

Indoor Positioning and Navigation with Camera Phones

This low-cost indoor navigation system uses off-the-shelf camera phones to determine user location in real time by detecting unobtrusive fiduciary markers.

GPS-based navigation systems have become very popular, mainly because they let people rapidly explore unknown areas. However, GPS works only outdoors, because the required satellite links are blocked or unreliable inside buildings. For indoor location sensing, many types of sensors exist, including ultrasonic, infrared, magnetic, and radio sensors. But they all require a permanent electronic infrastructure to facilitate measurements, and localizable objects relying on this infrastructure need special sensors or actuators. Practical problems such as power consumption, wiring, and overall infrastructure cost have inhibited such technology's deployment in entire buildings.

In previous work,¹ we demonstrated how to detect and decode square fiduciary markers in real time using off-the-shelf camera phones. Such markers contain a 2D barcode that provides a unique ID from the camera image, which the camera phone can use to estimate in real time its position and orientation relative to the marker. The indoor navigation system we describe in this article takes advantage of associating locations with markers to provide an inexpensive, building-wide

orientation guide that relies solely on camera phones. Whereas previous work on barcode-based location tracking, such as QR Codes (www.qrcode.com), relies on non-real-time "snapshot" processing, our approach allows for continuously scanning an environment in real time (15 Hz or more) in search of navigation hints. Thus, navigation scales from sparse, strategically placed fiduciary markers to continuous navigation in 3D.

Real-time marker-based tracking of position is available in several desktop-based applications, but it has only recently become available in phones. Consequently, systems for continuous navigation using marker-based tracking haven't been studied extensively. Non-real-time marker-based recognition in existing location-based services usually takes several seconds, and the service providers typically place markers to highlight a particular location rather than to act as pure navigation landmarks (see the "Indoor Navigation and Localization" sidebar).

We've examined marker-based tracking's suitability for continuous navigation in mobile phones. We conducted a controlled user study to compare our system with a map without localization and with a GPS-like real-time localization. In this article, we provide an evaluation of subjective experiences about ease of use and location awareness. We also discuss experiences from deploying our software at four large-scale events, testing its usability under real-world con-

Alessandro Mulloni,
Daniel Wagner,
and Dieter Schmalstieg
Graz University of Technology

Istvan Barakonyi
*Imagination Computer
Services GesmbH*

Indoor Navigation and Localization

There is a large body of work on indoor navigation in robotics. Guilherme DeSouza and Avinash Kak provide a good overview.¹ These systems commonly harness a robot's controlled position and movements and try to detect its pose using natural-feature tracking based on cameras or range sensors. However, such algorithms exceed what is currently possible on mobile phones.

Among the first dedicated wearable location systems was Active Badge,² which consisted of infrared badges sending location information signals to a server. Its successor, the Bat system,³ used ultrasonic location estimation to provide more accurate position data. Another system for location tracking, PlaceLab,⁴ used signal strength of various wireless connections such as GSM (Global System for Mobile Communications), Bluetooth, and Wi-Fi. Accuracy strongly depended on the number of senders in the environment and ranged from 3 to 6 meters for indoor use. For a good overview of positioning technologies, see the survey by Jeffrey Hightower and Gaetano Borriello.⁵

Cyberguide was an early project targeting human indoor navigation and guidance.⁶ It used remote controls as low-cost infrared beacons, but the cost of the remote controls prevented deployment in larger areas or large numbers of users. The eGuide project and the Resource-Adaptive Mobile Navigation System use similar techniques.^{7,8} Davide Merico and Roberto Bisiani use inertial sensors to track user movements.⁹ Periodically, users must calibrate their position by choosing distance measurements in panoramic views of the environment on the device's screen. Naturally, creating these views is work intensive for large areas. Harlan Hile and Gaetano Borriello report an indoor navigation system based on the scale-invariant feature transform (SIFT) method.¹⁰ This system relies on a server to outsource the actual pose estimation work, providing limited scalability and long latency—a processing time of roughly 10 seconds per image makes this system unsuitable for large-scale deployment.

Recently, Tsutomu Miyashita and his colleagues presented a PC-based museum guide using natural-feature tracking.¹¹ As

with our system, localization works on only certain hot spots. But the markerless tracking approach makes it difficult for visitors to recognize those hot spots.

REFERENCES

1. G.N. DeSouza and A.C. Kak, "Vision for Mobile Robot Navigation: A Survey," *IEEE Trans. Pattern Analysis and Machine Intelligence*, Feb. 2002, pp. 237–267.
2. R. Want et al., "The Active Badge Location System," *ACM Trans. Information Systems*, Jan. 1992, pp. 91–102.
3. M. Addlesee et al., "Implementing a Sentient Computing System," *Computer*, Aug. 2001, pp. 50–56.
4. V. Otsason et al., "Accurate GSM Indoor Localization," *Proc. Int'l Symp. Ubiquitous Computing (UbiComp 05)*, LNCS 3660, Springer, 2005, pp. 141–158.
5. J. Hightower and G. Borriello, "Location Systems for Ubiquitous Computing," *Computer*, Aug. 2001, pp. 57–66.
6. G.D. Abowd et al., "Cyberguide: A Mobile Context-Aware Tour Guide," *Wireless Networks*, Oct. 1997, pp. 421–433.
7. J.L. Encarnação and T. Kirste, "Beyond the Desktop: Natural Interaction and Intelligent Assistance for the Everyday Life," *Proc. Computer Graphik Topics*, 2000, pp. 16–19 (in German).
8. J. Baus, A. Kruger, and W. Wahlster, "A Resource-Adaptive Mobile Navigation System," *Proc. 7th Int'l Conf. Intelligent User Interfaces (IUI 02)*, ACM Press, 2002, pp. 15–22.
9. D. Merico and R. Bisiani, "Indoor Navigation with Minimal Infrastructure," *Proc. 4th Workshop Positioning, Navigation and Communication (WPNC 07)*, IEEE Press, 2007, pp. 141–144.
10. H. Hile and G. Borriello, "Information Overlay for Camera Phones in Indoor Environments," *Proc. 3rd Int'l Symp. Location- and Context-Awareness (LoCA 07)*, LNCS 4718, Springer, 2007, pp. 68–84.
11. T. Miyashita et al., "An Augmented Reality Museum Guide," *Proc. 7th IEEE/ACM Int'l Symp. Mixed and Augmented Reality (ISMAR 08)*, IEEE Press, 2008, pp. 103–106.

conditions. Both evaluations suggest that marker-based navigation with camera phones works well and fits real-world requirements.

Human-Friendly Fiduciary Markers

Localization using fiduciary markers is a well-established mechanism in mobile applications. Unlike natural-feature tracking, detecting and decoding artifi-

cial markers is highly robust and works well under varying lighting conditions and with minimal computational resources. Moreover, fiduciary markers' unusual visual appearance makes them more noticeable, helping users identify information hot spots in visually cluttered environments as well as large environments where interest points are sparse. For example, the Yellow Arrow project (<http://yellowarrow.net>)

encourages users to place highly visible stickers with barcodes worldwide, each linking to online content.

Our marker-tracking software library can estimate, with 6 degrees of freedom (DOF), a camera phone's position and orientation with respect to markers. Because maps are only 2D representations of the world, we can limit the localization to use only 3 DOF, sufficient to estimate the phone's 2D



Figure 1. Development of fiduciary markers and conference posters at (a) the 2007 Microsoft Mobile and Embedded Developers Conference (MEDC 2007) and (b) Microsoft's TechReady7; and (c) use of our localization system at TechReady7.

position and orientation. This makes our system more practical to deploy because marker positions and orientations need only be specified in 2D.

We deployed traditional square markers at the 2007 Microsoft Mobile and Embedded Developers Conference (MEDC 2007), as Figure 1a shows. Over time, we introduced a new type of less visually obtrusive markers called *frame markers* (see Figure 1b). Frame markers encode the same amount of data (36 bits) as traditional square markers, but they arrange the data along the border: 9 bits are encoded at each marker side in small black and white squares, carrying 1 bit each. The frame marker's specifications arrange the code in clockwise order such that only one of the four possible rotations yields a valid code. The remaining part of the square isn't used to decode the marker and can contain arbitrary content, giving graphics designers freedom for customization.

To exploit markers, users must point the phone's camera at them (see Figure 1c). As soon as a marker appears in the camera's view, the system detects and decodes it in real time from the live video stream. For detection to succeed, each square describing a bit should be at least 2 pixels in the camera image. Our system can cope with tilt angles of up to roughly 70 degrees, and rotations around the camera's optical axis don't affect the system. These constraints are comparable to other marker-based tracking techniques.

Conference Guide Application

Large events such as conferences often challenge participants to find their way through vast multistory convention cen-

ters or hotels. Using our marker-tracking technology, we created a location-based conference guide called *Signpost*. We designed this application to work with sparse tracking, to limit the number of deployed markers to a manageable size. For example, we installed 37 markers at a conference site in the Venetian Hotel Las Vegas in April 2007, in an area of roughly 100 m × 200 m.

Although 6-DOF tracking can deliver centimeter-level accuracy when markers are tracked, presenting only the 2D location on a map reduces accuracy requirements significantly. This is important because conference organizers must consider the logistics of deploying and inspecting marker placement. The most efficient way, developed after consulting conference organizers, was to stick markers onto poster stands, which could be quickly deployed on site at preplanned locations (see Figure 1c).

Signpost combines a conference calendar and a navigation system. Users can query the conference calendar by day or conference session, or by using full-text indexing. Live RSS updates received over the air ensure that the schedule reflects the latest changes. All calendar entries are linked to locations, so users can plan their fastest route from the current location (the last seen marker) to the desired lecture hall. Signpost displays the results on a map that users can freely pan, rotate, and zoom. Alternatively, in live-tracking mode, the system automatically aligns the map as soon as a marker is detected.

We implemented Signpost atop the Studierstube ES (Embedded Systems) framework,¹ as Figure 2 shows. The system runs on Windows Mobile phones (Figure 3), independent of screen resolution and form factor. Signpost can impact the device's battery life, mainly because of its use of the camera and the network connection. However, when no one is using the Signpost application, the system automatically disables the camera. Hence, battery drain is limited to when users actively interact with the application. Furthermore, Signpost

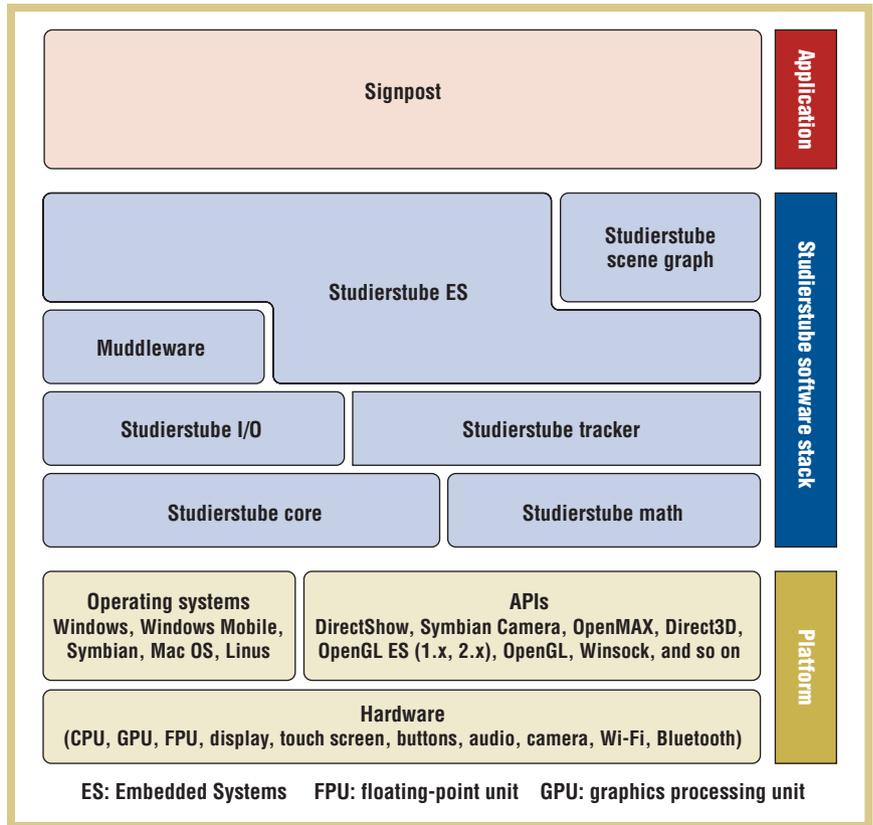
Figure 2. Software architecture of the mobile phone client. Our software framework, Studierstube ES, provides the application layer with an abstraction of all platform-dependent functionalities, allowing for an easier cross-platform development of applications.

uses the network connection only when users explicitly ask to update the conference schedule.

For large events in venues with multiple levels or buildings, a single map is no longer sufficient. Signpost, therefore, supports multiple maps linked to a 3D overview or, alternatively, an interactive 3D representation showing the current and target locations' global geographic relationship.

Deploying the system to a new location consists of three steps:

- *Create a map and database of marker locations.* On the basis of sketches or CAD data, create one or more 2D images using maps of the target location. The system uses bitmaps rather than vector graphics images



for the maps, allowing reuse of existing maps and artworks. On the basis of gatherable information from the

2D maps, select preliminary marker locations and orientations and enter them into a configuration file. This

Figure 3. Examples of phones running Signpost and screenshots of the application: (a) HTC Cingular 8525, (b) Orange SPV E600, (c) HTC Touch Diamond, (d) HTC S710, and (e) Motorola Q. Signpost supports various form factors and resolutions of the screen, and different input capabilities (both keypad- and touch-based).

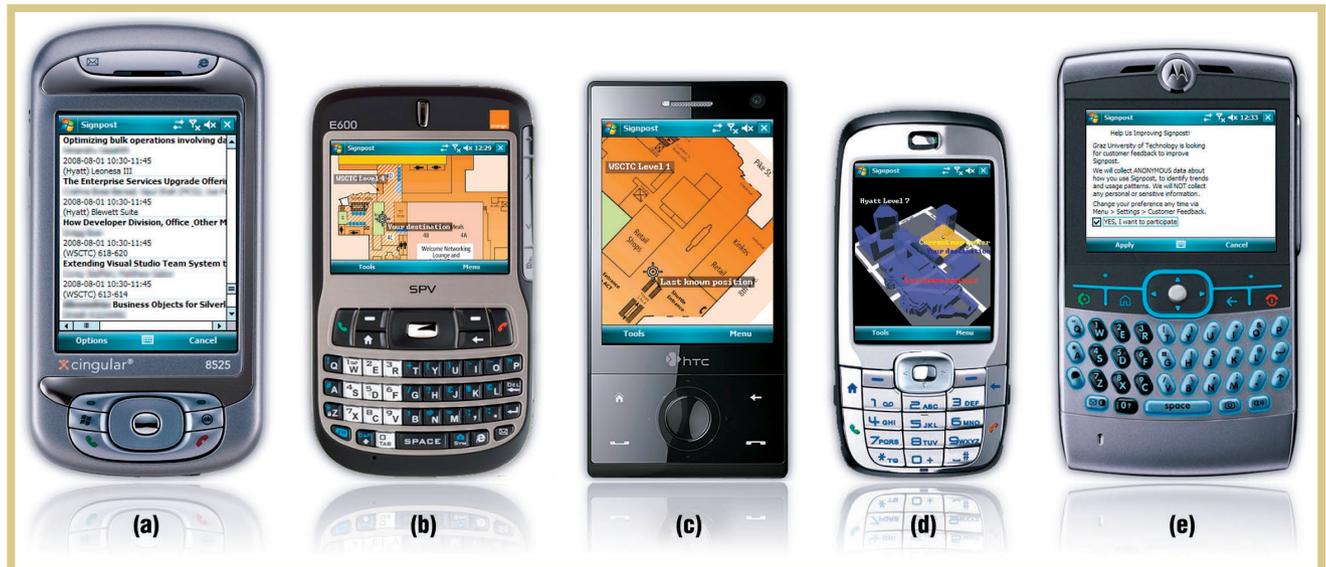




Figure 4. Positioning marker-based localization systems on an ideal localization continuum. This continuum spans from systems without localization to those with continuous localization.

step can take place offline in a planning office without access to the target location.

- *Deploy markers on site.* During event preparations, deploy markers on site. Depending on the required accuracy, a coarse deployment (± 50 cm) is often sufficient. For simplicity, markers' horizontal bearing is restricted to 45-degree steps, which is usually sufficient when mounting them on walls or poster boards. For those markers that can't be deployed as planned in the previous step, update the database accordingly.
- *Create a new software release.* Finally, create a new software release by bundling the software with the updated map and marker database. Changes to this data (such as repositioning markers or changes to the schedule) that become necessary later can be deployed over the air.

We envisioned a typical usage pattern to be as follows: First, a user browses the schedule, choosing a desired talk. Second, by selecting this talk, the user can see its location on the map. Third, when Signpost detects a marker, the application shows the user's current position, helping the user to decide how to reach the talk. Fourth, if desired, the

user can get a better understanding of the conference complex by looking at the 3D view.

Comparison of Localization Techniques

Even though many projects have used marker recognition, we aren't aware of any studies that compare marker-based localization with nonlocalized digital maps. So, we conducted a study to assess whether the effort of outfitting the environment with fiducial markers pays off in terms of improvements in user navigation. We compared marker-based navigation in Signpost with two conditions that represent the extremes of the localization continuum, shown in Figure 4: a digital map with no localization and a system with continuous real-time localization (similar to GPS-based navigation systems).

We hypothesized that the *continuous localization* system is significantly easier to use than the other systems, while providing the highest degree of location awareness. We also hypothesized that *discrete localization* is significantly more helpful in terms of location awareness than the *no-localization* condition, while probably requiring more learning effort. The study's overall goal was to show that marker-based localization is

a good solution for navigation when no GPS positioning is available.

We recruited 20 users with diverse cultural backgrounds and varying expertise in technology. The users were between 20 and 34 years old (average age of 25), with half male and half female. For each of the three conditions, we asked the participants to use the maps and the localization system as their only aids to reach a specific destination. The location of the study was the Department of Computer Science at Graz University of Technology (in Graz, Austria)—a complex comprising four buildings connected by several bridges. This complex contains many repeated features, with a general lack of clear landmarks. We considered it to be a significant example of a hard case for navigation in a new environment. To avoid biased results, we ensured that no user had previous experience with the buildings. We selected three different destinations, balancing their difficulty in terms of distance from the starting point, bridges to cross, and number of floors. The study used a *within-subjects design* (all test users tried all three conditions) with randomization (based on Latin squares) of the order of conditions and target rooms to avoid bias. We gave all users some time to famil-

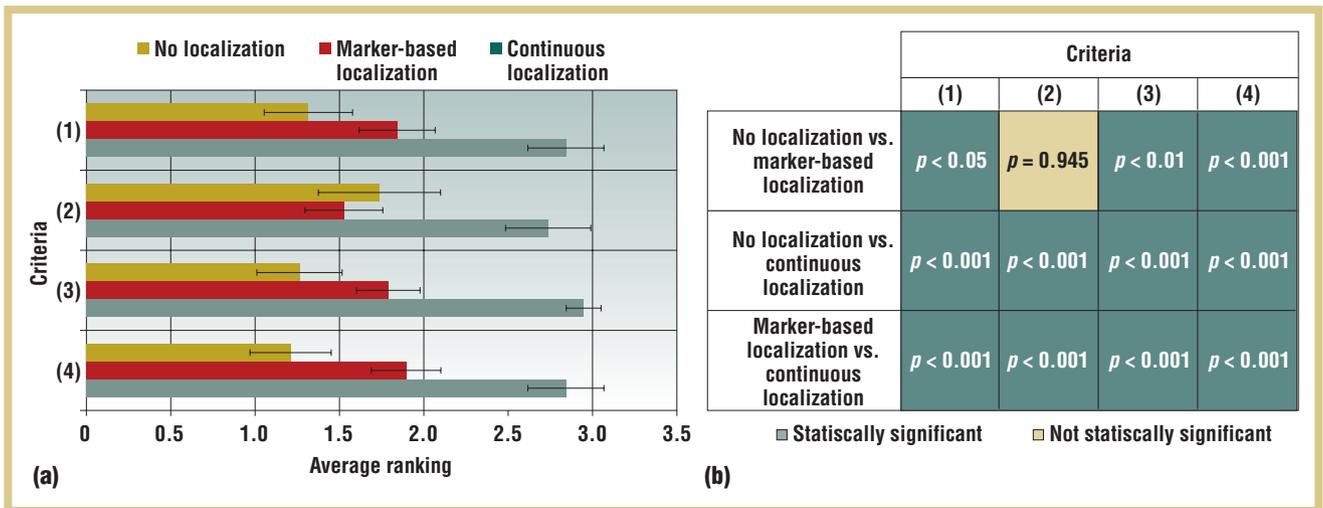


Figure 5. Subjective rankings of the three localization systems, based on four different criteria: (a) average ranking of the three localization systems (higher scores indicate better ratings); (b) statistical significance of pair-wise differences (as a result of Bonferroni post hoc analysis). The criteria were as follows: (1) the system is easy to use, (2) the system is easy to learn, (3) the system requires little attention, and (4) the system makes me confident I know where I am. Horizontal lines indicate 95 percent confidence intervals.

iarize themselves with the touch-screen-based application before starting the evaluation.

For the no-localization condition, we implemented a digital map viewer, which lets users pan the maps with a finger on the touch screen. In all conditions, each map showed the start and destination points with crosshairs. When such locations were outside the map’s view, we presented off-screen directions using labeled arrows. Users could access each floor’s maps using keypad shortcuts on the phone.

For discrete localization, we integrated our marker-based solution into the navigation map and presented the live camera video view in a screen corner. The map presented each marker’s position as a red dot. As soon as the system detected a marker, it automatically updated the user’s detected position and orientation, presenting it as a labeled icon on the map. However, it didn’t reposition and reorient the map itself automatically. Our pilot study revealed that users prefer to rotate and center the map manually when using a system that doesn’t provide continuous localization.

For the continuous-localization condition, we didn’t have an indoor equivalent to GPS available. Therefore, we used a “Wizard of Oz” approach: we introduced a hidden operator to manually control the map’s position and rotation using a second phone connected via Bluetooth. A crosshair on the map showed the user’s current position. This setup proved to be believable, and the users sufficiently concentrated on their devices such that no one noticed the trick. Continuous localization was a control condition useful only for comparison within our experiment, and the Wizard of Oz approach let us quickly build a running system. Even if this solution can’t scale to real-world situations, we considered its scope to be limited to the described user study.

After users completed all three tasks, we asked them to rank the three conditions from worst to best, according to four different criteria: ease of use, ease of learning, required attention, and confidence about the current location. For each rating, we assigned a score of 1 to the worst condition and 3 to the best condition. Figure 5 shows the average rankings and their 95 per-

cent confidence intervals. Friedman’s test shows that the effect is significant for every criterion ($p < 0.001$), with a high probability that differences in scores aren’t due to chance. Figure 5 also shows the results of the Bonferroni test for post hoc analysis.

As expected, the continuous-localization condition outperformed the other two conditions for all criteria. Surprisingly, though, users found discrete localization significantly easier to use than no localization, while requiring less attention. During the user study, we noticed that the localization information provided by markers helped users mentally register the view on the digital map with the real environment. Although in the no-localization condition users looked for matching landmarks in the environment with landmarks on the map, with discrete localization, the burden was reduced to registering the icon on the map with the user’s real position and orientation in the world. This might explain the ease of use reported for discrete localization. The results show no significant difference in the ease of learning between the two conditions,

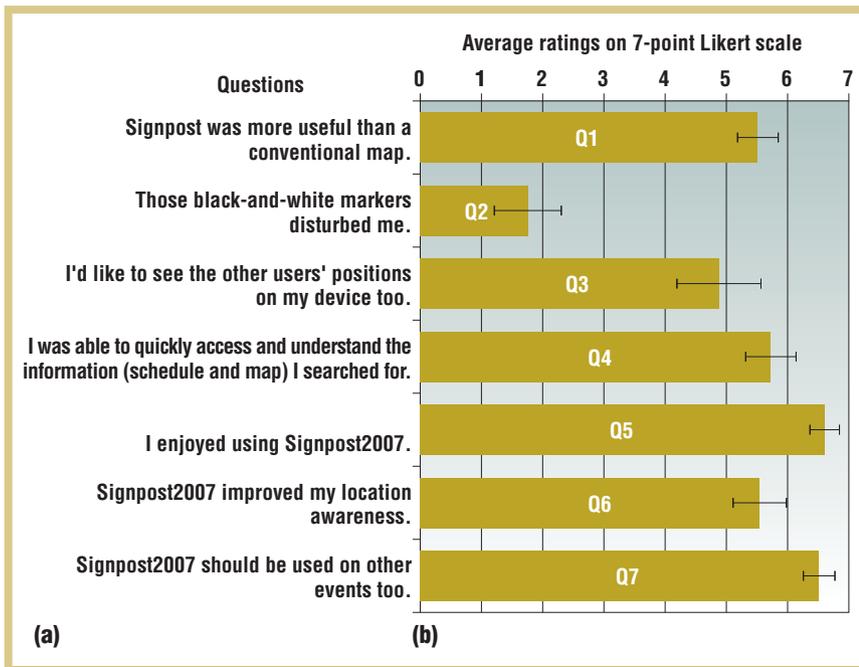


Figure 6. Results of 34 anonymous questionnaires from MEDC 2007: (a) questions given, and (b) average ratings for each question on a 7-point Likert scale. The vertical lines indicate the 95 percent confidence intervals.

although discrete localization scored slightly worse than no localization. In our application, we address marker-related learnability problems by introducing video tutorials or by providing on-screen hints (for instance, viewfinder frames, which are common in photo cameras).

Finally, user confidence was significantly higher ($p < 0.001$) for discrete (marker-based) localization compared to no localization. Thus, compared to a static map, users appear to feel a significant increase of location awareness when provided with a means of verifying their position on the map, even if coarse and discrete. Our observations during the user study support this conclusion. Although users generally didn't use markers intensively when they were going in the right direction, markers seemed fundamental for users who were lost in a wrong branch of the building so that they could remap their mental model with the real building and restructure their path accordingly.

Experiences from Real-World Deployment

We deployed Signpost at four international conferences: MEDC 2007

(April 2007), Microsoft Tech Ed 2007 (June 2007), TechReady6 (February 2008) and TechReady7 (July 2008). The number of distinct users that installed Signpost on their devices rose from 150 at MEDC 2007 to more than 1,000 at TechReady6. To our knowledge, this is the most widely used phone-based indoor navigation system ever deployed in a real environment.

At all four conferences, we introduced Signpost as the official conference guide endorsed by the conference organizer. All users were conference attendees who hadn't seen the application beforehand and who didn't know us. By deploying our application in previously unknown environments and with a large user base possessing untested hardware, we were able to collect data from many users in a natural environment via anonymous usage logs, questionnaires, on-field observations, and interviews. These real-world experiments complement the controlled study presented earlier.

Unfortunately, approaching all users directly wasn't possible. Yet, we managed to interview some of them

personally and to collect questionnaires and usage logs from many of them. Our overall aim was to determine how useful attendees found the application. More specifically, we wanted to learn what worked and what didn't, and which features were appreciated and which were missing. We also wanted to gain insight into other research areas, such as pedestrian navigation.

During MEDC 2007, we collected 34 anonymous questionnaires. Attendees marked their answers on a Likert scale from 1 (strongly disagree) to 7 (strongly agree). Figure 6 presents the results. All answers were consistent, with only a minimal standard deviation.

At TechEd 2007, we deployed the application at the request of the conference organizers, but we didn't manage to get in contact with any of the conference attendees.

At TechReady6, we systematically observed selected users to monitor their behavior patterns, and we performed recorded, semistructured interviews afterward. The interviews focused on several core topics, such as how Signpost changed the user's conference experience and organization, how well the navigation worked, and how much the small screen limited the application's usefulness.

Finally, during TechReady7, we collected usage logs from 74 anonymous users, covering a time frame of four days, to better understand how users employed Signpost. We identified the following core functions: display of 2D maps, visualization of the conference buildings' 3D models, live positioning using markers, browsing of the conference schedule, and full-

Figure 7. Usage statistics for five core functionalities (from TechReady7): (a) percentage of times each functionality was invoked, and (b) average usage count of each functionality per user (relative to the days since the user first ran the application).

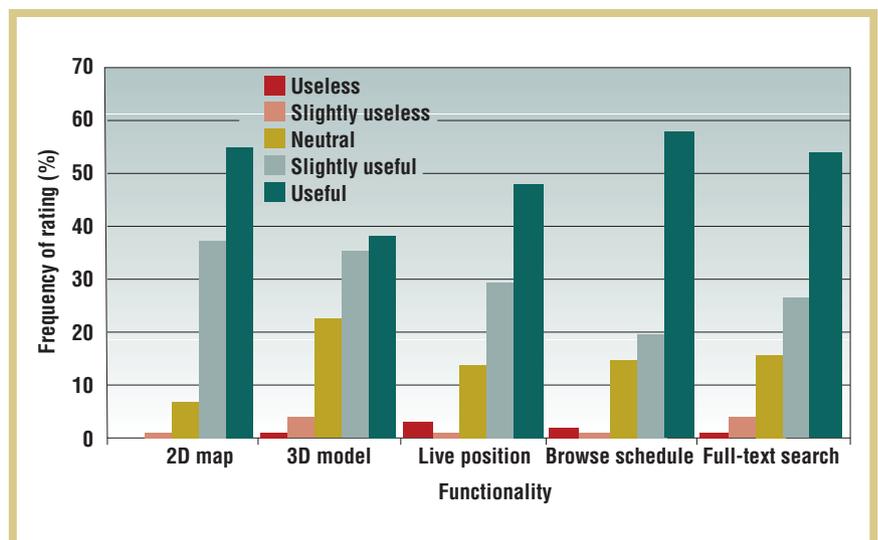
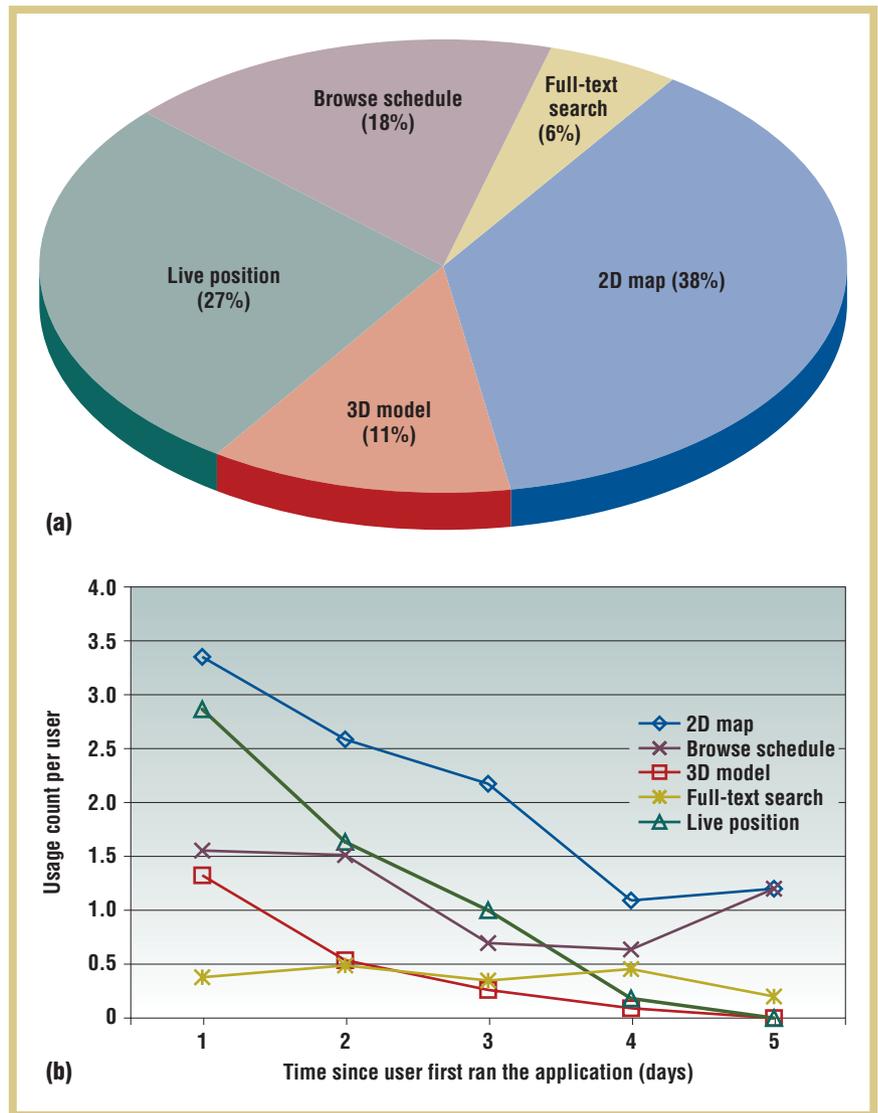
text search on the conference schedule. Figure 7a shows the percentage of times users invoked the various functions. At TechReady7, we also collected questionnaires from 64 users. In this questionnaire, users rated the usefulness of the five core functions compared with the printed conference booklet, on a Likert scale from 1 (useless) to 5 (useful), as Figure 8 shows.

In the following, we discuss the experiences gained throughout the four conferences.

Usefulness

At MEDC 2007, attendees found Signpost more useful than the conventional conference map that was part of the printed conference booklet (Figure 6a, Q1). At TechReady6, all subjects agreed that the application was very helpful. One user said, “Everything I looked for, I used the schedule. I used the map, since I didn’t know where to go in the Hyatt; that was very helpful.” From the questionnaires at TechReady7, we can also see generally high rates for Signpost’s usefulness. Most users quickly accessed the information presented on their device (Q4). Users consistently enjoyed the application (Q5), yielding an average score of 6.6 out of 7. Furthermore, the

Figure 8. Distribution of ratings on the usefulness of functionalities (compared to the printed conference handouts at TechReady7). All values are oriented toward positive ones (“useful”), but the ratings for the 3D model seem to be more spread and generally centered on “slightly useful.”



attendees strongly believed that Signpost should be used at other conferences as well (Q7), resulting in a score of 6.5 out of 7.

3D Overview Map

Overall, people found the 3D overview not very helpful. For example, one user at TechReady6 said, “It is cool. It is eye-candy, but for me it is not helpful. The 2D map is just fine.” Another said, “It was interesting, but I didn’t use it to try to find where I was going. I think the model is too simplistic.” This trend is also evident from the TechReady7 data, where the 3D functionality’s usage was very low, and the ratings for its usefulness were generally centered on 4 (“slightly useful”). A likely reason for this is that the conference area was too large. It spanned an area of four buildings, each with two to six levels. Still, all users voted to keep the 3D view, because of its “eye-candy” factor or because it gave them a better large-scale overview of the environment.

Small Screen

Most users had no problems seeing the overview, despite the small screen size, and they made comments such as, “That is perfect. I wouldn’t want to carry something bigger” or “I’d rather use my phone because I am used to [looking] at the

Navigation

To minimize the workload of adapting to new locations, we didn’t design Signpost to provide textual navigation instructions. Although users generally liked the way Signpost guided them, there were many suggestions for improvements. For example, one user said, “I think the biggest thing that would help me was if it would tell me steps: go down escalator, turn right ... like some of the car navigation things, but maybe not that precise.” As we expected from the controlled user study, users generally felt an improvement in their location awareness (Q6).

Tracking Accuracy

Although the marker-tracking system was quite precise, we decided to mount the markers only coarsely, to minimize the effort of mounting and measuring. Yet, users were generally satisfied with the tracking accuracy: “When I looked at it—immediately I thought wow, this is where I am.” One user referred to the tracking accuracy as “half a meter, for the purpose it was accurate enough ... 2 feet off the door versus 4 feet off the door really doesn’t matter.”

Fiduciary Markers

Using fiduciary markers in public areas always raises questions concerning visual clutter. However, most attendees didn’t complain about the mark-

with a black border (see Figure 1). In addition, they were mounted on special poster boards, so they were easy to spot. For TechReady6, we introduced the frame markers, branded with the design of the conference. Although they looked significantly more pleasing, they were far more difficult to spot. For TechReady7, we corrected this problem by using a different color scheme. At TechReady6 and TechReady7, the conference organizers deployed only 24 markers over an area of four buildings. Although the markers were placed prominently, they were small (15 cm). Surprisingly, all users commented that enough markers were available.

Privacy

With a system like Signpost, it would be possible to store all users’ current positions on a central server. Q3 asked about the users’ interest in seeing other users on their screen. Here, we noticed the highest variation among answers. Discussions with users confirmed our expectations that some users have concerns about their privacy being affected by such a feature. One user said, “At a conference it is OK ... Or you could just enter your conference ID.” Another suggested, “Oh, simply make it turn-off-able, when you don’t want to be localized.”

Other Use Cases

Users commonly agreed on alternative usage scenarios: “Inside buildings, malls, etc. There it would be useful. Outdoors you have GPS, which is already there, and people are used to it. The key thing would be inside.”

Technical problems

Most technical problems were related to erroneous camera drivers, which unfortunately are common on Windows Mobile phones. Moreover, to work on all Windows Mobile phones, Signpost must support all various screen resolutions, camera resolutions, and aspect ratios (see Figure 3).

Although users generally liked the way Signpost guided them, there were many suggestions for improvements.

small screen anyway.” Yet, another user said “Getting an overview is the toughest part, in the level of detail that is required.” As expected, some users found it difficult to browse a big map. Interestingly, this seemed to wear off as users gained a better feeling for the site.

ers, giving Q2 (“Those black-and-white markers disturbed me.”) a score of 1.7. Perhaps fiduciary markers don’t affect conference sites, which are already densely decorated with posters and screens, as much as other environments. At the first two conferences, the markers looked like a checkerboard

For MEDC 2007, we tried to supply configuration files for devices tested in advance, which turned out to be insufficient, owing to brand and model variety. For TechEd 2007, we implemented an automatic detection module for device capabilities, which improved the situation significantly. For TechReady6 and TechReady7, we also implemented a camera wizard that let users work around the most common driver bugs.

As we solved the most disruptive problems, new issues emerged. For example, power consumption due to camera usage initially didn't appear to be a problem. However, running Signpost with an active camera continuously in the background when the phone wasn't in use quickly drained the battery.

The combination of quantitative and multiple qualitative studies shows that marker-based indoor navigation provides advantages over simple, manually operated digital maps. These results also illustrate our approach's practicality in real-life use cases. To our knowledge, Signpost is the first indoor navigation system successfully deployed at several large-scale venues that runs on users' own mobile phones. We received encouraging feedback from our test audience, despite the reported technical difficulties.

Backed up by this positive feedback, Signpost is now a commercial product. It is adaptable to new events, thanks to authoring tools that make it easy to import existing floorplans. Tracking based on computer vision is cost-efficient, as it only requires placing a few posters with markers at the site rather than deploying an active beacon infrastructure. The use of commercial off-the-shelf camera phones lets users experience the application on their own devices, weaving navigation more intimately into everyday life.

In the future, we plan to compare our



Alessandro Mulloni is a PhD student at the Graz University of Technology. His research interests include 3D real-time graphics on handheld devices and human-computer interaction, especially user-centric design of interaction and visualization methods for handheld augmented reality. Mulloni has an MSc in computer science from the University of Udine. He is a student member of the IEEE. Contact him at mulloni@icg.tugraz.at.



Daniel Wagner is a postdoctoral researcher at the Graz University of Technology. His research interests include mobile augmented-reality technology, and real-time graphics and computer vision for mobile phones. Wagner has a PhD in computer science from the Graz University of Technology. He is a member of the IEEE. Contact him at wagner@icg.tugraz.at.



Istvan Barakonyi is a software developer at Imagination Computer Services GesmbH. His research interests include stationary and mobile augmented-reality applications, virtual reality, and embodied autonomous agents. He has a PhD in computer science from the Vienna University of Technology. Contact him at istvan.barakonyi@imagination.at.



Dieter Schmalstieg is full professor of virtual reality and computer graphics at the Graz University of Technology. He's also an advisor for the K-Plus Competence Center for Virtual Reality and Visualization in Vienna, deputy director of the doctoral college for confluence of graphics and vision, and director of the Christian Doppler Laboratory for Handheld Augmented Reality. His research interests include augmented reality, virtual reality, distributed graphics, 3D user interfaces, and ubiquitous computing. Schmalstieg has a PhD in computer science from the Vienna University of Technology. He is a member of the IEEE and the Austrian Academy of Science. Contact him at schmalstieg@icg.tugraz.at.

guidance system with paper maps. We see paper maps as a separate condition from the space (digital maps) we've thus far examined. Introducing another independent variable (paper vs. digital) will require a separate user study and far more test subjects. Further ongoing work includes integrating online marketing campaign material by using DataMatrix 2D barcodes as placeholders for Web links, and extending platform coverage. Our ultimate goal is for Signpost to evolve from a conference guide to a generic system for indoor navigation. The application we envision will support a broader range of venues and tasks, and will have an embedded scripting language to support direct downloads of new functions and scenarios to the system. ■

ACKNOWLEDGMENTS

This project was partially funded by the Austrian Science Fund (FWF) under contract no. Y193 and by the Christian Doppler Research Association (CDG).

REFERENCE

1. D. Schmalstieg and D. Wagner, "Experiences with Handheld Augmented Reality," *Proc. 6th IEEE/ACM Int'l Symp. Mixed and Augmented Reality (ISMAR 07)*, IEEE Press, 2007, pp. 3–18.

For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.