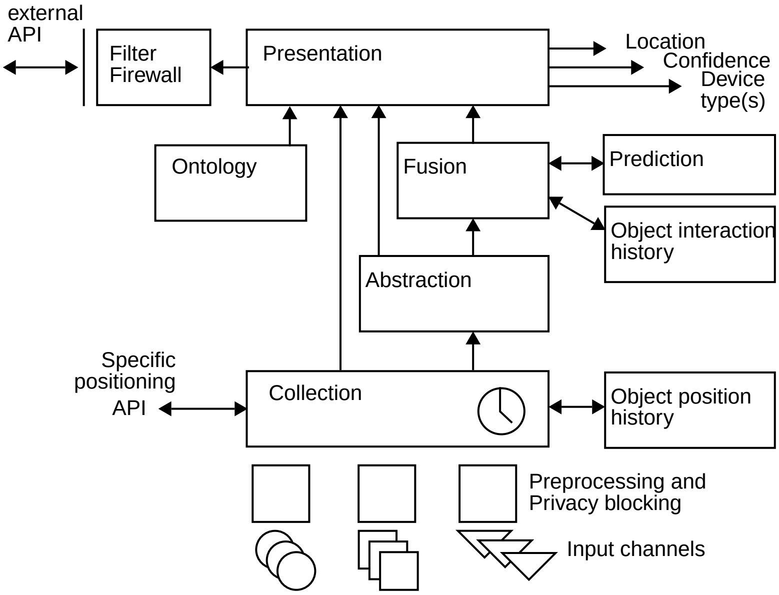

Fig. 2. Architecture for redundant positioning.

at stationary stations. While usage at bank ATMs should be withheld for security reasons, a more relaxed policy might be acceptable for multi-purpose campus cards (library, access control, petty cash).

Obviously, the best effects are expected from the most different kinds of principles of positioning, thereby being most complementary.

3.3. Architecture

In Fig. 2, an architecture for the approach of redundant positioning is presented, and discussed from bottom to top below. It features an arbitrary number of different *input channels*, each representing a large number of sensing devices of a common positioning technology, as discussed above.

The heterogeneity of these input channels requires *preprocessing* tailored to the specific technology. E.g. a technology generating large numbers of sightings might apply an early aggregation of data, outputting the counting of same or similar scans. The video tracking or face recognition might have the tracking and recognition process encapsulated there, outputting room vectors of sightings, or extracted standardised visual features.

In the preprocessing, *privacy filters* can be implemented. So I might switch a video tracking camera in a mode prohibiting face recognition or apply temporal restraints. RFID readers could contain a blacklist of items not to be scanned (or, more restrictively, a positive list of allowed items, like office property). The Wifi positioning software might respect non-sighting escape zones, e.g. the canteen.

The goal of that early filtering is to avoid protected data to accumulate in history buffers appearing in upstream modules.

The pre-processed data are forwarded to the *collection* module. Its most important task is to provide a ‘plug-in’ socket for each input channel. Two basic types of modules can be distinguished:

- a module that uses a specific location systems API to collect data from it,
- a module that has an API that a location system uses to pass in data.

Scans will be *time* stamped here, if they have not been stamped in their originating positioning system. This time stamp is crucial to compare data from different sources, and for applications to determine if there is a location determination lag.

The format of the data is not unified at this point—it would be normal to expect multiple, incomparable formats.

On the collection level, a *history of object positions* is important to be kept. This preserves valuable data coming from individual location systems for later comparisons, detection of typical patterns and routines. Beyond this point data can/will be modified.

The *abstraction* module changes the location data (if necessary) into a site specific form as a default operation. The primary goal is to present the data in a unified, object-oriented view. The abstraction module consists of elements that transform between data types. A limited number of transformers can be chained together to perform any transformation, avoiding the need for *n*:*n* individual transformers.

The *fusion* module combines data from multiple sightings on multiple location systems. Employing the provided redundancy at this level leads to potentially improved accuracy. The fusion should take into account that some location systems have a built-in level of accuracy. Potentially, the fusion module could have different fusion algorithms—and could chain them to perform multiple fusions—in order to provide different perspectives to the presentation layer, which then can tailor the perspective to the heterogeneous applications.

At the fusion level, it is possible to store a *history of object interactions*, i.e. preliminary assignments of objects to other objects and objects to people, to observe regular patterns and routines, e.g. what items a person regularly wears, what valuables typically go with which owner.

This history is the prerequisite for the *learning* process, which then can lead to *predictions* of probable interactions in the future. Prediction is a supplement to fusion. It uses history information to make predictions of where a user should be or is going. People tend to retrace their steps due to routines they may have or sheer habit. Prediction could aid accuracy and provide for interesting applications such as predicting that a doorway or hall will be very busy in the near future and, therefore, should be avoided.

The *presentation* level must provide everything an application could possibly need. In coding location based services it is often equally important to locate a device as well as the person using it. Previously proposed approaches of fusion, as discussed in Section 2.2, obscure/abstract location data to a level where useful information is lost (e.g. the type of devices that have been located and have been part of the fusion process, or the reason for assigning a certain confidence level).

This hinders the application to ask different questions, e.g. the question “Where is Tom?” might be answered by an aggregation of sightings of items he regularly wears. Having such sightings fused with details lost would make it impossible to answer Tom’s question “Where is my PDA?”. Avoiding this problem, our approach allows the presentation level to communicate with pre-fusion levels, such as collection and abstraction.

Patterns of joint movement are part of the object interaction history, and can be employed by applications where the authorisation of movement can be observed—e.g. a car key moving with somebody else than the owner, or appearingy alone, could raise attention.

On the output side, beside the location itself, the confidence level, and the type of confidence expressed in the type of the providing device are presented.

So far, the system has been restricted to (timed) location as a subclass of context. Whether introducing other subclasses and background knowledge of the surrounding world—as indicated with the *Ontology* module—improves the positioning, has to be investigated.

For any external access, a *filter and firewall* module restrict access to an allowed amount of information, which can be adjustable to the degree of authentication the querying party provides. Further details about security and privacy are discussed in Section 3.6.

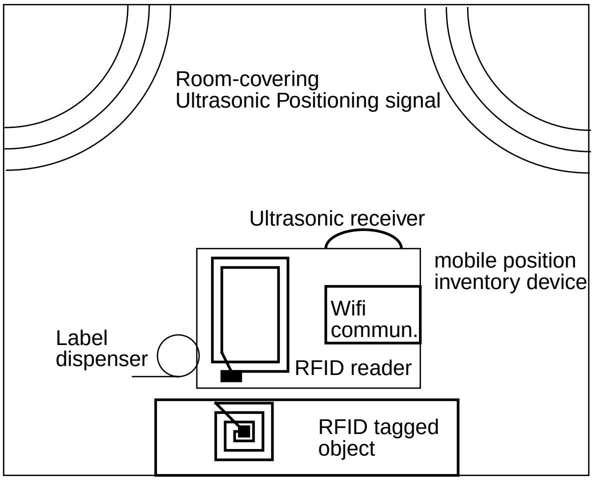

Fig. 3. Mobile Position Inventory Device: location aware RFID dispenser.

3.4. Location calibration

When positions are derived by inference from more or less stationary equipment, it has to be inventarised in respect of its precise location in the room. This would be a tedious process when done manually and raises questions of manageability: nobody would like to type lots of positions in databases. A location aware registration tool, e.g. in the form of an RFID label dispenser, would be very helpful.

Combining a precise indoor positioning device, e.g. Ultrasonic/UWB positioning (such as Cricket compass [21] or Active Bat [20], UbiSense [14]) or a laser range-finder, with a near-range RFID reader results in a position inventory tool, as depicted in Fig. 3.

The RFID, already on the object and thereby assigned to the name of the item, or freshly dispensed on location, is read, while simultaneously the precise position within the room is determined. Via the Wifi interface, the location and if necessary the just inventarised item description is communicated towards a database.

Re-calibration of the positions of equipment after rearrangements is just a short walk through the office, touching each item with the tool.

3.5. Inter-domain communication

It is logical to split geographic regions up for the purposes of ease of management, for example, local government or councils. The same logic applies to location information systems. By making a location information system responsible for a specific area it limits the number of devices it has to constantly locate, thus reducing the stress on the system. Other location information systems can then be established to cover other geographic locations, or other property domains. If the systems are aware of each other and can query device locations from one another, a peer to peer location information network forms, as shown in Fig. 4.

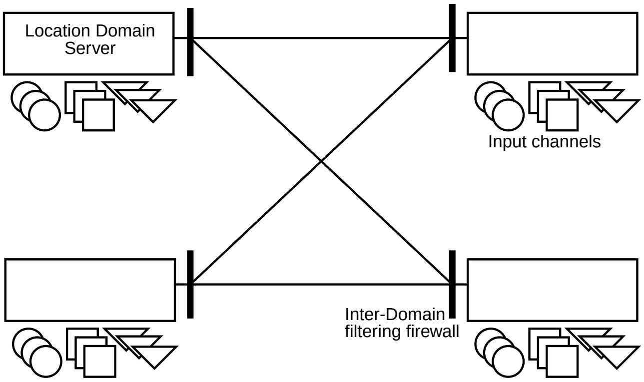

Fig. 4. Inter-Domain location Information exchange.

Thereby, they operate as a group of 'loosely coupled cooperating domains'.

The information exchange requires a message format and protocol, which considers the heterogeneity of the data being exchanged. Examples for such data possibly exchanged are:

- Geographical, room-relative or map-specific positions,
- characteristics of objects (RFID, visual image suitable for recognition, visual sub-image, e.g. logo, extracted visual features in the recognition process, measurements, voice pattern, acoustical features of the voice in the recognition process, etc.)
- biometric data sets,
- meta-information about the objects or object interactions (e.g. by whom typically carried).

As the communication might comprise different administrative domains, different degrees of private or public services have to be considered. Therefore, a filtering firewall allows to tailor the data to the authorisation level of the other party.

3.6. Security and privacy

Security issues arise in particular when location information is exchanged between independent location domains. Privacy and ethical implications have become a topic of hot discussions [3] with few solutions so far. Location information firewalling and filtering, as well as early avoidance of collecting certain data are used to solve these issues.

The vast amount of information being processed for the redundant positioning might be acceptable within an environment of open communication, i.e. within a department, a small company, a family, or among friends. For any external questions, this information needs to be filtered, and only explicitly specified details, or highly aggregated data would be provided.

The home domain (home office or residential home server) of a user might, e.g. publish some identity information about the user on the level of an ID card, comprising a public photograph, some biometrics, a cryptographic hash of his RFID number, and a digital signature. Further information, e.g. about objects typically worn/carried might be limited to certain trusted domains, or withhold completely in respect to relevance and privacy. Aggregated position information might be externally provided, e.g. during business hours. The filter rules can be adjustable for differently authenticated users (friends, cooperating partners, strangers).

The filtering and blocking algorithm built into the preprocessing before data collection avoids unwanted data to accumulate in the system.

4. Research and implementation

Section 3 has described principles of a new quality of highly redundant fusion for positioning and authentication in location aware systems. This approach is being followed from the theoretical as well as the practical side.

4.1. Data model

A location is much more than a set of geographic coordinates. Coordinates can only contribute as a method to uniquely identify locations [37]. Each location usually has a role such as representing a position in a specific street, being the entrance to a building or a meeting room. In indoor positioning any location can be defined by more than its coordinates as room, floor and building, giving a complete address, or relative positions within the room.

The semantic of location information depends on the application domain. The area of interest determines in

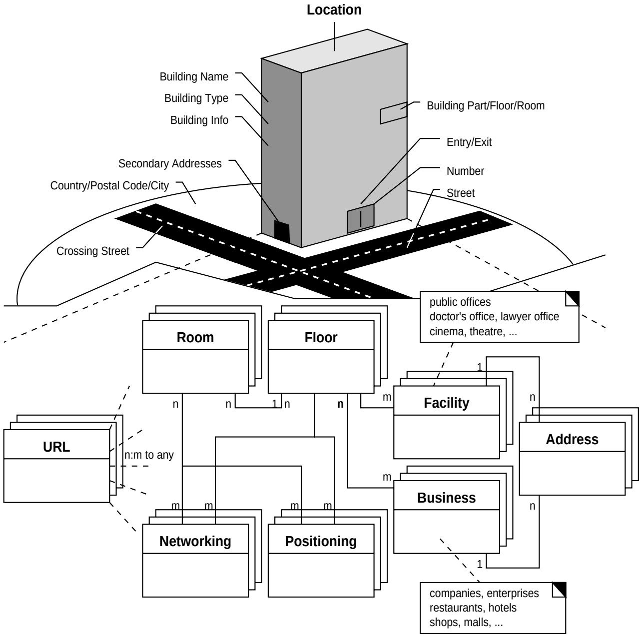

Fig. 5. Area of Interest and Position data modelling illustration [35].

which form the final position is being expected. A number of geographical data models—as illustrated in Fig. 5—have been discussed in literature. For example, [36] breaks ‘Location’ into four levels of granularity, namely *Position* (a single fixed place of an object, measured in relative or absolute coordinates), *Address* (a structured, textual description how to find a Location), *GeographicArea* and *Structure* (representing buildings, floors, rooms, and other entities).

In any location system, regardless of the type of coordinates used, there is likely to be a set of impossible coordinates. These represent areas where it would—be physically impossible for a device to be. Knowing that certain coordinates can be discounted will aid accuracy and influence decisions.

The data model to be developed needs to consider the possible requirements of the application as well as the heterogeneous sources of positions from the redundant supply.

4.2. Testbeds and experiments

As part of the NOMAD project [38], a wide scale infrastructure to support mobile applications has been deployed on the WIT campus. This infrastructure includes Wifi positioning technology that is both software and hardware based. So far it has been used for psychological experiments of group interaction among mobile device users, equipped with augmented PDAs.

Each access point in this network is being polled for the list of devices currently associated to it. This gives very general device sightings for general areas within the campus. At certain spots of interest, the more accurate hardware based Wifi triangulation provides higher precision.

Based on a existing partnership with the mobile GSM/GPRS operator O_2 Ltd in Ireland, sub-cell location data can be obtained for an area of the city including the campus. Given the nearly complete penetration of GSM/GPRS handsets among the student body, a significant amount of data contributes to the system.

RFID reader technology is being installed at numerous points around the campus, RFID tags are being dispensed in a large number (to handsets of participating students, to office property, to teaching material, etc.).

Visual tracking cameras are being deployed in selected labs only for current privacy concerns.

A dual Magnetic Swipe and Smart Card card of the WIT campus is currently being used for access control, library check-out and petty cash card. The debiting is processed in a central database, which will deliver data about the specific cash terminal used as positioning information.

The patch-panel database where fixed private IP addresses are assigned to room locations is being used for login-based positioning at workstations.

Exploiting all sources discussed above will provide the critical mass of redundant data to implement and test the system architecture as described in Section 3, and preliminary experiment with the data flow, while more theoretical aspects of the approach are being investigated.

4.3. Analysis, modelling and simulation

The redundant positioning approach processes a vast amount of data. For the algorithms in fusion and decision, for the continuity and the timeliness of the arriving data, appropriate models are being designed and evaluated in a mathematical analysis, and tested in simulations, before they can finally be implemented. Some approaches of handling the data stream can also been derived from the Auto-ID consortium, e.g. the Savant software [29].

The modelling also needs to determine how other context classes can contribute to the improvement of positioning.

5. Summary

Analysing technology for positioning available commercially and in the laboratories from a usability perspective, an architecture has been designed that allows to combine an arbitrary number of dedicated positioning systems and location inference methods to contribute in a massively redundant way. This allows the reliable determination of positions of objects and persons, at the time present and in predictable near future, as well as their interactions and the elimination of wrong and misleading information by self-learning, self-adapting and self-healing.

The presented architecture has purposely been restricted to positioning (including aspects of recognition and authentication), as a subclass of context. To what extent the interaction with the other subclasses of context would increase the usability of the system has to be investigated.

References

[1] Alois Ferscha, Coordination in pervasive computing environments, Proceedings of the 12th IEEE International Workshop on Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE’03), 2003.

[2] T. Buchholz, A. Küpper, M. Schiffers, Quality of context: what it is and why we need it, Proceedings of the 10th HP-Open View University Association Workshop, Geneva, Switzerland, July, 2003.

[3] Jürgen Bohn, Vlad Coroama, Marc Langheinrich, Friedemann Mattern, Michael Rohs, Social economic, and ethical implications of ambient intelligence and ubiquitous computing, in: W. Weber, J. Rabaey, E. Aarts (Eds.), Ambient Intelligence, Springer, Berlin, 2004.

[4] Jeffrey Hightower, Gaetano Borriello, A survey and taxonomy of location sensing systems for ubiquitous computing. Technical Report UW-CSE 01-08-03, University of Washington, Department of Computer Science and Engineering, Seattle, WA, Aug. 2001.

[5] Jeffery Hightower, Gaetano Borriello, Dieter Fox, The location stack. White Paper IRS-TR-03-008, Intel Research, Seattle, 2003.

[6] David Graumann, Walter Lara, Jeffrey Hightower, Gaetano Borriello, Real-world Implementation of the Location Stack: The Universal

Location Framework, Proceedings of the 5th IEEE Workshop on Mobile Computing Systems and Applications (WMCSA 2003), 2003.

[7] Ulf Leonhardt, Jeff Magee, Multi-Sensor Location Tracking, Proceedings of the 4th ACM/IEEE International Conference on Mobile Computing and Networking, Dallas, TX, 1998 pp. 203–214.

[8] Jeffrey Hightower, Barry Brumitt, Gaetano Borriello, The Location Stack: A Layered Model for Location in Ubiquitous Computing, Proceedings of the 4th IEEE Workshop on Mobile Computing Systems and Applications (WMSCA 2002), Callicoon, NY, 2002 pp. 22–28.

[9] Albrecht Schmidt, Ubiquitous Computing—Computing in Context. PhD thesis. Schöneiche, Germany: Multicon, 2003, ISBN: 3-930736-04-7.

[10] P. Bahl, V. Padmananbhan, RADAR: An In-Building RF-based User Location and Tracking Systems, Proceedings of IEEE Infocom Tel-Aviv, Israel, 2000.

[11] C. Randell, H. Muller, Low Cost Indoor Positioning System, UbiComp 2001: International Conference on Ubiquitous Computing, Atlanta, Georgia, USA; Springer LNCS 2201, 2001.

[12] Steve Poizner, Karissa Todd, Extending GPS capabilities, Wireless Review 16 (1999) 9 Overland Park, KS, 1 May 1999.

[13] W. Hirt, Ultra-wideband Radio Technology: Overview And Future Research, Computer Communications, vol. 26, Elsevier, Amsterdam, 2003. 1; pp. 46–52.

[14] Ubisense Ultra-wideband (UWB) positioning technology, http://www.ubisense.net/, 2004.

[15] A. Harter, A. Hopper, Distributed location system for the active office, IEEE Network 8 (1994) 1 Jan/Feb. 1994, IEEE Computer Society, pp. 62–70.

[16] R. Want, Andy Hopper, V. Falcao, J. Gibbons, The Active Badge Location System, ACM Transactions on Information Systems 10 (Jan) (1992) 1.

[17] EIRIS Infrared Localization System, System Manual, ELPAS Electro-optic Systems Ltd, Raanana, Israel, Jan 1999, see also: http://www.elpas.com/.

[18] T. Pfeifer, D. Elias, Commercial Hybrid IR/RF Local Positioning System, KiVS'2003, Kommunikation in Verteilten Systemen, Leipzig, Germany, Feb 2003, in: K. Irmscher, K-P. Fähnrich (Eds.), KiVS 2003. Kurzbeiträge, Praxisberichte, VDE Verlag, Berlin, Offenbach, 2003, pp. 119–127. (ISBN 3-8007-2753-6).

[19] M.D. Addlesee, A. Jones, et al., The ORL active floor, IEEE Personal Communications 4 (1997) 5 IEEE, New York, Oct. 1997, pp. 35–41.

[20] A. Ward, A. Jones, A. Hopper, A new location technique for the active office, IEEE Personal Communications 4 (1997) 5 IEEE, New York, Oct. 1997, pp. 42–47.

[21] Nissanka Priyantha, Allen Miu, Hari Balakrishnan, Seth Teller, The Cricket Compass for Context-Aware Mobile Applications, Proceedings of 7th ACM MOBICOM, Rome, Italy, July, 2001.

[22] Trevor Darrell, Vision-based Perceptive User interfaces. Presentation on September 30, 2001, Workshop on Sensing and Perception for Ubiquitous Computing at UbiComp 2001, International Conference on Ubiquitous Computing, Atlanta, Georgia, USA, 2001.

[23] P. Bahl, V. Padmananbhan, A Software System for Locating Mobile Users: Design, Evaluation, Lessons, Microsoft Research Technical Report TR-2000-12, Feb. 2000.

[24] Gregory D. Abowd, Robert J. Orr, The smart floor: a mechanism for natural user identification and tracking. Extended Abstracts of the 2000 Conference on Human Factors in Computing Systems, The Hague, The Netherlands, ISBN 1-58113-248-4, pp. 275–276.

[25] Infineon: News.-http://www.infineon.com/cn/news/20030715-01.htm (Jul. 2003) and http://www.infineon.com/cn/news/20040114-01.htm,Jan.2004.

[26] Ryan Cassel, Christophe Collet, Tracking of Real Time Acrobatic Movements by Image Processing, Gesture-Based Communication in Human–Computer Interaction: 5th International Gesture Workshop, GW 2003, Genova, Italy, April 15–17, 2003, Springer, LNCS 2915/2004, ISBN 3-540-21072-5.

[27] SafeTzone Child locating system. http://www.safetzone.com/.

[28] AeroScout Location system.-http://www.aeroscout.com/, 2003.

[29] AutoID standards summary, http://www.autoid.org/standards/project_Summary.htm, March 2004.

[30] Electronic Product Code (EPC) Network: EPCglobal, http://www.epcglobalinc.org/, 2004.

[31] S. Long, R. Kooper, G.D. Abdowed, C.G. Atkeson, Rapid Prototyping of Mobile Context-Aware Applications: The Cyberguide Case Study, Proceedings of the 2nd ACM International Conference on Mobile Computing and Networking, MobiCom, 1996.

[32] 3GPP Technical Specification Group Radio Access Network. Stage 2: Functional Specification of Location Services in UTRAN, Technical Specification 3G TS 25.305 V3.1.0, 3rd Generation Partnership Project (3GPP/TM), March 2000, http://www.3gpp.org.

[33] Tim Kridel, E-911, mandate of opportunity, Wireless Review 15 (1998) 23 Overland Park, KS, 15 Nov. 1998.

[34] Radu Popescu-Zeletin, Tom Pfeifer, A Modular Location-Aware Service and Application Platform, Proceedings of The Fourth IEEE Symposium on Computers and Communications, ISCC'99, Red Sea, Egypt, 1999.

[35] Mario Schuster, An Open Service Platform for Context and Location Aware Applications, Diplomarbeit (MSc Thesis), Technical University of Berlin, 2001.

[36] Shared Information/Data (SID) Model, Addendum 5PR-Physical Resource, Business Entity Definitions, GB922-Addendum-5PR, Member Evaluation Version 3.0, NGOSS Release 3.5, June, 2003.

[37] OpenGIS, Location Services (OpenLS): Core Services, Open GIS Consortium, 2004, http://portal.opengis.org/files/?artifact_id=3418, http://xml.coverpages.org/OpenGIS-LS03006r3.pdf.

[38] NOMAD project: SToRC (Dundalk IT), CCTA (Dun Laoghaire IADT-DL), TSSG (Waterford IT), Ireland, http://www.nomadireland.org/.

Tom Pfeifer had joined the TSSG research group of Waterford Institute of Technology (WIT) as a Principal Investigator, Director of research in Ubiquitous and Pervasive Computing, in 2004, funded by the Science Foundation Ireland. As visiting professor, he teaches at WIT as well as at Universities in Limerick and Dublin.

From 1992 until 2003 he had worked in Berlin in a variety of senior research, industry consulting, project management and lecturing positions in the Fraunhofer Institute of Open Communication Systems (FOKUS), the GMD Computer Science research labs, and the Technical University of Berlin, Faculty of Computer Science and Electrical Engineering.

Following lectures in the 1990s in different areas of High Speed Networking and Multimedia Technology, he had recently developed a module of lectures and seminars on Ubiquitous and Pervasive Computing.

Tom Pfeifer received his Dr.-Ing. (PhD) in Computer Science at Technical University of Berlin, 2000. He holds a Dipl.-Ing. (Masters) degree in Electrical Engineering / Electronics from the Technical University of Dresden, Germany, 1989.

He has published over 30 technical papers, is member of the editorial boards of the Journals Computer Communications and Pervasive Computing and Communications.

He has been invited to numerous scientific talks throughout Europe, and served in more than a dozen Programme Committees.