

Collaborative Augmented Reality for Outdoor Navigation and Information Browsing.

Gerhard Reitmayr and Dieter Schmalstieg,
(reitmayr|schmalstieg)@ims.tuwien.ac.at
Interactive Media Systems group,
Vienna University of Technology

Abstract

Augmented reality (AR) can provide an excellent user interface for visualization in a mobile computing application. The user's view is augmented with location based information at the correct spatial location, thus providing an intuitive way of presenting such information. In this work we demonstrate the use of AR for collaborative navigation and information browsing tasks in an urban environment. A navigation function allows one or more users to roam through a city and guides them to selected destinations. Information browsing presents users with information about objects in their surrounding. Both functions feature support for collaboration. The developed system does not only concentrate on the user interface aspects but also provides a scalable infrastructure to support mobile applications. To this end we developed a 3-tier architecture to manage a common data model for a set of applications. It is inspired by current Internet application frameworks and consists of a central storage layer using a common data model, a transformation layer responsible for filtering and adapting the data to the requirements of a particular applications on request, and finally of the applications itself.

1 Introduction

Many researchers believe that augmented reality (AR) is an excellent user interface for mobile computing applications, because it allows intuitive information browsing of location referenced information. In AR the user's perception of the real world is enhanced by computer generated entities such as 3D objects and spatialized audio. The interaction with these entities happens in real-time to provide convincing feedback to the user and give the impression of natural interaction. Augmented reality as a user interface becomes particularly powerful when the computer has access to location based information so that it can merge the virtual entities with real world objects in a believable manner.



Figure 1. A user is roaming a historical location in the City of Vienna.

Two common applications for mobile computing and location based services are navigation and information browsing based on location. The work described in this paper demonstrates the potential of using augmented reality as a user interface for such tasks. A tourist guide scenario motivated the specification of the requirements of the prototype (see section 3). Our work integrates a number of advanced user interface features demonstrated in earlier research systems into a single coherent application. Moreover, it also supports collaborative work for human communication among several mobile users.

Such applications require a detailed model of the user's environment. The model needs to include geometric representation of the environment and also semantic and contextual elements. Compared to conventional location based systems, AR not only requires integration of a wider variety of data sources to build interesting applications, it also creates new types of content. Geometric models of real objects are not only used for visualization purposes, but for computing occlusions, rendering of shadows, interaction, and vision-based tracking of real objects.

The complexities of AR content resulted in limited demonstrators that only worked within a well defined environment. In order to scale the working environment to be comparable to a real environment as experienced by tourists, we also had to address the complex modeling and data handling needs of ubiquitous augmented reality applications. Consequently, we present the concept of a three-tier application architecture based on XML as the enabling technology (see section 4). A central XML-based database stores a common model used for all applications. Using common XML tools, data is transformed and imported from different sources. Once in the database, the data can be maintained more easily and application or domain specific preprocessing operations can be applied. Then the data is transformed to a form directly useful to a single application. These transformations will often cull the model to return only those aspects of the data relevant to the application.

2 Related work

Navigation and information browsing are two common themes used for demonstrating wearable computing and mobile augmented reality. The Touring Machine by Feiner et al. (1997), the work described by Thomas et al. (1998), and the Context Compass by Suomela et al. (1992) show how pedestrian navigation can benefit from heads-up displays and AR interfaces. Information browsing was first demonstrated by Rekimoto (1997) with the Navicam and has since become a popular topic of AR applications. A notable example of a wearable tourist guide application is the GUIDE project described by Davies et al. (1999) that presents roaming tourists with location based information on Lancaster.

Augmented reality has also been identified as an important means of computer supported collaborative work (CSCW) as demonstrated by Billingham et al. (1996) in the Shared Space project. Collaboration between stationary and mobile users has been investigated by different groups in the MARS project (Höllner et al., 1999) and Tinmith project by Piekarski et al. (1999). Mobile users could also create world models in the BARS project (Baillot et al., 2001) but only in well controlled indoor environments. Our work extends these attempts to collaboration between several mobile users in an outdoor setting.

All recent AR demonstrations work with small data sets that have typically been entered manually and do not require data warehousing. Related work by Julier et al. (2000) addresses the issue of selecting appropriate data for display, from a user's point of view rather than that of the application. Höllner et al. (2001) describe a database and description logic used to store a model of the floor of a building which is annotated with meta-data for navigation target selection. The

sentient computing project described by Addlesee et al. (2001) uses a CORBA run-time infrastructure to model a live environment as distributed software objects where locations and attributes of objects are updated permanently. Newman et al. (2001) describe a set of AR applications based on this infrastructure. Another important research project is the Nexus project described by Rothermel and Leonhardi (2001) which deals specifically with the problem of establishing a world model for location- and context-aware applications developing an efficient solution to providing location based data. In contrast, our work focuses on modeling the complex information needs of AR applications and providing powerful tools for using such complex data by simplifying the transformations of the general model into application specific data structures.

3 Tourist guide application

The needs and requirements of a tourist are a suitable starting point for testing location based applications. A tourist is typically a person with little or no knowledge of the environment. However, tourists have a strong interest in their environment and also want to navigate through their surroundings to visit different locations. Guided tours present also a common practice for tourists. In such a situation a single person navigates a group of people and presents information.

Consequently, we have chosen a tourist guide application for the City of Vienna as an example scenario for an augmented reality application that integrates a large amount of data from different sources. It provides a navigation aid that directs the user to a target location and an information browser that displays location referenced information icons that can be selected to present more detailed information in a variety of formats. Both functions support collaboration between multiple mobile users.

3.1 Augmented reality system

Our current setup uses a notebook computer with a 2GHz processor and an NVidia Quadro4Go graphics accelerator operating under Windows XP. It includes a wireless LAN network adapter to enable communication with a second mobile unit. A Trimble Pathfinder Pocket differential GPS receiver is used to determine the position of the system in outdoor applications. All the equipment is mounted to a backpack worn by the user. We use a Sony Glasstron optical-see-through stereoscopic color HMD fixed to a helmet as an output device. An InterSense InertiaCube2 orientation sensor and a PointGrey Research Firefly camera for fiducial tracking are mounted on the helmet (see Figure 1).

We use *Studierstube* (Schmalstieg et al., 2002) which is based on Open Inventor (OIV) (Strauss and Carey, 1992) as a software platform for developing AR applications. It provides a multi-user, multi-application environment, and supports a variety of display devices including stereoscopic HMDs. It also provides the means of 6DOF interaction, either with virtual objects or with user interface elements displayed in the user's field of view. Applications are developed as scene graphs that can be described with the declarative OIV file format. *Studierstube* is a very capable rapid prototyping system, but does not incorporate any database functions beyond a scene graph based runtime data structure.

Collaboration between different users requires distribution of the application's state among different setups. To simplify development of such distributed applications, we implemented an extension to Open Inventor – Distributed Open Inventor described by Hesina et al. (1999) – that provides shared memory semantics on the scene graph data structure. Changes to a distributed part of the scene graph are transparently communicated to other instances of the application without exposing this process to the application programmer.




Figure 2. a) An overlay of the building model over the real world. b) 2D user interface components and HUD.

3.2 User interface

The system presents information to the user on the head mounted display. Such information is either presented as graphical objects rendered to fit into the natural environment or as text, images and 3D objects providing a heads-up display. The graphical objects are drawn to enhance and complement the user's perception of the natural environment. They can represent abstract information, alternative representations of real objects or highlighted real structures. The heads-up display is used to provide a graphical user interface consisting of typical 2D widgets such as buttons, lists, and text fields and to provide other status information. Figure 2 shows a typical view through the users display.

The user can control a cursor within the 2D user interface part in the upper right corner with a touchpad that is either worn on the belt or handheld. She can switch between different modes of the application such as navigation, information browsing and annotation. Each mode presents a number of individual panes to provide control of parameters and other options related to the current task. A general heads-up display (HUD) at the bottom of the view presents generic information such as the current location, selected target location, distance to the target and an orientation widget modeled after a compass.

3.3 Navigation

In navigation mode the user selects a specific target address or a desired target location of a certain type such as a supermarket or a pharmacy. The system then computes the shortest path in a known network of possible routes. It is interactive and reacts to the user's movements. It continuously re-computes the shortest path to the target if the user goes astray or decides to take another route.

The information is displayed as a series of waypoints that are visualized as cylinders standing in the environment. These cylinders are connected by arrows to show the direction the user should move between the waypoints. Together they become a visible line through the environment that is easy to follow (see Figure 3a). The user can enable an additional arrow that points directly from her current position to the next waypoint. Buildings can clip the displayed geometry to enable additional depth perception cues between the virtual information and the real world (see Figure 3b). Finally, simple directional information is displayed, if the user is not able to perceive next waypoint because she is looking into the wrong direction.

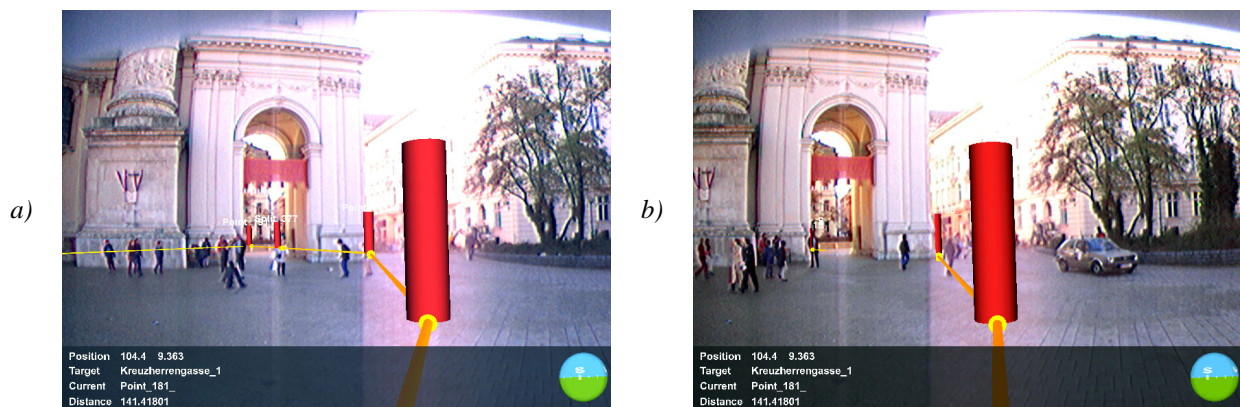


Figure 3. a) A visualization of the path to the selected target without clipping on known objects. b) The same path clipped at an object.

If two or more users are present, a number of collaborative interactions are possible. The interface will present a list of all users that have joined the collaboration session. Every user can select another user and specify an interaction mode:

Follow. The user can decide to follow the selected user. The navigation display will update the target location to always coincide with the waypoint closest to the selected user.

Guide. The user can guide the selected user by setting the destination point of the selected user. The navigation system of the selected user will then behave as if that user had selected the target herself.

Meet. This mode supports to meet halfway with the selected user. The navigation system calculates the meeting point to be halfway between the two waypoints the users are closest to. Then the destinations of both users are set to this new target. Each user can still change the common target to a more suitable location if desired.

3.4 Information browsing

The information browsing mode presents the user with location based information. Location referenced information icons appear in her view and are selected by looking at them. They then present additional associated information. The application conveys historical and cultural information about sights in the City of Vienna.

The information icons can have any shape for display as well as for ray intersection. In the current application we use geometric representations of parts of buildings to annotate these with cultural information. The icons appear to the user as outlines of the building parts. A virtual ray is cast through the center of the display and intersected with the information icons. The first icon that is hit triggers the display of information that is associated with it (see Figure 4).

The information consists of images and text which were taken from a guide book. These are shown to the user in the heads-up-display. The user can also select a subset of the icons to be active and visible by choosing a set of keywords the icons should relate to. The reduction of visible icons avoids visual clutter.

The information browsing mode supports multiple users. Users can choose to share their selection of topics, or alternatively, tour guides can control the selection for a group of guided users. A user can also trigger the highlighted information on another user's display.

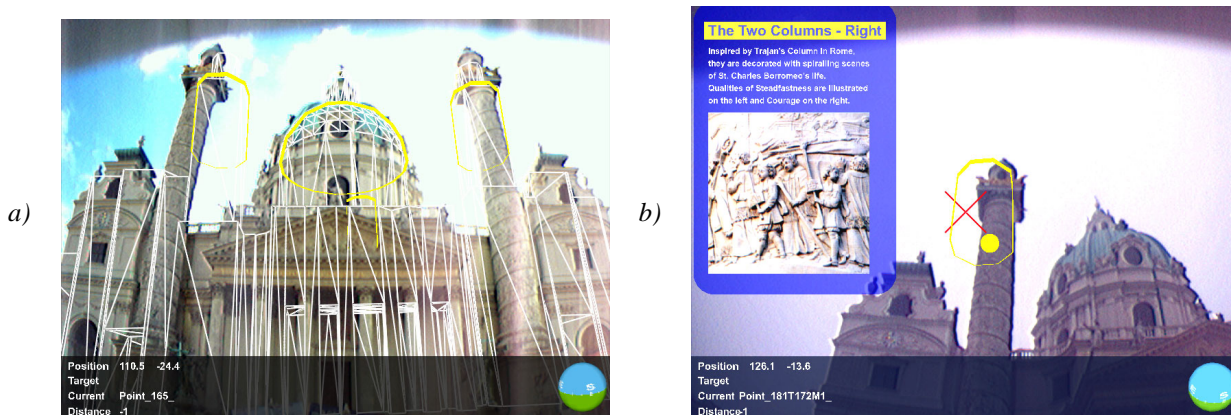


Figure 4. a) Different parts of the building are highlighted to show possible additional information. b) The user selects the column by looking at it and the content is displayed.

3.5 Annotation

In addition to pure browsing, users can also annotate the environment by placing virtual icons of different shapes and colors on structures and buildings in their surroundings (see Figure 5). Again, a virtual ray is cast through the cross hair in the heads-up display and intersected with the geometry of the buildings. If an intersection is detected a yellow sphere is placed at the intersection point to visualize it. Then the user can place a predefined 3D icon at the intersection point. The icon is oriented parallel to the tangent surface in the intersection point. The user can select between different predefined shapes and colors to use and can also choose which kinds of icons to display.

The virtual icons are shared between different users in a collaborative session and are annotated with the name of the user who created them. In this case a user can also include the name of the icons creator into the selection of visible icons.

The virtual markers can help to point out features on distant structures such as building facades. Users can attach and discern different meanings associated with markers by assigning different styles. They support collaborative work styles because they are shared information and can help users to communicate information about individual locations in the surrounding.

4 Data management

Any significant real-world application of mobile augmented reality will require a large model of location referenced data. While it may appear that a natural approach is to develop application specific data formats and management strategies, we have found that such an approach actually prevents reuse of the data and ultimately produces additional complexity in developing and maintaining the application. Therefore we developed a general model schema and a 3-tier architecture to maintain a central data store and application specific reuse through dedicated transformations. For a more detailed treatment of the data management concepts see Reitmayr and Schmalstieg (2003).

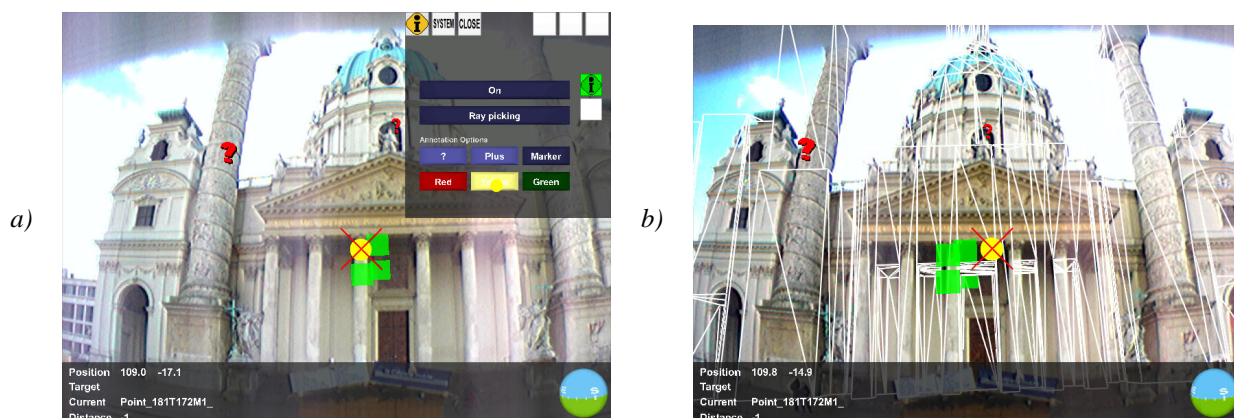


Figure 5. a) Annotation options and icons. b) The building geometry used for ray intersection is overlaid.

4.1 Three-tier architecture

The architectural concept is based on a typical 3-tier model. The first tier is a central database storing the overall model. The second tier mediates between database and application by converting the general model from the database into data structures native to the respective application. The applications themselves are the third tier and only deal with data structures in their own native format.

This architecture provides the usual advantages of the n-tier model. A common data model and store reduces the amount of redundancy in the data required for different applications and allows centralized and efficient management of this data. The middle layer separates the presentation issues from the actual data storage and the applications. Thus the applications can be developed using efficient data structures independent of the actual storage format. Moreover, the transformation can be adapted to changing data formats or processes without touching either the application or the storage back-end, because it is a distinct entity separated from both.

We built such an architecture based on XML technology, leveraging recent developments in the web development community where it is commonly deployed and directly supported in open-source or commercial products. The use of XML has a number of advantages for our task. An XML based model schema is powerful enough to capture the complex aspects of the data model. Transformation tools such as XSTL defined in Clark (1999) allow rapid prototyping and development of import, transformation and export tools to and from the data model. Parsers and generators exist for a wide range of programming languages, and allow applications and tools to use the most appropriate language for the task.

4.2 Modeling

At the heart of our architecture lies a data model that encompasses the requirements of all applications. Care was taken in keeping the model extensible so that new data could be integrated during the development. This data model is described by an XML schema.

The model is based on an object-oriented approach using a type hierarchy of entities. The root type is called *ObjectType* and contains an id and a generic annotation subelement that can be used to store any XML tree. All data types defined in the model are derived from this type. The *SpatialObjectType* adds pose information and a geometrical representation to the super class. We further derive the *SpatialContainerType* that adds a children subelement to aggregate entities of type *ObjectType* for hierarchical composition.

From the three base types, we derive a number of application specific types that are used in the actual data files. The *Object*, *SpatialObject* and *SpatialContainer* elements are used for general purpose data and correspond directly to the base types. Applications can define additional types derived from the base types to provide more specific information. For example, the *Waypoint* element has a specific sub element to define neighboring waypoints connected by a path. Because elements refer back to their base type, an application can always provide a reasonable fallback behavior if it encounters an unknown derived application elements.

4.3 Data handling

Having defined a model and data format, there are a number of tasks and tools necessary to fill the database, transform and manipulate the data and finally make it available to the user through the development of appropriate applications. The typical tasks include import from other data formats, maintenance of the model and export to application specific data structures. In the last step transformations are applied to retrieve application data from storage and generate data structures for the applications. As described in section 3.1, applications are implemented as Open Inventor scene graphs. Each application uses a custom XSLT style sheet to directly generate the required scene graph and additional data structures from the general model.

The use of a separate step to transform the data into the application format has a number of advantages. It separates the general data format from the application specific data structures and provides an interface between both. It also provides a point for customizing the presentation independently of the application, similar to the way traditional cascading style sheets work for HTML. As the style sheet generates the actual graphical content, it can adapt it to different output requirements or profiles.

4.4 Data acquisition

The described application requires diverse data from a variety of sources. A 3D model of a part of Vienna was obtained from the cartography department of the city administration in 3D Studio Max format. This model was created as a reference model of Vienna, and is part of the official general map of the city (Dorffner and Zöchling, 2003). The department of Geoinformatics at Vienna University of Technology supplied us with a network of accessible routes for pedestrians, delivered in GML2 format defined in Cox et al. (2001), an XML based file format used in the geographic information systems (GIS) community. This model was derived from the general map of Vienna and is represented as an undirected graph. Each node in this graph is geo-referenced and used as a way point in the navigation model. For each building, a so-called address point was defined and included into the path network to be able to construct a path to this address.

Furthermore, annotation information such as businesses located at certain addresses was derived from the general map of Vienna. This information is connected to address points in the spatial database. Cultural information taken from a guide book was included at various places to provide more detailed data for a tourist application. Finally, we placed the icons as spatial representations of this information into our model.

5 Results

The described system was outfitted with data surrounding the area of Karlsplatz and adjacent building blocks down to Gußhausstraße in Vienna. The area of Karlsplatz proved to be a good

testing ground because it is open enough to allow reception of GPS signals for positioning, has a somewhat complex network of foot paths through the Resselpark and a number of famous tourist attractions such as the Karlskirche are situated at its border. Finally, it is close enough to our institute to allow frequent visits for development and testing purposes. The images throughout the paper were captured during such test runs. We use a video see-through approach and render the image of the video camera into the background to simulate the experience of a real user.

During November 2003 we showed the prototype to a group of colleagues in an informal demonstration session. While they have been exposed to AR environments before, they did not know about the user interface of the tourist guide demonstrator. After an initial acclimation phase, users found the user interface of looking at objects and walking through the environment simple enough. The interaction with the 2D heads-up display interface was useable, but can be improved because the controls were not tailored specifically to the tasks.

The accuracy of the GPS positioning is not satisfying and the lack of precision leads to jumps in the graphics overlay. The blocking of satellites by high buildings prohibits the use of the system in the smaller side streets and close to buildings. Still it is the only technology available that provides the minimal required accuracy and robustness.

6 Conclusion and future work

The presented work tries to give an outlook into possible future user interfaces for location based services. Although many of the implemented user interface features have been demonstrated before, our work exceeds former work in two areas. Firstly, the collaboration features add another dimension to the possibilities of such systems by supporting groups of users in their tasks. While the features of the system are implemented, we need more tests of the collaborative aspects to determine useful and interesting extensions. Secondly, the integrated approach to handling the data required by a location based service allows the system to scale to realistic environments. The use of a flexible data model also simplifies extending the system with future applications that will have new requirements for data to be stored in the repository.

The application of a generic AR software framework such as *Studierstube* allowed us to leverage the graphical and interaction possibilities of modern 3D rendering libraries and simplified development to a great extent. Thus we did concentrate on the problems of scaling the data management and implementing collaborative features into the system. Moreover, we could enhance the user interface at every step during development because changes to the presentation and interaction can be implemented and tested within short turn around cycles. Still we did not fully explore the design possibilities because of the overwhelming number of options that presented themselves to us.

We plan to use DGPS in the future and distribute RTCM correction data over a GPRS connection from an Internet server to the mobile setups. Additionally a Kalman filter should reduce the noise in a static position while still reacting quickly to changes when the user is moving. Another interesting possibility is the use of velocity and direction information to auto-calibrate the orientation tracker. Under the assumption that the user is looking into the direction she is moving, we can update an error description of the orientation tracker from the direction information provided by the GPS receiver.

Acknowledgements

This work was sponsored by the Austrian Science Foundation *FWF* under contracts no. P14470 and Y193, and Vienna University of Technology by *Forschungsinfrastrukturvorhaben*

TUWP16/2002. We would like to thank the Vienna City Administration for providing the 3D model of Vienna, Prof. Stephan Winter for the navigation model and Joe Newman for setting up the DGPS system.

References

- Adlesee, M., Curwen, R., Hodges, S., Hopper, A., Newman, J., Stegges, P., and Ward, A. (2001). A SENTIENT COMPUTING SYSTEM. In IEEE Computer: Location-Aware Computing.
- Baillet, Y., Brown, D., and Julier, S. (2001). AUTHORIZING OF PHYSICAL MODELS USING MOBILE COMPUTERS. In Proc. ISWC 2001, pages 39-46.
- Billinghurst, M., Weghorst, S., and Furness, T. (1996). SHARED SPACE: AN AUGMENTED REALITY INTERFACE FOR COMPUTER SUPPORTED COLLABORATIVE WORK. In Proc. Collaborative Virtual Environments Workshop '96, Nottingham, Great Britain.
- Clark, J. (1999). XSL TRANSFORMATIONS (XSLT) VERSION 1.0. <http://www.w3.org/TR/xslt>.
- Cox, S., Cuthbert, A., Lake, R., and Martell, R. (2001). GEOGRAPHY MARKUP LANGUAGE (GML) 2.0. <http://opengis.net/gml/01-029/GML2.html>.
- Davies, N., Cheverest, K., Mitchell, K., and Friday, A. (1999). 'CACHES IN THE AIR': DISSEMINATING TOURIST INFORMATION IN THE GUIDE SYSTEM. In Proc. WMCSA'99, New Orleans, LA, USA. IEEE.
- Dorffner, L. and Zöchling, A. (2003). DAS 3D MODELL VON WIEN - ERZEUGUNG UND FORTFÜHRUNG AUF BASIS DER WIENER MEHRZWECKKARTE. In Proc. CORP 2003, pages 161 - 166, Vienna, Austria.
- Feiner, S., MacIntyre, B., Höllerer, T., and Webster, A. (1997). A TOURING MACHINE: PROTOTYPING 3D MOBILE AUGMENTED REALITY SYSTEMS FOR EXPLORING THE URBAN ENVIROMENT. In Proc. ISWC'97, pages 74-81, Cambridge, MA, USA.
- Hesina, G., Schmalstieg, D., and Purgathofer, W. (1999). DISTRIBUTED OPEN INVENTOR : A PRACTICAL APPROACH TO DISTRIBUTED 3D GRAPHICS. In Proc. ACM VRST'99, pages 74-81, London, UK.
- Höllerer, T., Feiner, S., Terauchi, T., Rashid, G., and Hallaway, D. (1999). EXPLORING MARS: DEVELOPING INDOOR AND OUTDOOR USER INTERFACES TO A MOBILE AUGMENTED REALITY SYSTEM. *Computer & Graphics*, 23(6):779-785.
- Höllerer, T., Hallaway, D., Tinna, N., and Feiner, S. (2001). STEPS TOWARD ACCOMMODATING VARIABLE POSITION TRACKING ACCURACY IN A MOBILE AUGMENTED REALITY SYSTEM. In Proc. AIMS'01, pages 31-37, Seattle, WA, USA.
- Julier, S., Lanzagorta, M., Baillet, Y., Rosenblum, L., Feiner, S., and Höllerer, T. (2000). INFORMATION FILTERING FOR MOBILE AUGMENTED REALITY. In Proc. ISAR 2000, Munich, Germany. IEEE
- Kretschmer, U., Coors, V., Spierling, U., Grasbon, D., Schneider, K., Rojas, I., and Malaka, R. (2001). MEETING THE SPIRIT OF HISTORY. In Proc. VAST 2001, Glyfada, Athens, Greece. Eurographics.
- Newman, J., Ingram, D., and Hopper, A. (2001). AUGMENTED REALITY IN A WIDE AREA SENTIENT ENVIRONMENT. In Proc. ISAR 2001, pages 77-86, New York, New York, USA. IEEE and ACM.
- Piekarski, W., Gunther, B., and Thomas, B. H. (1999). INTEGRATING VIRTUAL AND AUGMENTED REALITIES IN AN OUTDOOR APPLICATION. In Proc. IWAR'99, San Francisco, CA, USA. IEEE CS.
- Reitmayr, G. and Schmalstieg, D. (2003). DATA MANAGEMENT STRATEGIES FOR MOBILE AUGMENTED REALITY. In Proc. STARS 2003, pages 47-52, Tokyo, Japan.
- Rekimoto, J. (1997). NAVICAM: A MAGNIFYING GLASS APPROACH TO AUGMENTED REALITY. *PRESENCE - Teleoperators and Virtual Environments*, 6(4):399-412.
- Rothermel, K. and Leonhardi, A. (2001). MAINTAINING WORLS MODELS FOR CONTEXT-AWARE APPLICATIONS. In Wells, W. R., editor, Proceedings of the 8th International Conference on Advances in Communication and Control (COMCON-8).
- Schmalstieg, D., Fuhrmann, A., Hesina, G., Szalavari, Z., Encarnação, L. M., Gervautz, M., and Purgathofer, W. (2002). THE STUDIERSTUBE AUGMENTED REALITY PROJECT. *PRESENCE - Teleoperators and Virtual Environments*, 11(1).
- Strauss, P. and Carey, R. (1992). AN OBJECT ORIENTED 3D GRAPHICS TOOLKIT. In Proc. SIGGRAPH'92.
- Suomela, R. and Lehtikoinen, J. (2000). CONTEXT COMPASS. In Proc. ISWC 2000, Atlanta, Georgia, USA. IEEE.
- Thomas, B. H., Demczuk, V., Piekarski, W., Hepworth, D. H., and Gunther, B. (1998). A WEARABLE COMPUTER SYSTEM WITH AUGMENTED REALITY TO SUPPORT TERRESTRIAL NAVIGATION. In Proc. ISWC'98, pages 168-171, Pittsburgh, USA. IEEE and ACM.

Gerhard Reitmayr is a PhD candidate at the Vienna University of Technology, Austria. His PhD topic is software design for mobile collaborative augmented reality. He is involved in developing mobile augmented reality systems for indoor and outdoor use. His current research interests are data models for augmented reality, software framework support for tracking, rapid prototyping of user interfaces and distributed AR applications. He received a Dipl.-Ing. degree of engineering mathematics from Vienna University of Technology in 2000. He serves as lecturer for the virtual reality lab exercise and supervises lab work and master thesis students on a regular basis.

Dieter Schmalstieg is an associate professor at Vienna University of Technology, Austria, where he directs the "Studierstube" research project on collaborative augmented reality. His current research interests are virtual environments, augmented reality, three-dimensional user interfaces, and distributed graphics. He received Dipl.-Ing. (1993), Dr. techn. (1997) and Habilitation (2001) degrees from Vienna University of Technology. He is author and co-author of over 60 reviewed scientific publications, editorial advisory board member of computers & graphics, local chair of EGVE'99, general chair of ISMAR 2002 and program chair of ISMAR 2003. In 2002, he received the START award, one of five 6-year career grants annually awarded by the Austrian Science Fund to excellent young researchers.