

An Integrated Platform for the Management of Mobile Location-Aware Information Systems

Anthony Savidis^{1,2}, Manolis Zidianakis¹, Nikolaos Kazepis¹, Stephanos Dubulakis¹,
Dimitrios Gramenos¹, and Constantine Stephanidis^{1,2}

¹ Institute of Computer Science, Foundation for Research and Technology – Hellas

² Department of Computer Science, University of Crete

{as,zidian,kazepis,dubulak,gramenos,cs}@ics.forth.gr

Abstract. We present an integrated platform comprising a set of authoring and management tools for mobile location-aware information systems. The development of the platform was targeted in supporting large-scale systems with very crowded use sessions, at the scale of hundreds of simultaneous visitors, addressing information delivery for exhibits with proximity down to a couple / few meters. The key platform features are: (i) spatial content editing with mixed-mode administration, either mobile (on-site with a PDA) or non-mobile (off-site, using a PC); (ii) system-initiated location-triggered information delivery combined with free user-initiated data exploration; (iii) applicable both indoors and outdoors; (iv) very efficient device renting processes through barcode readers; and (v) multiple location sensing technologies, prioritized according to precision trust (includes WLAN, GPS, and infrared beacons). Currently, the platform is being installed at the fifteen main museums and archeological sites of Greece (including Acropolis, Olympia, Delphi, Knossos and Mycenae), encompassing a total of five thousands mobile devices (see acknowledgements).

1 Introduction

Mobile location-aware information systems are capable to deliver position-dependent information over a portable device for users that are primarily on the move. In our work we consider multiple sources for user-location information, including explicit positioning as coordinates in a 2D plane (e.g. X , Y and optionally a DIR vector) or implicit positioning as the result of interpretation data coming from sensory modules. The latter is managed by various technologies like special tags (e.g. infrared or radio), sonar methods and computer vision. The former is possible mainly via GPS, for outdoor environments, and WLAN positioning engines, relying on signal fingerprints, for both indoor and outdoor setups. We adopted the dual use of explicit and implicit positioning, while supporting the statically-prioritized (i.e. invariant during runtime) use of location sensing technologies based on the known precision of the generated location data. The reported system is currently being installed at the fifteen major museums and archeological sites of Greece (see acknowledgements), encompassing a total of five thousand mobile devices (PDAs).

1.1 Contributions

The key technical contributions of the reported work, concerning the support for content administration, user navigation, and runtime management of user sessions, are listed below:

- Spatial location-oriented data editing with a direct-manipulation editor, offering mixed-mode administration: (a) mobile, on-site, mainly to administer location data; and (b) non-mobile, off-site, mainly to administer typical semantic content.
- System-initiated location-triggered information delivery combined with user-initiated on-demand content exploration.
- Unified infrastructure applicable both to indoor and outdoor setups.
- Very efficient and intuitive device renting facilities through barcode readers.
- Multi-channel statically-prioritized location sensing, currently implemented to deploy WLAN positioning, GPS, and infrared beacons.

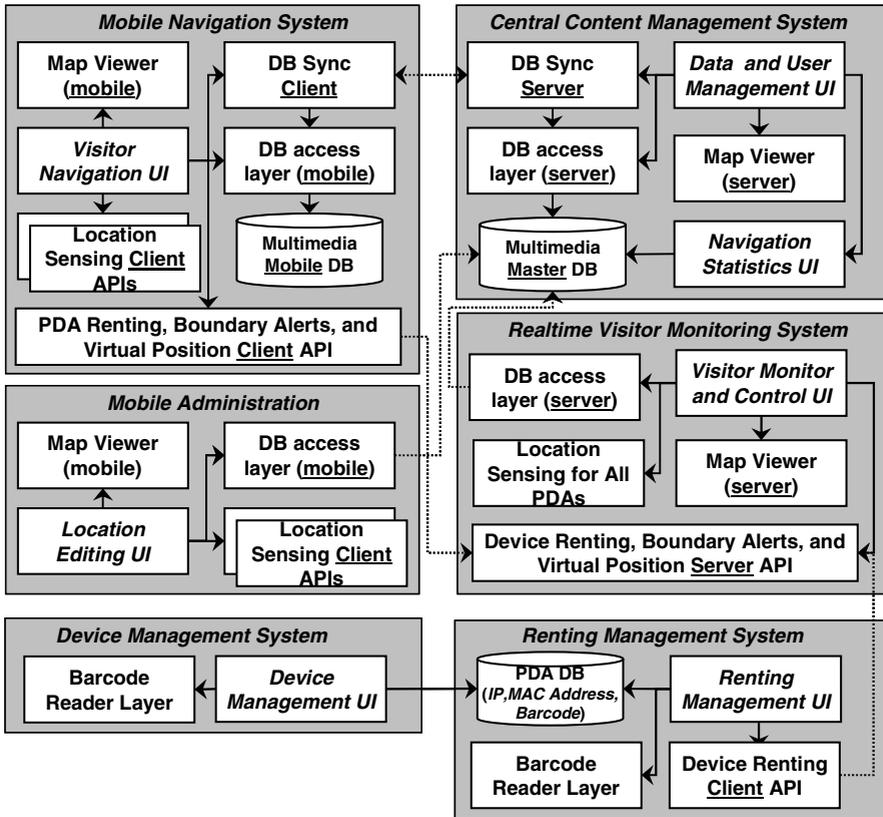


Fig. 1. Overall platform architecture

1.2 Architecture

The overall system architecture is provided under Fig. 1 at two levels of decomposition detail: (i) primary applications and tools (shaded rectangles); and (ii) their respective key constituent components (white rectangles and buckets). Solid arrows indicate intra-process method invocations (e.g. deployment of *Map Viewer* component), while dashed arrows denote inter-process RPCs (e.g. *DB Access Layer*). Also, with underlined labels we mark different versions of a component, such as: mobile / server DB

Access Layer, mobile / server Map Viewer, client / server API categories, and mobile / master DBs. Some other important sub-systems, not to be elaborated in this paper, are also incorporated in the architectural diagram of Fig. 1:

- On-demand synchronization of the master DB to the mobile DB on the PDAs (i.e. *DB Sync Client / Server*). While one expects this facility to be offered by the DBMS, we had to implement it from scratch as the mobile edition of the MS SQL Server restricts the size of the mobile DB to at most 100 MB.
- Recording data for navigation sessions (e.g. time spent at an information item, information items reviewed, elapsed time of use, etc.) supporting various queries (e.g. most popular information item in a selected period, total time of use for all visitors, average exploration time for visitors, information items with simultaneous visitors up to a threshold, etc). The latter concern the *Navigation Statistics UI*.
- Security policy relying on the definition (during content administration) of the legal navigation boundaries over area maps, out of which specific alert procedures are triggered (e.g. PDA locking with message to the user, explicit alert pop-ups on the monitoring console, visitor identification from the PDA and notification from the museum audio system, etc.). The latter concern the *Visitor Monitoring System*.

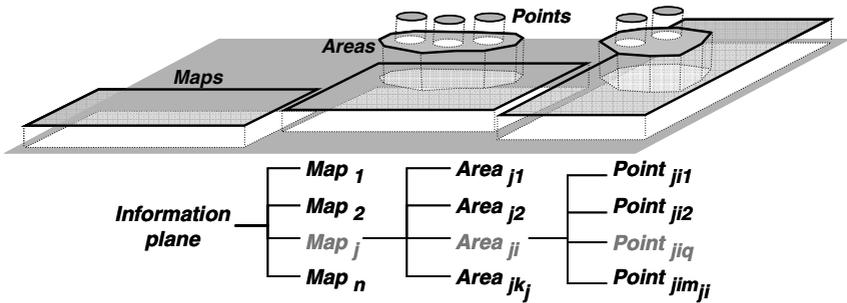


Fig. 2. Split of the information plane to maps, areas and information points

Our platform relies on a generic location-sensing interface enabling sensing APIs to be loaded dynamically (as DLLs). Our method splits the global information plane into independent maps, where maps encompass polygonal information areas, which enclose the actual information points, i.e. the real exhibits (see Fig. 2). In this context, the location sensing interface allows distinct technologies to return either a point (higher precision) within the currently active map plane, or alternatively the identifier of a polygonal area (lower precision) that is associated to a particular physical area (e.g. room, hall, corridor, corner, stairs, etc.). Following our approach, the adopted technologies need not be merely point-based, but alternative techniques like infrared beacons or radio tags can be deployed as carriers of area identification information (e.g. an infrared beacon can be programmed to simply emit the logical identifier of an area denoting a specific room).

Effectively, when the latter is combined with typical point-based sensing technologies it helps to resolve possible ambiguities or precision problems, or overcome other types of practical barriers. The capability to support multiple prioritized sensing technologies is crucial, since, in some situations, the particular characteristics of

the installation site, or other types of restrictions regarding possible physical interventions, may directly exclude specific technologies. For instance, in our project, we were forbidden to put any type of equipment (i.e. no tags, no power supply) inside archeological sites, except of the wireless network antennas perimetrically installed. Additionally, the prioritized use of location sensing APIs reflects the differing reliability of the position outcome in different technologies.

2 Related Work

Location sensing technologies are the cornerstone of location-aware information systems. Some technologies support indoors localization through specialized, expensive and usually inflexible hardware installations. A typical recent example is the Active Bat system (Harle et al., 2003) relying on signal distance measurements to determine user's position. C-MAP (Sumi et al., 1998) was among the first exhibition-tour systems exploiting the Active Badge System for location awareness. Other systems based on distance signal measurements are the Cricket Location Support System (Priyantha et al., 2000), Spot-On system (Hightower et al., 2000), and EasyLiving (Krumm et al., 2000), the latter with vision methods via motion-tracking cameras.

Recent location sensing systems rely on standard wireless networking hardware, by measuring signal intensity and attenuation to determine a user's location. The RADAR system (Bahl & Padmanabhan, 2002) was amongst the first tracking systems based on IEEE 802.11, while in (Savarese et al., 2002) an algorithm is increasing the precision of the estimated users' positions. A commercial system in this category is the Ekahau Positioning Engine¹ (EPE) supporting laptops, PDAs and other Wi-Fi enabled devices, accomplishing floor, room, and sometimes door-level accuracy, while working indoors and outdoors. In our platform, we deployed EPE as one of the point-based location sensing technologies. The official specifications of the EPE indicate that it is capable to reach an average precision 1 meter, though we observed that in real practice the average precision is around 3 meters. Besides point-based methods, technologies like passive RFID tags or active infrared beacons support proximity-based localization as carriers of land-marking information. Such tags require explicit installation and programming, while they are not very intuitive in use: (a) infrared requires line of sight; and (b) passive radio tags that do not need power-supply require very close proximity, sometimes virtually like physical contact or touch as in MoVIS (Schwieren & Vossen, 2007) - the reason is that radio antennas on PDAs are not powerful due to low-power demands enforced to enable autonomous PDA use for a few hours. Advanced systems like PlaceLab (LaMarca et al., 2005) support metropolitan area coarse-grained positioning (at the level of 20 meters), relying on the variety of preinstalled stationary beacons, ranging from WiFi antennas to Bluetooth tags, roughly estimating users' position in indoor / outdoor urban areas. A probabilistic method to fuse results from various types of sensors for location estimation, relying on particle filters, is proposed in (Hightower & Borriello, 2004); the technique displays increased accuracy as the number of sensors is increased, while having outcomes close to the most accurate of the sensing technologies deployed.

¹ <http://www.ekahau.com>

In our platform, targeting for a unified solution across indoor and outdoor settings, no single technology suffices. Our emphasis is put to enable integration of solutions with acceptable estimations, rather than to introduce a new method for location sensing. Currently, outdoor information systems lack genericity (no indoors), practicality (satellite signal can be lost) and precision at the level of information points (coarse-grained positioning). On the other side, most indoor navigation systems employ tag technologies, like infrared beacons or radio tags, with known issues as previously mentioned. The latter technologies implement context-sensing rather than position tracking, so user monitoring and trajectory recording is minimal. Naturally, better positioning may be gained by placing a large number of tags around; however, this is hardly a globally acceptable option.

Employment of multiple location-sensing technologies has been originally proposed in (Nord et al., 2002), encompassing virtually all sorts of location sensing methods. They propose a resolution policy based on area intersections, however, it is not mentioned how ambiguities are resolved (i.e. when the intersection is empty) and they do not encompass a prioritization of technologies according to trust on localization precision. Bluetooth is also proposed, known to suffer from severe problems: unacceptable stalls upon device discovery (scan inquiry takes around 15 seconds), signal traveling through walls causing ambiguous results (Savidis & Stephanidis, 2005), and significant power demands.

A high-level architecture for redundancy of adopted technologies to derive positioning information is suggested in (Pfeifer, 2005). The idea is to exploit virtually any type of medium from which positioning information can be extrapolated. Our approach is closer to this philosophy, however, not targeted on algorithms to fuse signals for deriving better estimates (i.e. enhanced location sensing), but on ways to make systems deploy reliable sensing technologies with known trust values with the least implementation dependencies (i.e. enhanced software engineering).

3 Mobile Navigation System

The key novel features of our visitor's navigator system to be discussed are: (i) prioritized *multi-channel location sensing*; (ii) combined *system-initiated and user-initiated navigation*, supporting predefined tours; and (iii) smooth map zoom using 3D rendering with *interactive control on the information detail*. The deployment of a WLAN positioning system providing continuous affine location estimations together with RFID tags to get more precise localization is adopted in (Blache et al., 2003), although in the reporting RFID integration was not yet applied. We rejected the use of RFID tags for usability reasons: (i) in typically crowded museum rooms tags are hardly locatable; (ii) since PDA-powered readers require touch-like proximity, a social protocol is needed in case multiple visitors intend for the same exhibit, something that is less preferred due to the significant cultural diversity among visitors; and (iii) they require technical equipment proximate to exhibits, which may not be allowed (e.g. the new Acropolis museum was designed with a glass-rooms, allowing extra equipment to be mounted only on the ceiling corners).

Extensible and Prioritized Location Sensing. To support location sensing, we have implemented an infrastructure supporting multiple concurrent sensing channels as dynamically installed and managed API instances, each given a specific static priority relating to practical precision characteristics (i.e. trust) of its respective location-sensing approach. Such an infrastructure essentially allowed our mobile navigation and administration systems to accommodate dynamically alternative location sensing APIs, in the form of runtime loaded components with an XML configuration file listing the API components to be loaded. Such scalability and extensibility was accomplished by making APIs pertain to the same generic interface, while loaded dynamically. In particular, the statically assigned trust-based priority is taken by invoking a specific method (*GetPriority* – see Fig. 3). This software architecture approach allows modularly incorporate alternative sensing technologies without modifying the original administration and navigation systems. The super-classes are provided under Fig. 3, defined in C++ for platform generality (the original specification is in C#).

```

typedef unsigned mapid_t;
typedef unsigned areaid_t;
typedef unsigned priority_t;
typedef pair<unsigned, unsigned> point2d_t;
typedef pair<mapid_t, point2d_t> location_t;
class LocationSensingAPI {
public:
    virtual string          GetId (void) const = 0;
    virtual priority_t      GetPriority (void) const = 0;
    virtual void            SetPollingInterval (unsigned t) = 0;
    virtual bool            IsConnected (void) const = 0;
    virtual void            StartUp (void) = 0;
    virtual void            CloseDown (void) = 0;
};
class PointBasedAPI : public LocationSensingAPI
{ public: virtual location_t GetLocation (void) const = 0; };
class AreaBasedAPI : public LocationSensingAPI
{ public: virtual list<areaid_t> GetAreas (void) const = 0; };

```

Fig. 3. The generic interfaces for dynamically loaded location sensing APIs

There are two derived super-classes for location sensing components, reflecting our classification of alternative technologies in two families: (i) those relying on user-position tracking and estimation (like Ekahau, GPS, radio tags / badges, and vision modules); and (ii) those relying on context marking (land-marking), like infrared beacons and passive / active radio tags. In our system, we referred to the technologies estimating directly user's position as *point-based* methods, while those relying on land-marks as *area-based* methods.

As we discuss latter, this separation is also very important for the mobile administration system where we needed to recognize conflicts when placing distinct land-marking tags: when at a particular position multiple beacons / tags are detected, though carrying different context identifiers, then we have a conflict (implying imperfect placement of the beacons / tags). Up to now, we already employed three specific technologies for location sensing, two of them being point-based and one being area-based, discussed below with increasing priority value:

- *WLAN* point-based location sensing, through the Ekahau Positioning Engine, both indoors and outdoors, with precision of 3-5 meters, requiring no line of sight.
- *GPS* point-based location sensing for outdoors environments. Average precision is 2-4 meters, requiring visual contact among the GPS card and the satellites.
- *TAGS* (currently mostly infrared beacons) area-based location sensing for indoors, and outdoors (when applicable) environments. For infrared beacons the precision at area level is virtually perfect, requiring the receptor to fall inside the active ray-diffusion conic space. Practical barriers concern the required line of sight.

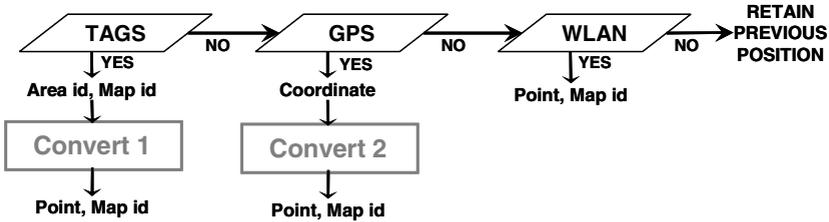


Fig. 4. The location sensing technologies deployed ordered by trust-based priority

The location-sensing pipeline is illustrated under Fig. 4, showing the way the APIs are actually deployed with decreasing priority. One should consider that this sort of processing is generalized for the case of N distinct prioritized location sensing API components. Following Fig. 4 the infrared receiver API is used to check if a signal is detected. The ray actually emits a single number in which we encode an area and a map identifier. From the area id its respective polygon and centre are extracted (indicted as *Convert 1*), the latter is returned together with the original map id as an approximation of the current users' position. It should be noted that the conversion of area identifiers to planar coordinates is not part of the area-based API implementation itself, but is a standard component of the overall location sensing API management logic, that we automatically apply to area-based APIs.

In case no signal is detected and the GPS is on, the current geographical coordinate is taken, used to identify the current enclosing map from the list of maps – maps are stored in the database with their geographical boundaries, and turn the geographical location to a raster coordinate inside this map (indicted as *Convert 2*). This conversion is included in the GPS point-based API implementation, returning the point and the map id as the current user's position. Finally, and alternatively, if the WLAN is on, the local EPE client API is used to inquire the current PDA position and the respective map id. As expected, if none of the previous technologies is active, the old user position is retained. The adoption of this particular implementation style to cope with varying location sensing technologies has been proved to be a key complexity-reduction factor in our project. For instance, although we generally dropped the deployment of Bluetooth tags (beacons) due to practical barriers, we also retained a respective API implementation as there are even at this point in time a couple of cases where we intend to use it as a back-up. In particular, in two specific archaeological sites, there are tombs whose entrances and main rooms are blind spots for both WLAN and GPS, while we do not yet have clearance to install infrared beacons around. In this context, one option we investigate is to install Bluetooth tags using the

power for the lights, close to the entrance and down to the floor, to at least be able recognize presence of visitors on site.

Supporting Multiple Navigation Scenarios. Our platform supports the authoring of arbitrary navigation scenarios, i.e. different types of tours that visitors may take, edited as sequences of selected information spots associated with a short descriptive title. During use, the navigator explicitly marks all information spots of the initially selected (by the visitor) navigation scenario on the map, displaying also their relative order within the navigation scenario. The authoring of navigation scenarios is shown under Fig. 5: the left column (Selected Items) is the ordered list of information points (titles) for the currently edited navigation scenario, while the right column (Remaining Items) is the list of all information points (i.e. entire information plane), offering typical editing facilities (add / move / reorder).

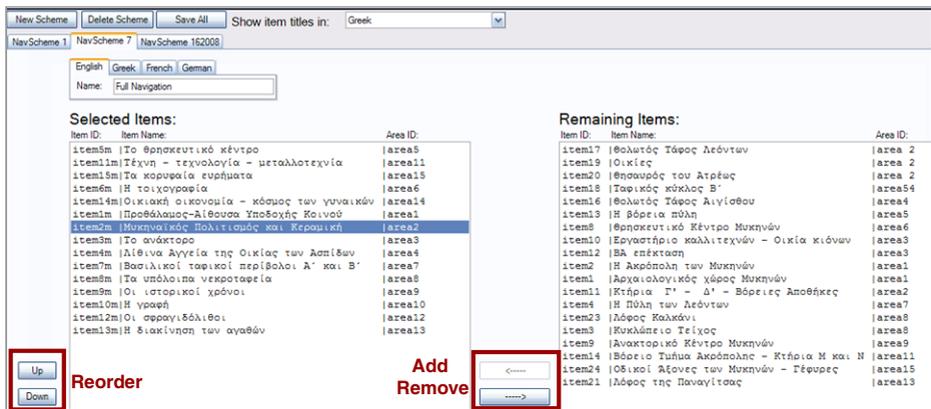


Fig. 5. Authoring of navigation scenarios

Mixed Initiative for Information Presentation. During the requirements analysis phase it came out that whether using a personal guide or navigating with a leaflet, visitors prefer to deviate from predefined paths to visit alternative proximate or distant information spots. Thus, the navigator had to guide users across the scenario, allowing them to freely deviate from it, effectively enabling them to keep track of their overall path. Additionally, experienced guides reported that during tour sessions, many visitors use leaflets or books to review information details for previously visited items, and in some cases even for items to be met later during the tour, since in many situations the exposed items are directly related to each other. This gave us an alternative insight on location-aware information delivery: besides automatic location-based information provision, users should be enabled to freely review all information points, whether visited or not. In our case, the latter is possible either via previous / next navigation buttons, or by pointing directly with the pen over items on the displayed map. A snapshot of the navigator's map view is provided in Fig. 5, showing how users may interactively control the display of extra details for information items, performing sequential transition among the following display options in a circular manner:

- Display no extra detail (default)
- Display the order of non-visited items that are part of the navigation scenario (see labels indicated with thick circles in Fig. 6)
- Display only the titles for all items

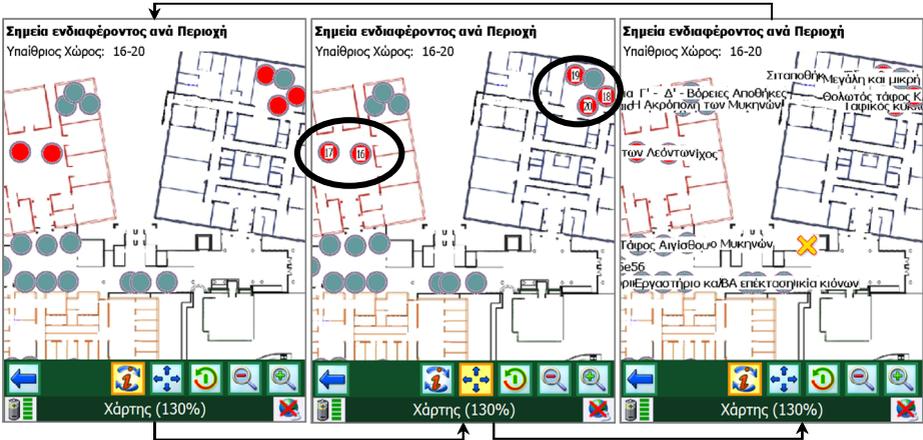


Fig. 6. Controlling level of overview detail for information items

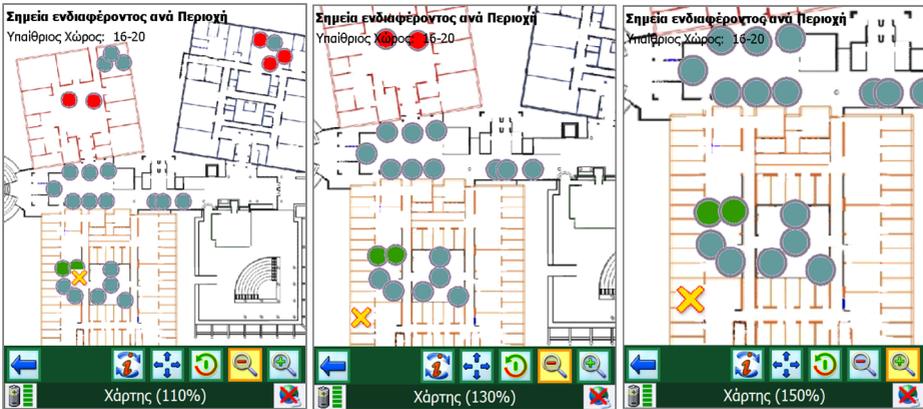


Fig. 7. Real-time smooth zooming. (110%, 130% and 150% are the zoom factors).

The user position is indicated with an X (see Fig. 6, middle of rightmost screen) while color encoding is used to distinguish among visited and non-visited items, with extra encoding for the latter to separate items belonging to the selected navigation scenario from those that do not. Additionally, real-time smooth zooming is supported as shown in Fig. 7; the latter was accomplished by implementing the map rendering functionality using Direct3D. In every map, the number and orders of unvisited items from the navigation scenario is briefed (top part of maps in Fig. 7). In many cases exhibits must to be placed very close to each other, with relative distances that cannot

be handled by the precision of the positioning systems. For instance, it is common to place ancient helmets, shields and arcs close to each other, each having a separate information unit. For such scenarios, our platform supports the grouping of multiple information items into a single chunk called an information area, physically being defined by a polygon encompassing the locations of all respective information items.

During navigation (see Fig. 8), once the visitor enters an information area, a menu of all information items is displayed (screen 1). The initial screen of an information item (screen 2) displays a representative image, together with two options, to review textual and multimedia content (screen 3). Video playback with a brief panoramic tour for an ancient castle is shown in screen 5. Besides automatic location-based information delivery, manual exploration of all information areas and points is allowed (next / previous arrows of screen 2). Information elements already reviewed by the visitor are drawn with a tick marker in their respective menu.

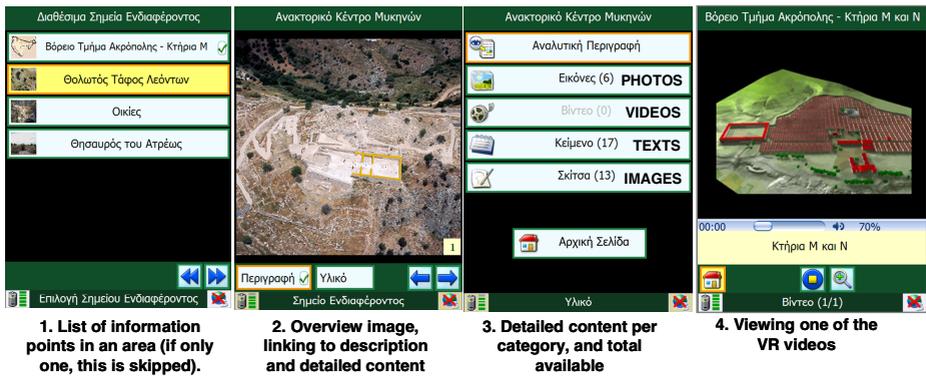


Fig. 8. Snapshots of the mobile navigation system during use

4 Mobile Administration System

The novel feature in this context is actually the *explicit provision of a mobile administration system for location data editing*, except from the traditional non-mobile content management and from-based data entry that is also offered. The support for mobile administration of position data is crucial for indoors and outdoors location-aware information systems relying on fine-grained positioning. While such a facility may seem redundant, assuming the common practice of GIS where content management is done off-site by positioning data over maps, as we explain below, the need for such a tool is mandated by various practical constraints. Firstly, location sensing may display varying accuracies at different physical areas. For instance, within indoor settings, WLAN accuracy is affected by many factors like material, room structure, crowd, and signal interference. During real-life installations, there have been rooms or corridors where we gained 2 meters accuracy, and there have been halls and open rooms where we never got something better than 5 meters. Secondly, the ability to use infrared beacons as wraparounds to overcome the accuracy problems point-based sensing is limited by a

number of practical factors, such as light reflections, room structure (ceiling height), signal diffusion, power supply points, and the practical barrier of line-of-sight.

Similarly, when it comes to outdoor environments, the dual use of GPS and WLAN positioning, to rectify GPS and WLAN blind spots, necessitated precise on-site programming and validation of position data. Overall, while performing content administration, authors should be enabled to practically evaluate the actual accuracy of location sensing accomplished during navigation sessions, so as to more accurately define: (a) the physical polygon of an information point that will trigger information delivery (when the user falls inside); and (b) whether multiple information points need to be collected into a single enclosing area (when we cannot get precision down to non-overlapping smaller areas). Additionally, testing of signal collisions (conflicts) among different infrared beacons is supported. The mobile administration system (see Fig. 9) offers the following features:

- Map display, showing the spatial distribution of information areas and points, the current administrator position as tracked by the active point-based location sensing technologies on site (WLAN and GPS in our case), while offering smooth zooming and scrolling facilities.
- Definition of areas and information points, editing of geometrical data for areas including point editing and “drag & drop” facilities – the latter, also offered by the spatial data editor of the content management, are elaborated in the next Section.
- Start trajectory recording and displaying of various user paths (way points) from both GPS and WLAN location-sensing technologies (see Fig. 9, left part).
- Testing of infrared tag signals and marking corresponding active areas, detecting conflicts, i.e. multiple signals at a certain position that encode different area identifiers (see Fig. 9, middle part).

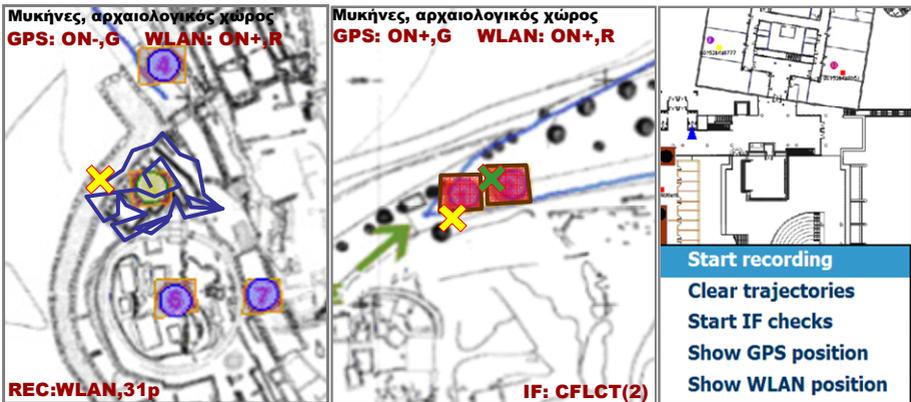


Fig. 9. Mobile administration in use (big arrow in middle part belongs to the map image)

The message displayed on top of the PDA screen explains whether GPS or WLAN are active at the current position (label *ON / OFF*), if position display was chosen (label *+ / -*), and the respective location marker colour (label *G* for green, *R* for red). The message prefixed with an *IF* concerns information for detected infrared signals,

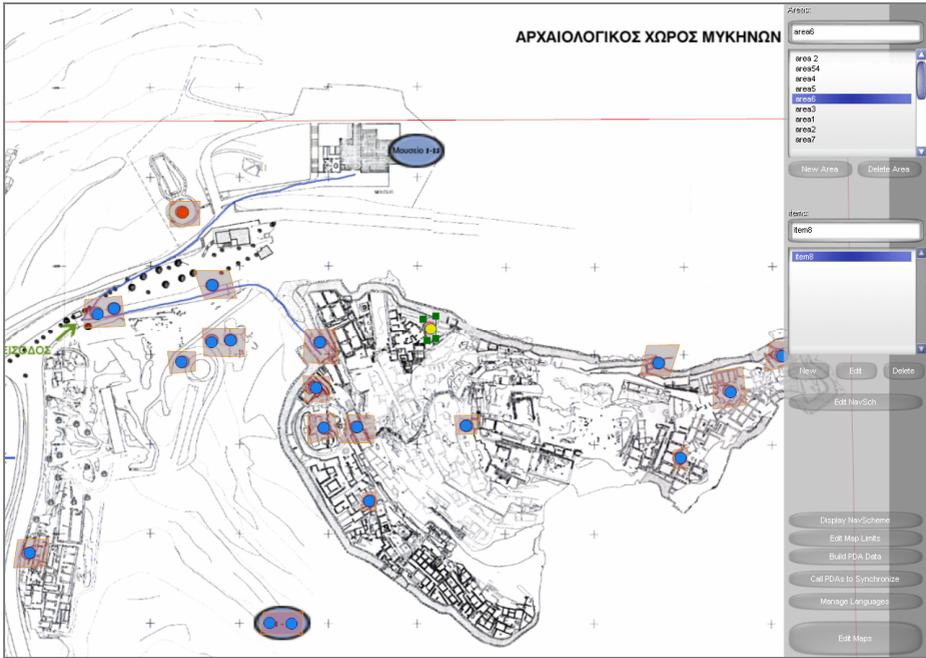


Fig. 10. Spatial content editing (Mycenae site); the horizontal line on top is part of the trespassing limits defined by the administrator for security reasons

CFLCT(2) meaning a conflict because of 2 different incoming signals. Additionally, for infrared signals, the respective areas are also highlighted (see Fig. 9, middle part - areas are drawn with semi-transparent red polygons). When recording user positions (*REC* indication), the sensing system chosen is shown (e.g. *WLAN*), together with the number of different successive positions returned, e.g. *31p*. Finally, the main control menu, activated by clicking on the map, is provided in Fig. 9, right part.

User Interface Automation. A very important remark is that the mobile administration system relies entirely on the dynamically loaded location-sensing API components. For instance, the specific identifiers displayed on the PDA, like “*GPS*”, “*WLAN*” and “*IF*”, as shown within Fig. 9, are not hard-coded in our implementation but are extracted during runtime from the APIs via *GetId()* so as to produce messages like “*GPS: ON-, G*” and “*WLAN: ON+, R*”. Moreover, the construction of our mobile administration User Interface is quite automated and parameterized. More specifically, given any point-based dynamically loaded API, following the definitions of the loading configuration file, we automatically add: (i) in the main menu a toggle option to tune position display like “*Show[Hide] <GetId(> position*”; (ii) a toggle option in the “*Start recording*” sub-menu as “*<GetId(> On [Off]*”, and (iii) an option “*<GetId(>*” in “*Clear trajectories*” sub-menu. In a similar way, for area-based loaded APIs we add a toggle option in the main menu to perform area identification (polling) and multiple signal (conflicts) checks as “*Start [Stop] <GetId(> checks*”.

5 Content Management System

They key feature of our content management system is that *content administration is facilitated through spatial data editing*, providing a direct-manipulation graphical editor to administer the semantic content and the geometrical data of information areas and information points directly over the displayed maps. Information entities are firstly positioned over the map using the mobile administration system or the spatial editor of the content management system. Such dual administration mode, to spatially position items both ‘on the desk-top’ and ‘on the move’ is not met in other systems. The two methods are used in cooperation as follows: (i) the content management system is used to define an initial approximate positioning and distribution of information areas and information items over the map; (ii) the mobile administration system is deployed later to fine-tune the positioning on-site, through the actual data provided by the location sensing systems.

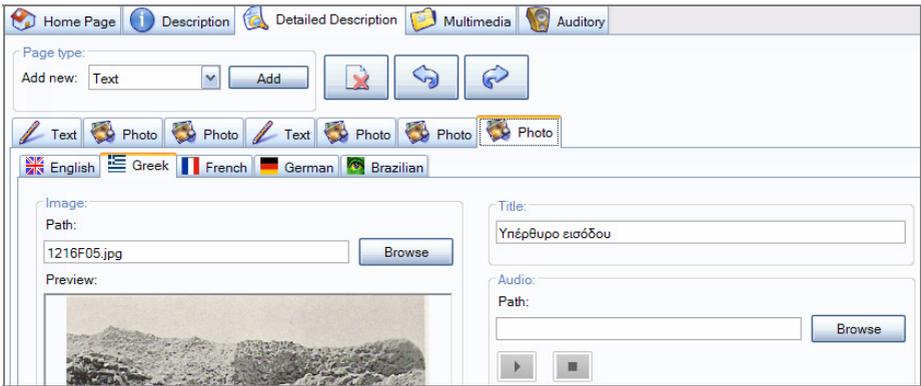


Fig. 11. Content editing opened from the spatial administration pane

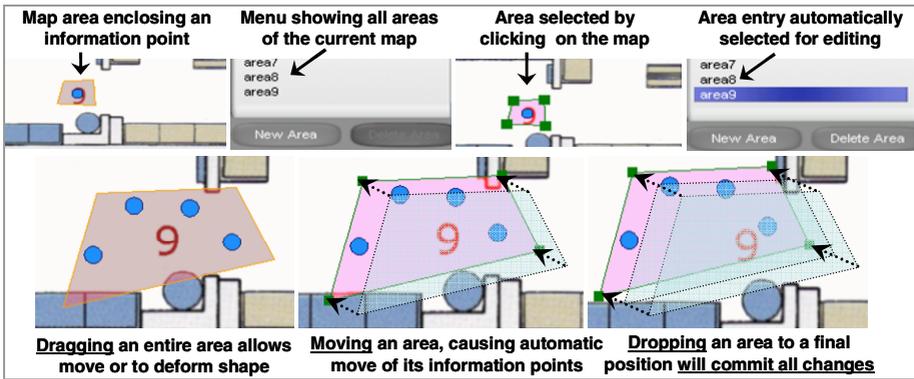


Fig. 12. Direct area manipulation with *drag & drop* facilities

The main window of the content management system is shown during deployment under Fig. 10; together with the map, all respective information areas and information points are also rendered. The control panel at the right side provides operations for opening the content editing forms, for all sorts of information types, as well as to edit maps and navigation scenarios. The content editing panel, for authoring of multimedia and multilingual information elements, is shown under Fig. 11.

The positioning adjustments, apart of typical displacements and polygon editing, may require radical changes, such as moving information points outside particular areas or joining multiple information points together into a new containing area. Under Fig. 12 we show some of the real-time features for editing the geometrical aspects of information areas, in particular selection, and drag & drop. More specifically, when an entire area is moved, all encompassed information points are automatically moved. In case that movement is not cancelled in the middle, the final positions of the information points are committed directly to the underlying content database. Such direct-manipulation facilities, offered by both the desk-top and the mobile versions of the content management systems, proved to be very valuable for practical administration of content spatial data. Item misplacements on maps and imperfect area definitions were far more frequent than originally anticipated, and it turned out that without automatic facilities, any approach relying on manual corrections directly on the database would be clearly impractical for a project of this scale.

6 Device and Renting Management System

Existing location-aware mobile information systems do not deal with the issue of explicit device renting and returning processes. Typically, users are pre-assigned a particular device (organization-wide use) or they are required to input individual information after deployment via on-line forms (sporadic trusted use). The latter are not applicable for protected, ad-hoc, commercial use, displaying high demands for quick servicing, as in museums, galleries and exposition settings. Our platform supports *very fast device management and renting processes* (charging of visitors is optional). It has been design to meet specific requirements coming from the Greek Ministry of Culture: ‘the system should support massive inquiries, e.g. incoming groups of 50 to 100 visitors; the delay for renting is considered unacceptable when it exceeds twice the delay of issuing a ticket receipt’.

The option of letting the system run on devices all the time was impractical: it wasted precious WLAN resources (bandwidth), it provided useless data to the visitor surveillance system, and it introduced fake entries (‘noise’) to the statistics component. So, practically, we had to find an error-prone way for unlocking, application activation, and language selection in less than 15 seconds. In this context (see Fig. 13) we adopted a variant of the process met in typical shopping transactions (right part): (i) WLAN restart is activated by pressing a large on-screen button on the PDA (not shown); (ii) the barcode reader is used to get the PDA’s barcode; and (iii) by consulting a database mapping barcode to IP addresses, connection with the PDA is internally established and a message is sent to it for unlocking and launching the information system with the selected target language. The latter required the explicit registration of PDAs (labeled as ‘Device management’ in Fig. 13), as triplets of barcode, IP and MAC addresses, something that is supported by the device management system.

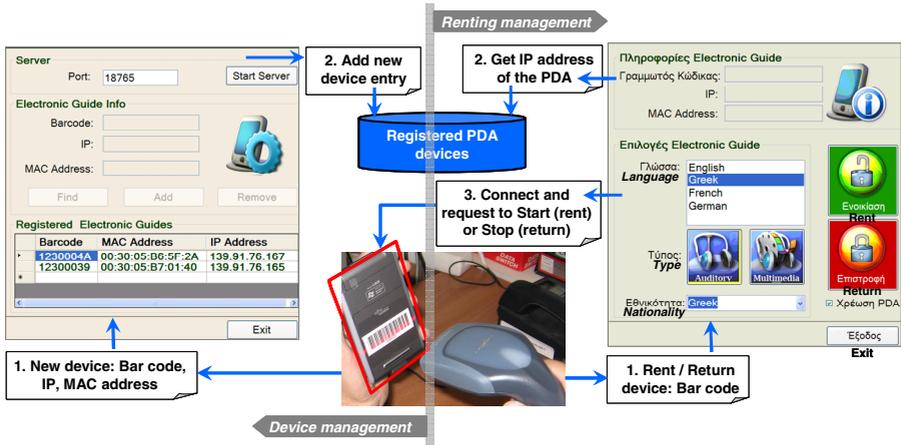


Fig. 13. Device and renting management systems and processes

Our contribution in this context, besides the use of barcode readers as a way for quickly renting the PDAs, is the enhancement of a well-known process by turning the actual purchased products, i.e. the PDAs, to active items capable of internally communicating with the database to self-initialize and self-configure for real use.

7 Discussion

The development of our platform was driven by specific objectives that justify some choices which look suboptimal at a first glance. In particular, we do not adopt the most powerful location sensing methods around (there are commercial solutions with claimed precision down to a few inches). This is due to our primary target: *being able to turn an existing infrastructure to a location-aware information context with minimal technical interventions, very low installation cost, and quick turnaround time*. We assumed as a rationale common denominator for the target environments presence of WiFi indoors and use of GPS outdoors, with availability of the respective chipsets on the PDAs. Also, we decided that digital tags like infrared or radio beacons may be sporadically employed as land-marking emitters. We selectively installed infrared beacons, using a model with up to six LEDs per device, placed at the top corners with a 45 degrees angle downwards (see Fig. 14).

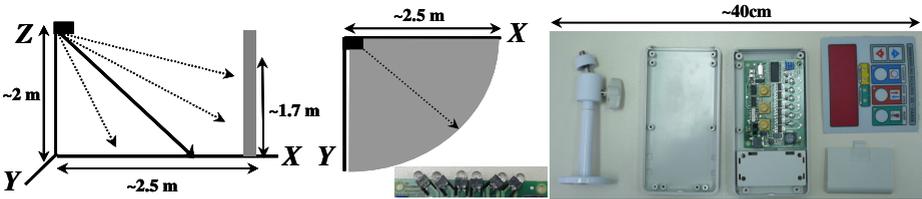


Fig. 14. Typical placement and boxing of infrared beacons in room corners

We focused in producing a tool to augment a variety of controlled environments into location-aware information contexts: exhibitions, organizations, museums, universities, hospitals, hotels, etc. This technical focus was proved to be the decisive factor enabling our platform to be applicable across the diverse museums and archeological sites of Greece. Practically, we have put a separation line with third-party location sensing technologies, treating everything from installation, configuration and calibration, to API licensing as something external to our platform. The only contact side with our system was the need for compliance to specific generic APIs, via which our platform became capable to deploy existing technologies without particular implementation dependencies. In any case, one cannot avoid dealing with the low-level details of such technologies to guarantee maximal gains, since the trust value that is incorporated in the respective location-sensing API should directly reflect real-life experience after performing numerous tests and experiments in the real field.

Problems Faced. Since we had five thousand PDAs across different museums with our software installed, we needed to address the issue of release updating after requested changes and bug fixing. For this purpose we added in the database all executables and DLLs, thus handled automatically by our synchronization system, and implemented a software updater to run locally at the PDA. This way, version upgrading became fully automatic and centralized. In outdoor settings, we observed considerable delays (sometimes tens of seconds) when a group of more than 20 visitors exits a WLAN blind spot; the delay was attributed to connection reestablishment at the system level with the WLAN. We also had some unexpected GPS precision issues at some areas, with accuracy frequently falling down to 8-10 meters. Fortunately, information provision on archeological sites was decided with information spots very distant to each other; however, besides information delivery, a few visitors observed this issue when reviewing their position on the map. Finally, in outdoor settings, having more than hundred PDAs connected to the same access point causes severe delays on WLAN positioning updates; these temporary stalls were bypassed since the GPS has highest trust value. Nevertheless, there are still scenarios where GPS low accuracy intervals appear either at network blind spots or at sites where the sporadic installation of access points turns WLAN positioning quite inaccurate to alternatively trust it.

8 Summary and Conclusions

We have presented key components of an integrated platform for the authoring and uptake of location-aware mobile information systems, putting emphasis on specific novel features reflecting requirements that emerged in the course of large-scale real-life deployment. The development of the overall platform required three years (2004-2007), while authoring and installation for the 15 major museums and archeological sites of Greece started at February 2007, and is completing during May 2008, with pilot use scheduled for summer 2008. Due to the considerable size of the project, both in terms of equipment (five thousands PDAs), as well as in terms of divergence of physical exhibition areas (e.g., hills, tombs, protective roofs, covered or open amphitheatres, halls, corridors, glass rooms, forests, open stadiums, etc.), we had to work in close cooperation with Ekahau for calibration sessions as well as for engine improvements, and with

the manufacturers of infrared beacons for optimized customized solutions (both in terms of h/w type, as well as in terms of SDK support). All computing equipment, peripherals, networking and constructions are delivered by SIEMENS.

Amongst the most technically challenging issues was the accomplishment of satisfactory location-sensing across all different settings. In this context, the conclusions from our technical decisions discussed in the paper are: (a) the unified location-sensing infrastructure for indoors and outdoors simplifies maintenance and extensibility; (b) trust-based statically-prioritized multi-channel location sensing solves many precision problems increasing significantly the robustness the position tracking infrastructure, even though it puts an extra authoring overhead; and (c) when tags are introduced to augment the location-sensing approach, they should be programmed to identify locations or areas, i.e. return geometrical information, rather than transmitting internal database keys or indices of particular information items. Regarding content authoring, the support for spatial data editing and mobile on-site editing and adjustment of position data, enabling to query and test location values, proved to be very helpful. The latter is not only due to easier administration of location-dependent content by being on site, but most importantly because of the ability to rectify the frequent precision problems. Our location adjustment tool works on top of, and in cooperation with, the Ekahau calibration system (site survey tool). The latter prepares the ground for EPE location sensing, and the former allows adapting initial positioning of exhibits to the real output of the location system after its calibration phase.

Acknowledgements

Part of this work is funded by the Greek Ministry of Culture, via a subcontract from SIEMENS to FORTH, for installation and deployment of the platform at fifteen museums and archeological sites across Greece, including Acropolis, Olympia Knossos, Delphi and Mycenae. The total project budget is 11 Million Euros, with installations and pilot deployment scheduled to complete before summer 2008.

References

- Bahl, P., Padmanabhan, V.N.: RADAR: An In-Building RF-Based User Location and Tracking System. In: Proceedings of IEEE INFOCOM 2002, vol. 2, pp. 775–784 (2002)
- Blache, B., Chraïet, N., Daroux, O., Evennou, F., Flury, T., Privat, G., Viboud, J.P.: Position-Based Interaction for Indoor Ambient Intelligence Environments. In: Aarts, E., Collier, R.W., van Loenen, E., de Ruyter, B. (eds.) EUSAI 2003. LNCS, vol. 2875, pp. 192–207. Springer, Heidelberg (2003)
- Harle, R.K., Ward, A., Hopper, A.: Single Reflection Spatial Voting: A Novel Method for Discovering Reflective Surfaces Using Indoor Positioning Systems. In: Proceedings of ACM MobiSys 2003, International Conference on Mobile Systems, Applications and Services, pp. 1–14 (2003)
- Hightower, H., Want, R., Borriello, G.: SpotON: An Indoor 3D Location Sensing Technology Based on RF Signal Strength, UW CSE 00-02-02, University of Washington, Department of Computer Science and Engineering, Seattle, WA (February 2000)

- Hightower, J., Borriello, G.: Particle Filters for Location Estimation in Ubiquitous Computing: A Case Study. In: Davies, N., Mynatt, E.D., Sio, I. (eds.) *UbiComp 2004*. LNCS, vol. 3205, pp. 88–106. Springer, Heidelberg (2004)
- Krumm, J., Harris, S., Meyers, B., Brumitt, B., Hale, M., Shafer, S.: Multi-Camera Multi-Person Tracking for EasyLiving. In: *Proceedings of the 3rd IEEE International Work-shop on Visual Surveillance (VS 2000)*, pp. 3–3 (2000)
- LaMarca, A., Chawathe, Y., Consolvo, S., Hightower, J., Smith, I., Scott, J., Sohn, T., Howard, J., Hughes, J., Potter, F., Tabert, J., Powledge, P., Borriello, G., Schilit, B.: Place Lab: Device Positioning Using Radio Beacons in the Wild. In: Gellersen, H.-W., Want, R., Schmidt, A. (eds.) *PERVASIVE 2005*. LNCS, vol. 3468, pp. 116–133. Springer, Heidelberg (2005)
- Nord, J., Synnes, K., Parnes, P.: An Architecture for Location Aware Applications. In: *Proceedings of the 35th Hawaii International Conference on System Sciences - HICSS 35 2002*, IEEE, Los Alamitos (2002) in CD-ROM
- Pfeifer, T.: Redundant Positioning Architecture. In: *Computer Communications*, vol. 28(13), pp. 1575–1585. Elsevier Science Publishers B.V, North-Holland (2005)
- Priyantha, Nissanka, B., Chakraborty, A., Balakrishnan, H.: The Cricket location-support system. In: *Proceedings of the ACM MOBICOM 2000*, 6th International Conference on Mobile Computing and Networking, pp. 32–43 (2000)
- Savarese, C., Rabaey, J., Langendoen, K.: Robust Positioning Algorithms for Distributed Ad-Hoc Wireless Sensor Networks. In: *Proceedings of the General Track: 2002 USENIX Annual Technical Conference*, pp. 317–327 (2002)
- Savidis, A., Stephanidis, C.: Distributed Interface Bits: Dynamic Dialogue Composition from Ambient Computing Resources. *ACM-Springer Journal on Personal and Ubiquitous Computing* 9(3), 142–168 (2005)
- Schwieren, J., Vossen, G.: Implementing Physical Hyperlinks for Mobile Applications using RFID Tags. In: *11th International IEEE Database Engineering and Applications Symposium (IDEAS 2007)*, Banff, Canada, September 6-8, pp. 154–162 (2007)
- Sumi, Y., Etani, T., Fels, S., Simonet, N., Kobayashi, K., Mase, K.: C-MAP: Building a Context-Aware Mobile Assistant for Exhibition Tours. In: Ishida, T. (ed.) *Community Computing and Support Systems*. LNCS, vol. 1519, pp. 137–154. Springer, Heidelberg (1998)