

A Rapid Prototyping Software Infrastructure for User Interfaces in Ubiquitous Augmented Reality

Christian Sandor, Gudrun Klinker
Technische Universität München
Institut für Informatik
(sandor,klinker)@in.tum.de

Abstract

Recent user interface concepts, such as multimedia, multimodal, wearable, ubiquitous, tangible, or augmented reality-based interfaces, each cover different approaches that are all needed to support complex human-computer interaction. Increasingly, an overarching approach towards building what we call Ubiquitous Augmented Reality user interfaces which include all of the just mentioned concepts will be required. To this end, we present a user interface architecture that can form a sound basis for combining several of these concepts into complex systems. We explain in this paper the fundamentals of DWARF's user interface framework, an implementation of this architecture. Finally we present several examples that show how the framework can form the basis of prototypical applications.

Keywords Augmented Reality, Ubiquitous Computing, Tangible User Interfaces, Multimodality, Software Architectures, Frameworks, Mobile Systems

1 Introduction

One of the major challenges of current computer systems is to provide users with suitable means to plan, model, and control complex operations which are composed of many inherently interdependent processes. For example, control rooms of industrial plants, surgery preparation rooms, cockpits of airplanes, and consoles of modern cars are typically equipped with many different physical or electronic input and output devices.

Recent user interface concepts, such as *multimedia*, *multimodal*, *wearable*, *ubiquitous*, *tangible*, or *augmented reality-based* interfaces, each cover different approaches. We believe all of these approaches are necessary to support increasingly complex human-computer interaction. Increasingly, an overarching approach towards building *Ubiquitous Augmented Reality* (UAR) user interfaces which include all of the just mentioned concepts might be required. An instance of an UAR user interface can be seen as an aggregation of these conceptual aspects of interaction design. But which aspects should be used for a task at hand? To allow interaction designers to quickly change the aspects that are used, we implemented a software infrastructure that allows the rapid exchange of interaction styles.

1.1 Current User Interface Paradigms

Current user interface research addresses several different issues of human-computer interaction: multi-channel communication between users and computers, user mobility and the three-dimensional combination of virtual and real worlds. Each of these issues describes a dimension in a design space of future human-computer interfaces.

Multi-channel Communication

To provide human-computer interaction beyond traditional WIMP-based user interfaces [45], various communications channels are being explored that correspond more naturally to the human visual, auditory and tactile senses, both for gathering user input (speech, gestures, special 3D input devices) and for providing output (sound, graphics, haptics) to the user. According to Nigay et. al. [29], multimodality is the capacity of a system to communicate with a user along different types of communication channels and to extract and convey meaning automatically.

Both multimedia based and multimodal systems use multiple communication channels. Research on *multimedia* based user interfaces focuses on handling the vast amount of data that is required to gather raw input streams and to generate raw output streams in real-time. *Multimodal systems*, on the other hand, are defined at a higher level of abstraction. They strive towards associating semantic meaning with media streams. They are expected to help users control systems more easily by combining several interaction modalities into more powerful interaction paradigms than any single modality would be able to provide on its own [35, 54]. There are four different types of multimodality (exclusive, alternate, concurrent, and synergistic), depending on the combined or independent fusion of several interaction channels, as well as on the sequential or parallel use of multiple modalities [29].

Although multimedia based and multimodal systems have much in common, they cannot be described as one being a subset of the other. Many of today's internet browsers and email systems provide multimedia based functionality without being multimodal. On the other hand, multimodal systems focus more on the synergistic high-level interpretation of a few combined and parallel input tokens.

Mobility

Current trends towards mobile systems enable users to communicate with their computers while they are far away from their desks. Such mobility requires light-weight and untethered solutions that allow users to roam freely in a wide area. There are two approaches that address these requirements.

Wearable user interfaces [44] strive towards providing users with light-weight, portable or wearable computer equipment that they can become part of their daily attire. Wearable functionality can be provided on palmbased computers or mobile phones. It can be attached to a belt or woven into users' clothes. Using a personal area network (PAN¹), computing power can be connected to wearable gadgets, such as headsets, microphones, head-mounted displays (HMDs), 3D mice, and portable keyboards.

Ubiquitous [55, 26], *ambient* [10] and *pervasive* [14] interfaces to computers have been proposed by Weiser and others with the goal of providing computing power to people in such a pervasive manner that the computer in itself becomes a secondary (virtually invisible) issue. Large-scale environments such as buildings are equipped with networked computers and multi-channel user interfaces such that users are always surrounded by them. Research in this field focuses on developing proper system layouts for such large-scale computer networks, requiring high data bandwidths and system adaptivity to changing user demands. Adhoc interoperability of services is needed [23] in order to build context-aware smart spaces into which wearable, smart appliances can be integrated to provide users with personalized and context-adapted information.

Wearable and ubiquitous computing are two different approaches along a continuum of options in the mobility spectrum. They are not mutually exclusive, but tend to favor different trade-offs between the low-power, individualized computer use on a wearable computer and high-performance computing in a server-like, community-based stationary environment. Current trends begin combining both approaches.

Interaction Embedded within the Real World

User mobility provides the unique opportunity to let users communicate with their computer system while remaining involved in their three-dimensional, real-world environment. This provides the chance to let users communicate with computers via the interaction with position-tracked real physical objects. To this end, two complimentary approaches have emerged: tangible user interfaces and augmented reality.

Tangible user interfaces (TUIs) build on the observation that century-old, very well-designed physical tools exist, e.g. in craftsmanships, that have been fine-tuned for years towards very specific usage. Based on humans' spatial and motor skills, each such tool is directly suited towards fulfilling specific tasks. The purpose of each tool is immediately obvious to a trained craftsman. It represents a unique combination of input and output behavior

¹<http://grouper.ieee.org/groups/802/15/>

that is directly suited to the task it has been designed for. The TUI community strives towards providing similarly powerful tangible user interfaces for computers and their interaction with virtual information. Ishii's work with the MIT Tangible Media Group has produced a large number of creative tangible user interfaces, e.g. [50]. A formal model of TUIs is provided in [51].

Augmented Reality (AR) focuses on presenting information in three dimensions with respect to the user's current position. Users can thus see, explore and manipulate virtual information as part of their real world environment. In his classical definition of AR, Azuma [2] states three requirements: real-time performance, user registration in three dimensions and a combined presentation of both virtual and real information. In Milgram's taxonomy [28], AR is seen as a subset of *mixed reality*.

Tangible user interfaces and augmented reality overlap considerably with respect to the manipulation of real objects. Yet, for AR, this may be a pure consequence of augmentations issuing instructions provided by an AR application, e.g. for machine maintenance or repair, while TUIs consider such tracked object manipulation as the main means for users to control the system. On the other hand, such TUI-based interaction does not necessarily have to result in geometrical (visual or aural) augmentations of the 3D world. Illustrations may be provided in a lower-dimensional space, such as on a 2D table top or on a wall. AR, on the other hand, focuses on analysing and presenting information in three dimensions.

1.2 Convergence of Paradigms

Many of the above-mentioned research directions are currently broadening their spectra to include important aspects of other concepts, thereby generating a confluence of the individual fields. Examples of those broadened approaches are e.g. Tangible Augmented Reality [20] and Multimodal Augmented Reality [19]. We refer to this emergent area of user interfaces Ubiquitous Augmented Reality (UAR).

The idea to focus on the overall user experience by selecting the right interactions and visualizations has already been proposed by Buxton [8]. He calls this approach holistic, referring to the idea of the whole being more than the sum of the parts.

Obviously this overarching, holistic approach introduces new challenges to both the interaction designers and a supporting software infrastructure. Interaction designers have to choose from a broader variety of possible interaction styles to use. On the other hand, a supporting software infrastructure has to be capable of dealing with a wide range of interaction styles. Also it has to enable the interaction designer to quickly try out new interaction styles. Additionally a good software infrastructure should be extensible so that new interaction styles can be added later on. Apart from the interaction style mentioned in the last section, several other styles are currently being incorporated into our infrastructure: Zoomable User Interfaces [36], Attentive User Interfaces [30, 53] and Perceptive User Interfaces [49].

1.3 A Supporting Software Infrastructure

To progress in this direction, we have built a framework for UAR user interfaces, which is based on a tangible tool metaphor and allows the multimodal construction of higher-level interaction metaphors. The framework offers a tool chest containing a set of tools, each providing a specific functionality to the user. By composing more complex tools out of the simple basic toolset higher level functionality can be achieved. This allows users to manage any complex, inter-related processes, using a number of physical objects in their surroundings.

The framework can be used for single-user as well as multi-user applications. The system state is presented as a three-dimensional augmentation that is embedded within the real environment of a user. Scene views are provided via several both personal and ubiquitously available displays, accounting for different options in user mobility and privacy. Some views are common to all users (e.g: in the form of projections on a table or wall), others are restricted to subgroups of users (shown on portable display devices) or to a single user (shown on a head-mounted display). The displayed content depends on the current position of a display in the scene, representing its current viewpoint.

Users are provided with several tangible objects which they can manipulate together in a group or individually to influence the system state. Some of these objects are owned by specific users, and the interaction style with these objects can also be tailored to the preferences of individual users.

1.4 Organization of the Paper

The main design goal of our user interface framework is rapid prototyping, collection and reuse of different interaction elements. The next sections explain how this is achieved. Section 2 provides an overview of our approach, introducing the technical requirements for UAR user interfaces, the DWARF framework and our user interface architecture which lay the foundation for our work. Section 3 presents four prototypical examples showing increasingly complex combinations of user interaction concepts. Section 4 provides a summary and discusses future directions.

2 Our Approach

In the last section the three main characteristics for UAR user interfaces were presented: mobility, multi-channel communication and interactions that are embedded in the real world. Based on these characteristics, we now discuss the implications for a supporting software framework that addresses UAR user interfaces. First we discuss the technical requirements in more depth and build an analysis model on the software level to have a reference frame for further explanations. Then we point out how our DWARF framework (Distributed Augmented Reality Framework²) generically addresses UAR user interfaces. Next, more details on the specific user interface concepts within DWARF are presented. At the end of this section, we briefly discuss the benefits and limitations of our approach.

Probably the most similar approach to ours is the iRoom [18] project from Stanford University. Their work shares some common concerns. They have built a highly dynamic software infrastructure for mobile and distributed user interfaces based on tuplespaces [17]. However they do not address tangible interactions and augmented reality interfaces. Instead they focus on more conventional input like pen-based interactions, whereas the output they are mainly concerned with is wall-sized displays. As a result, iRoom does not couple real world objects to visualizations and thus does not provide a data flow framework for continuous integration as introduced in the next section.

2.1 Technical Requirements for Ubiquitous Augmented Reality User Interfaces

It is common practice in software engineering to build an analysis model of a problem domain. After the main requirements are established and the analysis model has been formulated, this model serves as a reference frame for discussions. The general functional decomposition that is presented here is the result of our analysis of the problem domain.

Mobility The main requirement for a mobile system running in ubiquitous environments is flexibility. Resources in the environment have to be connected dynamically to resources a user might carry around, e.g. palmtop computers or wearables like MITHRIL³. This implies a highly modular architecture, whose components should be dynamically reconnectable.

Multi-channel Communication To address multimodality, a system has to be able to deal with several input channels. The user intention has to be extracted from the input that is received over these channels. Complementarily, a multimedia based system has to have a coordination instance that distributes the content to be presented to the user leveraging the available output channels.

Interaction Embedded within the Real World For AR user interfaces and TUIs, a proper three-dimensional registration between real and virtual objects has to be established. To this end mobile real world objects have to be tracked, i.e. their position has to be determined continuously. A typical example for AR involves tracking a user's head and using its pose to set the viewpoint of a virtual scene as seen in a head-mounted display. Similarly, TUIs often couple real world objects to virtual representations of objects.

Based on these requirements, we propose a generic functional decomposition of UAR user interfaces. A large number of AR frameworks have recently been analyzed (see [40]). As most of these frameworks also support multi-channel communication and mobile systems, the findings made in that analysis helps establish a foundation for an analysis model covering UAR user interfaces. We have consolidated the recurring patterns into this generic

²<http://www1.in.tum.de/DWARF/>

³<http://www.media.mit.edu/wearables/mithril/>

functional decomposition. Figure 1 (inspired by [27]) shows the relevant subsystems and components within them. It is important to note that the subsystems (Input Devices, Media Analysis, Interaction Management, Media Design and Output Devices) are general purpose and generic, however the components within them are just examples. Similarly DWARF is one possible implementation of a framework enabling UAR user interfaces. Other implementations adhering to the functional decomposition in figure 1 would be possible.

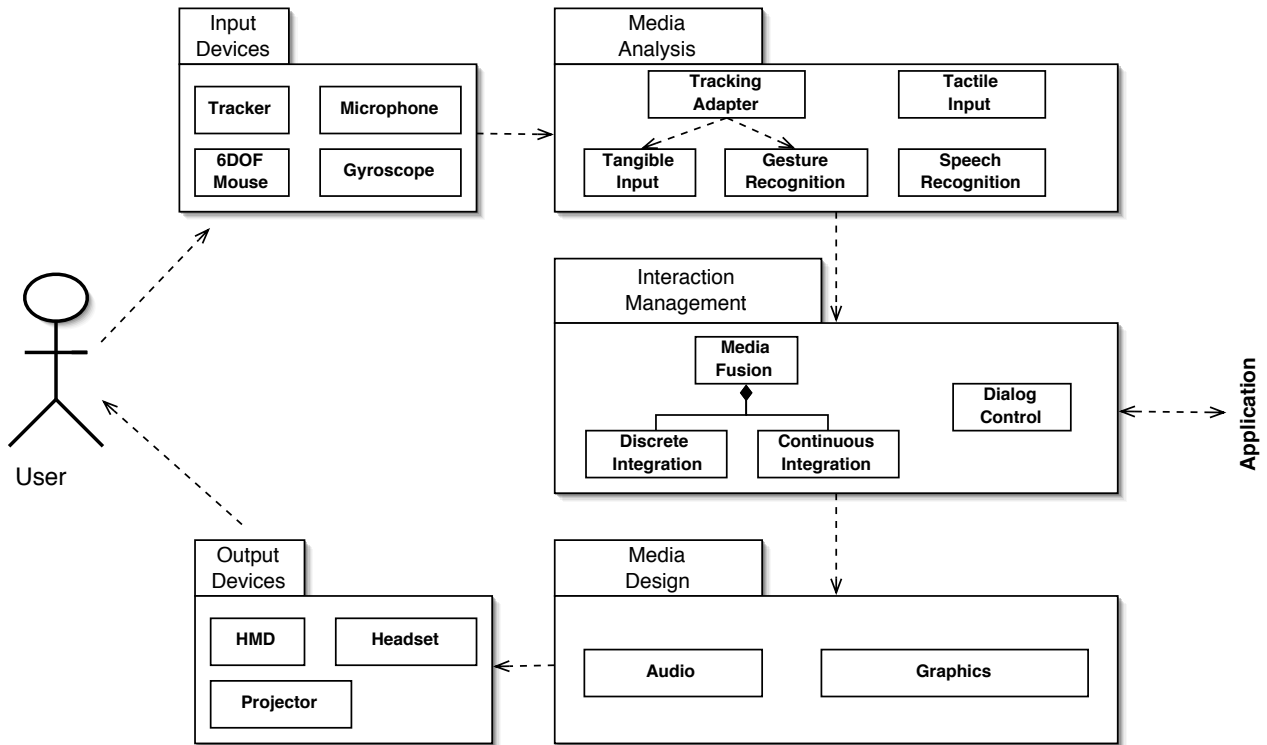


Figure 1. A generic functional decomposition of Ubiquitous Augmented Reality User Interfaces

The *Input Devices* subsystem contains input devices that are used to receive commands from the user. Each of these devices offers an individual input modality to be evaluated by the multimodal user interface. The *Output Devices* subsystem renders the signal on the specified output devices. For multimedia based systems several output devices are used at the same time.

Media Analysis is the process of turning physical user input into abstract tokens [34] handed over to the subsequent parts of the system – this can be compared to the task performed by the lexical analysis of a compiler. Separate classes, such as gesture analysis, speech analysis, and tangible input analysis deal with the specific properties of different input modalities of the input devices. The software components that present content to the user over any of the cognitive channels, e.g. visual and aural are contained within the *Media Design* subsystem. The *Interaction Management* subsystem determines which output is presented to the user. Current challenges for interaction management are performance, flexibility, adaptivity, usability and an efficient error management. The *Media Fusion* component takes the tokens of several input channels and infers user intention from them. In this component, two different ways for combining different input channels under respective boundary conditions are considered. *Continuous Integration* combines tokens that can take real values in a certain range, e.g. rotations can take an infinite number of different values between 0 and 360 degrees. Example input devices that deliver these kinds of tokens are mice, trackers and gyroscopes. A variety of frameworks [32, 41, 39] already exist that ease the task of building data flow networks for continuous integration, however they do not take into account the distribution over several hosts and the continuous dynamic reconfiguration that is necessary for UAR. *Discrete Integration* refers to the integration of devices like speech recognition that deliver only discrete values like the word

that was recognized. The most similar approach to our model for Discrete Integration is [15], however in contrast to us they focus on static setups. Finally, the *Dialog Control* component selects the presentation medium and what to present over it.

2.2 The Distributed Wearable Augmented Reality Framework (DWARF)

For the past few years, we have been building a general, reusable and easily (ad-hoc) configurable Distributed Wearable Augmented Reality Framework called DWARF [3, 4]. DWARF describes different contributing system parts as separate components that are able to connect with one another across a dynamically configurable peer-to-peer network of distributed processes. Whenever new components are in reach, they get connected automatically into the network of communicating components. The connectivity structure of components is not fixed at startup time. In fact, it can be changed arbitrarily at runtime.

DWARF is suitable for building highly dynamic, flexible system arrangements within which mobile users, who carry mobile sensors and devices, can be connected on demand to stationarily available, ubiquitous resources that are provided within intelligent environments of the future.

DWARF makes it possible to develop mobile applications, because of the following features:

- Flexible architecture** Because of the fine granularity of components and the loose coupling between them, DWARF systems are highly flexible.
- Fast** Several communication protocols are implemented for the communication between components. Some of them are especially well suited for real-time applications, e.g. shared memory and CORBA events.
- Distributed** The components that form a DWARF system can be a combination of local and remote devices. Distribution is completely transparent to the components.
- Adaptivity** With the inherent options for ad-hoc connections and reconfiguration of components, DWARF systems are also inherently adaptive.
- Operating System independent** To allow deployment among a variety of devices, DWARF has been designed to be independent of a specific operating system. We have successfully implemented DWARF systems on Linux, Windows and Mac OS X platforms.
- Programming language independent** Similarly, DWARF supports three programming languages so far: JAVA, C++ and Python.

We have already built around ten systems based on DWARF. For a thorough list, please refer to our projects webpage⁴. It can be downloaded from our webpage⁵. Based on DWARF's flexible nature, we have developed a user interface architecture with a supporting set of components that are described in the next section.

2.3 The User Interface Architecture in DWARF

DWARF was designed as a research platform combining wearable systems with ubiquitous environments. Its component model and architecture can be used in several different research areas. In this section we explain several architectural principles and components that make up the user interface framework within DWARF. An overview of the architecture can be seen in figure 2. An important distinction for communication channels is the frequency with which messages are passed. *Discrete Events* are typically sent every few seconds, whereas *Continuous Events*, such as tracking data, are sent at very high frequencies.

⁴<http://www1.in.tum.de/DWARF/ProjectsOverview>

⁵<http://www1.in.tum.de/DWARF/>

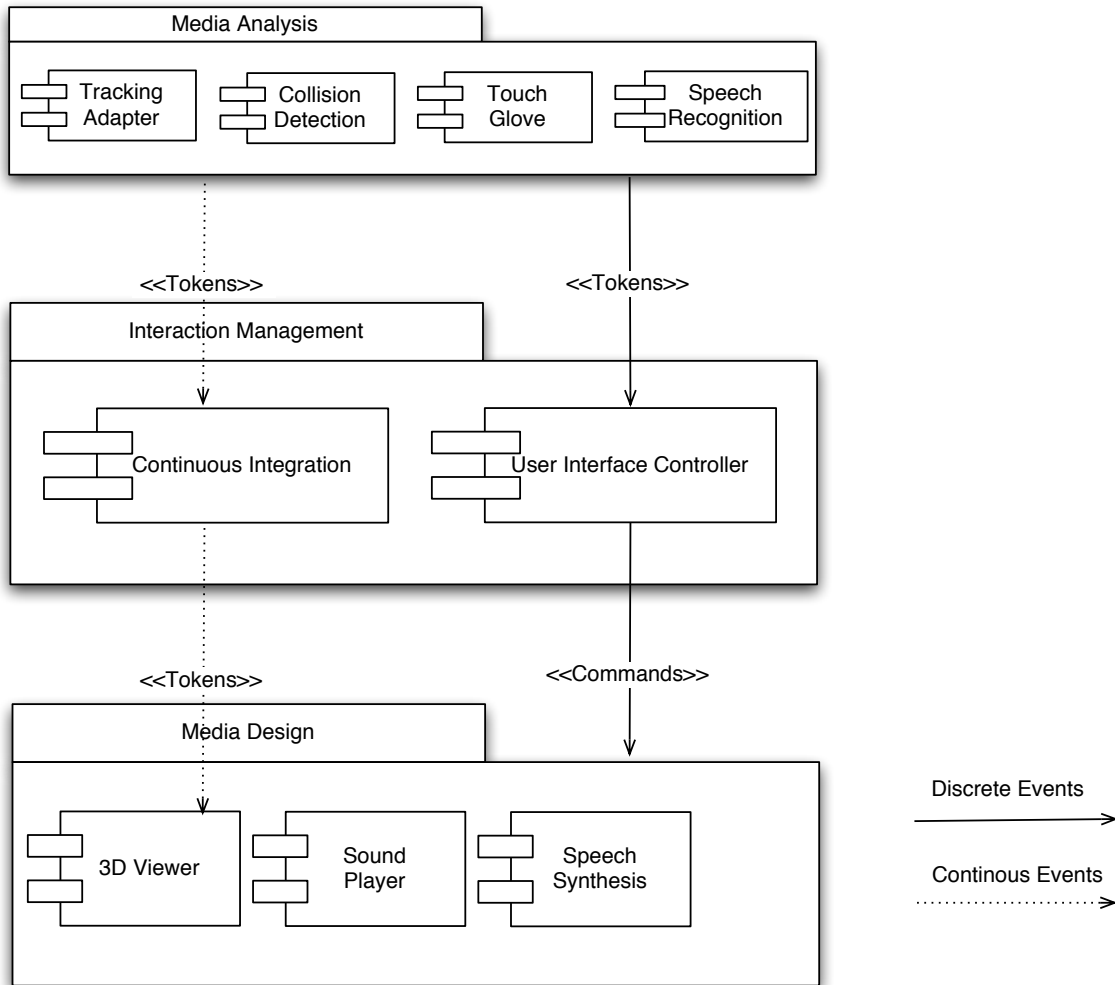


Figure 2. Functional decomposition of DWARF specific user interface components

Layering and Device Abstraction

We arrange the UI components in three layers according to figure 2. Most data flows linearly from the *Media Analysis* layer which contains input components to the *Interaction Management* layer where the tokens are interpreted. From there the data flow continues to *Media Design* layer where the output components reside.

We have developed a standardized format for tokens that are sent from the input components to the Interaction Management layer. Input tokens can be decomposed into four different types: analog values that can be either within a limited range (e.g., rotations) or an unlimited range (e.g translations) and discrete values that can be either booleans (e.g., pressing a button) or text strings (e.g., the output of a speech recognition process). Due to this standardized format, we can exchange one input device for another – as long as they emit the same type of tokens. For example, a speech recognition component listening for a set of words could be interchanged transparently with tangible buttons with the same set of labels.

Similarly, the Interaction Management layer sends commands to the Media Design layer. This setup corresponds to the *command pattern* described by Gamma et. al. [13]. The commands consist of actions that have to be executed by the output components (e.g. by presenting the question “yes or no ?” to the user). One can interchange output components in the same way as with input components. Due to this flexible DWARF component model, the exchange of I/O components works even at system runtime.

Efficient Coupling of Real and Virtual Worlds

One of the recurring input modalities used by all tangible objects is their location within the scene. For example, mobile displays present scene views depending on their current position. To ensure the most efficient transfer of pose data from trackers to 3D viewers, our system allows to establish direct connections between tracking and viewing components whenever no filter has to be involved.

For more complex visualizations e.g. non-linear input processing [37, 46], a pipe-filter component tree preprocesses the Pose Data emitted by the trackers before feeding it to the viewer. This approach has already been implemented by several other frameworks. In our approach, we add to this approach the concept that the filter components can be arbitrarily distributed and interchanged at runtime, once again using the flexible DWARF component model, allowing us to quickly experiment with different arrangements during system runtime. This feature turns out to be very beneficial to user interface development whenever we need to dynamically experiment with different options to interpret the tracking data [25].

In contrast, if the setup is known in advance, quite a few systems exist that are able to process and forward multi-sensory input data to a display system. An example is the OpenTracker system of the Technical University of Vienna [41]. In a joint collaboration [5], we have shown that the Open Tracker system can easily be integrated into DWARF by plugging the associated transformation hierarchy of the Open Tracker framework into the DWARF system. Similar approaches can be pursued for other systems.

Central Coordination with Petri Nets

Inside the *Interaction Management* layer, we decided to move all functionality into the DWARF User Interface Controller component (UIC), thereby combining the functionalities of *Dialog Control* and *Discrete Integration*.

It combines input tokens sent by the Media Analysis components and then triggers actions that are dispatched to components in the Media Design package. Note that the UIC must have an internal model of the state of the user interface. Otherwise context-sensitive commands could not be interpreted correctly. The internals of the UIC are explained using the example of the SHEEP game in section 3.3.

An interesting point here is that we chose Petri Nets [16] to specify the behaviour of a UIC instance, because the specification of multimodal interactions can be mapped very conveniently to Petri Nets. The UIC is based on the Petri Net framework Jfern⁶ which provides large parts of the functionality needed for this component: the Petri Net that models the multimodal interactions for a user are written in XML. From these descriptions JAVA, classes are generated. Jfern also allows the graphical display of the Petri Net and its current state. This is very useful during program development and also at demonstrations, because people can always see immediately the current state of the user interface.

User input is modeled as tokens (not to be confused with the tokens sent by the Media Analysis components) that are placed into the Petri Net. The rule-based evaluation of the user input is encapsulated into guards that check whether a transition is legal. Whenever a transition is triggered, events are sent to the Media Design components, which then adds, removes or changes the properties of parts of the UI.

Figure 3 shows the flow of events within a Petri Net (bottom row in figure 3 from the SHEEP game (see section 3.3)) and the according inputs to and outputs of the system (top row). The UIC receives two tokens from the *Media Analysis* package: one that represents the collision of wand and table and one for the “insert” speech command. These tokens are placed onto places in the Petri Net. After all places on incoming arcs of a transition are full, the transition is triggered which results in a creation of a new sheep.

Interactions and Petri Nets: One-to-One Mapping

To conform to our lightweight and distributed approach, we model each interaction in its own Petri Net. For example, all interactions found in section 3 are each realized as individual Petri Nets. When thinking about tangible interactions, this couples the functionality of each tangible object to a specific instance of a UIC. As a result, the visualization of the Petri Net’s state at runtime tells us the state of the coupled tangible interaction. Additionally this approach fits very well with the tool metaphor for interactions [4].

⁶<http://sourceforge.net/projects/jfern>

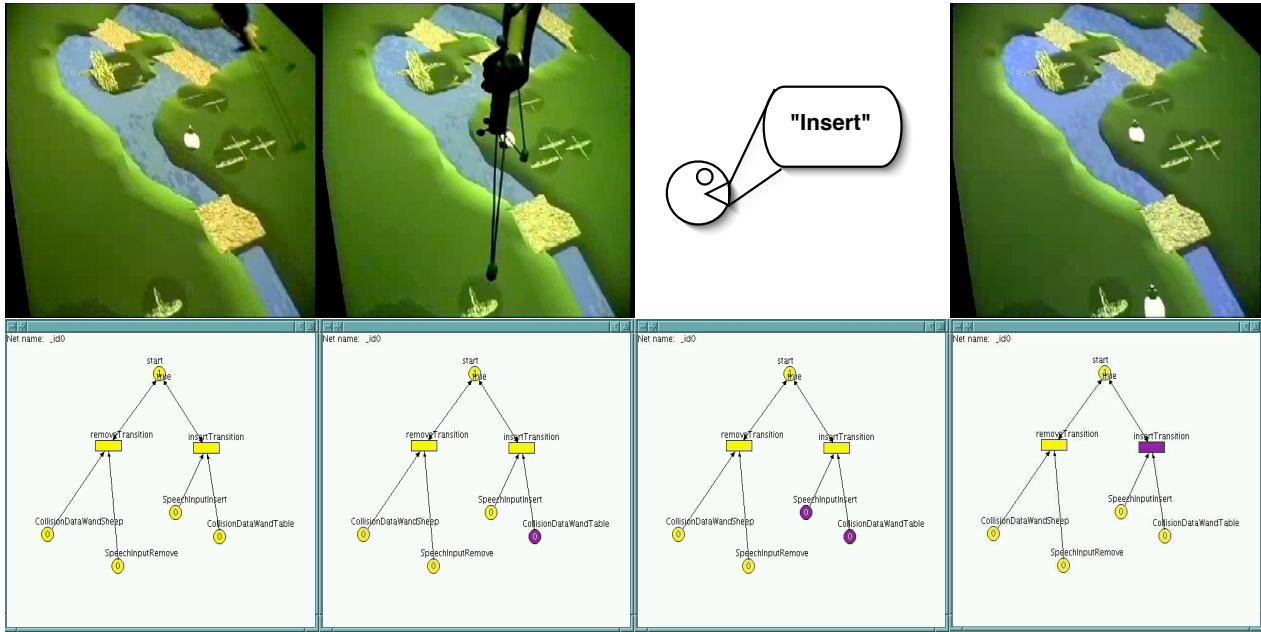


Figure 3. Realization of a multimodal point-and-speak user interface with a Petri Net within the SHEEP game

Lightweight and Stateless I/O Components

To address the flexibility requirement of user interfaces, we chose to keep as much state information as possible in the Interaction Management layer (to be explained in the next section). As a consequence, the I/O components were designed to keep as little state as possible. This allows us to add and remove I/O components conveniently at system runtime. However, this is not possible for some components. Currently, we are working on a persistence layer to be able to pass states between components.

Set of Reusable I/O Components

The available set of I/O components increases continuously. It is important to notice that reusing these components includes no programming of code, because they are generic and meant to be reused among different applications. To tailor components to a specific application, the components are configured via an XML file. Here is a short list of the most important I/O components:

Speech Recognition For commands that address the control of the system, usability studies have shown that a command language is preferable to tactile input [33]. We provide a Speech Recognition component. Especially in applications that require hands-free working (e.g. maintenance), this type of interaction gains importance. Internally, the Speech Recognition is a word spotter. It is configured via a context free grammar.

TouchGlove The TouchGlove is a special-purpose input device [6] developed at Columbia University. It is explained in some more detail in section 3.4. Input tokens that are emitted by this device can be fed into a DWARF user interface. Interestingly, this device can emit both continuous and discrete input data. It can thus be connected both to the UIC and directly to the Viewer.

Collision Detection The Pose Data emitted by the Tracking components is used by the *Collision Detection* component to detect collisions between objects. This includes collisions of real objects with virtual objects, real with real objects and virtual with virtual objects. The first two types of collisions must be considered to capture user input that is provided via movement of tangible objects. A common practice in TUIs is to couple real world objects with virtual objects. This explains the need for the last type of collisions that is

virtual with virtual objects. Consider a virtual slider that moves according to the tracked hand of the user. Whenever the slider is at its maximum value, a collision between the slider and the frame, in which it is displayed, is triggered.

Sound Player The Sound Player component is a simple JAVA application configured via an XML file. It contains the mapping between incoming commands and sound files that are to be played.

3D Viewer The 3D Viewer component displays specific views of the virtual world according to the current location and function of the tangible object. An important design goal is the ability to update the virtual parts of a 3D scene in real-time. This component turned out to be quite difficult to design and implement. Our current version accepts all important commands that are necessary for the display of dynamic 3D scenes in realtime. Additionally, several viewing modes are supported: videobackground for video see-through displays or visualization of AR scenes (e.g. figure 13b is a screenshot taken from the viewer), or support for a variety of stereo modes for different stereoscopic displays. Furthermore the display of head-fixed content (see [12]) is possible. The underlying technology is an open source implementation of Open Inventor [47]: Coin3d⁷.

2.4 Summary

A high-level description of the DWARF user interface framework is that it is a hybrid of the relatively old idea of an User Interface Management System (UIMS) [31] and a component-based toolkit for data flow networks. It is an UIMS approach, because of the application of a formal model (Petri Nets in the User Interface Controller component), clear layering and well defined tokens. It also has a toolkit character regarding the flexibly connectable filters that form data flow networks. Graphical widgets that are encapsulated in Open Inventor nodes can easily be reused and thereby are another feature of a lightweight toolkit. The combination of these two aspects fosters the advantages of both approaches. UIMS have nice properties regarding reusability and rapid prototyping. However they did not catch on [9], because of their limits regarding execution speed, difficulties to extend them for new interactions and their tendency to force programmers to use a certain specification model. The speed of interpretation of a formal model (in our case Petri Nets) is hardly an issue anymore these days. The last two concerns we hope to overcome by using a lightweight, component-based approach. Actually for rapid prototyping it is possible to run a DWARF user interface entirely without using Petri Nets and replacing them with simple Python components. After fine-tuning the parameters in the Python components, the logic can easily be ported back into the Petri Net model.

Our user interface approach has successfully been used in various systems⁸. During our experiments, we have observed that most user interface developers felt quite comfortable using Petri Nets to model their interactions. This is the first step towards developing a more general user interface based upon our approach. Once the functionality of the interaction has been specified, developers can experiment with different I/O devices to find the best setup [25].

The DWARF framework bears a lot of potential towards collaboration with other AR-based research activities and generating mutually integrated approaches. As a first step, we have worked towards integrating the STUDIERSTUBE system of the Technical University of Vienna with DWARF [5], thereby extending our space of available DWARF components. The results are very encouraging.

However, open questions remain. We have been able to achieve full flexibility for exchanging input devices at run-time. For output devices, however, we have only been partially successful. The underlying problem has been rather complex. It turned out that it is very difficult to define the semantic expressiveness of an output component, e.g. which auditory interfaces can be mapped to GUIs and which cannot? For input components the definition of expressiveness was relatively simple, because the receiver of emitted tokens is a *computer*. For output components, the receiver of content is a *human*. Perception and cognitive processing of information within the human mind are not understood well enough to come up with an ontology of output devices, yet.

⁷<http://www.coin3d.org>

⁸<http://www1.in.tum.de/DWARF/ProjectsOverview>

3 Example Systems

In this section we present several systems we have built so far. An overview of the types of user interfaces for these applications is given in figure 4. It locates various exemplary demonstration systems according to the classification of interface dimensions (multi-channel systems, mobility, integration into a 3D environment) that was suggested in the introduction.

Since all example systems are AR and TUI systems, the dimension *Interactions embedded within the three-dimensional real world* has been omitted. The abilities of the systems are shown on the two axes *Multi-channel* and *Mobile*.

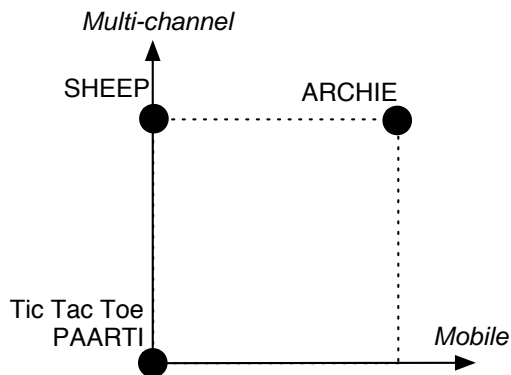


Figure 4. Overview of the types of user interfaces realized in the example systems

We present the example systems by increasing complexity and maturity. We start with the Tic-Tac-Toe system in section 3.1 which provides a tangible interface to a board game and accepts simple user gestures. Then we explain the PAARTI system (section 3.2). This is a very good example of extending a real worker’s welding tool into a tangible user interface. Adding multichannel presentation schemes, we proceed to SHEEP (section 3.3) and finally discuss the partly mobile ARCHIE system in section 3.4.

3.1 Tic Tac Toe

System Overview

An early example from our work is the Tic Tac Toe game [48, 22]. It was developed while the second author of this paper was at the Fraunhofer Institute for Computer Graphics, prior to the start of the AR research group at Technical University of Munich. It set some of the conceptual base for the later development of distributed tracking concepts [21], the DWARF framework and our approach to tangible and UAR user interaction.


The user sits in front of a Tic Tac Toe board with some chips. A camera on a tripod behind her records the scene, allowing the AR-system to track user motions while also maintaining an understanding of the current state of the game. The user and the computer alternate placing real and virtual chips on the board. The augmented scene is shown to the user on a nearby display.

Interactions

The user places a real chip onto the board (see figure 5a). When she has finished a move, she waves her hand past a 3D Go button (figure 5b) to inform the computer that she is done. It is important for the computer to wait for such an explicit "Go" command rather than taking its turn as soon as the user has moved a chip. After all, the user might not have finalized her decision. When told to continue, the computer scans the image area containing the board. If it finds a new chip, it plans its own move. It places a virtual cross on the board and writes a comment on the virtual message panel behind the game. If it could not find a new chip or if it found more than one, it asks the user to correct her placement of chips.



(a) Placement of a new stone



(b) Signalling the end of user action

Figure 5. Augmented Tic Tac Toe

Discussion

Augmented Reality Tic Tac Toe is an early and quite simple example illustrating the concepts behind our work. It does not address multichannel communication and rather focusses on the ease of use of tangible objects instead. The users generally liked the easy-to-understand and natural interactions within Tic Tac Toe. Mobile aspects have not been examined, since board games are inherently located in fixed locations.

When the Tic-Tac-Toe system was developed, the DWARF system did not exist yet. Thus, this system is a good prototype for the traditional way of developing a system. Focussed on the main event-loop of a C-program, the system used specially developed functions to use and interpret new video data, as well as event-oriented action items. Depending on the thereby resulting internal system state, the Tic-Tac-Toe game then proceeded to analyse the current viewer position, as well as the current state on the physical board game to decide upon its next move. The resulting motions of chips on the board were then transmitted to the users in an augmented, three-dimensional way.

3.2 PAARTI

System Overview

The PAARTI system [11] was a collaboration between our research group and BMW. It helps welders shoot studs with high precision in experimental vehicles. A prototype was built within six months to successfully demonstrate to BMW that it was possible to automatically guide a worker to place his welding gun with high precision at the correct locations on a car frame. The system has been tested by a number of welders. It shows significant time improvements over the traditional stud welding process. Our prototype is currently in the process of being modified and installed for productional use.

A common task in the automotive industry is to build prototype cars. A time consuming part of this process is the manual shooting of studs into a car frame. The studs have to be placed very accurately with required precisions in the millimeter range. The former process was to measure the points where the studs have to be placed with a high precision position locator (see figure 6a) and after that to shoot the stud with the welding gun (figure 6b).



(a) The old measuring device

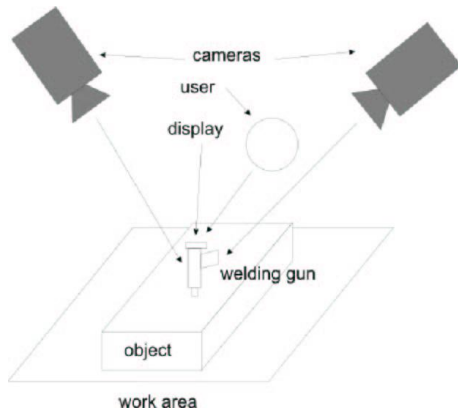


(b) The original welding gun

Figure 6. The old setup for stud welding

After examining the problem systematically (as described in more detail in [11]), we identified several different options for positioning the mobile, ar-equipped display: on the user (in form of an HMD), arbitrarily within the environment (e.g., on a near-by wall), or on the tool performing the welding operation.

Within the given context, the third solution appeared optimal with respect to the main issues, such as achieving well-specified and guaranteed levels of precision, as well as being intuitively usable. We thus decided to use the setup illustrated in 7a. It builds upon the concept of an *Intelligent Welding Gun*, a regular welding gun with a display attachment, a few buttons for user interactions, and reflective markers to track the gun position from stationary cameras (see figure 7b).



(a) Sketch of the system



(b) The intelligent welding gun

Figure 7. The new setup as realised in PAARTI

For *Tracking*, we used the commercial ART System⁹ which is a high precision (accuracy better than 1mm) optical tracking system. Several cameras are mounted around the area which contains tracked items. Near the cameras, there is an emitter for infrared light. The IR rays are reflected from the targets which allows the ART system to locate their positions. A target is composed of some retro-reflective spheres. Several targets can be tracked at the same time. We used the same tracking system for ARCHIE and SHEEP. Pictures of the targets can be seen in section 3.3, figure 9.

Interactions

While welders operate and move the gun, the display shows three-dimensional stud locations on the car frame relative to the current gun position (figure 8) Navigational metaphors, such as notch and bead and a compass, are used to help welders place the gun at the planned stud positions with the required precision. When a stud has been welded into the car frame and the welder is satisfied with the result, he presses a button on the gun as confirmation. Then he will be guided by the display on the gun to the next stud location. This process continues until the car frame is fully equipped with studs.



Figure 8. The improved welding process

Discussion

PAARTI is a stationary system, as the car frame and the intelligent welding gun have to be located in a high-precision tracking environment, which are so far only available as stationary systems. The interactions within PAARTI have been tested with several welders from BMW. The evaluation showed that the tangible interactions with the intelligent welding gun were liked very much by them. Multi-channel interactions were not necessary, because the relatively simple process of welding does not need any more information than the one delivered by the intelligent welding gun.

We think that an interesting lesson that can be learned from PAARTI is the way in which the tangible user interface was designed. A tool that was *already used* by the welders was enhanced with additional information. This pattern of augmenting already existing tools with additional information could be used in several other scenarios as well.

3.3 SHEEP

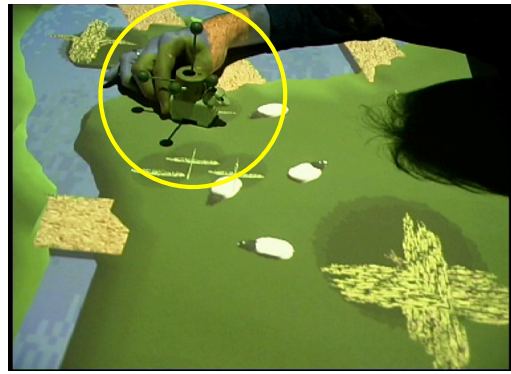
System Overview

One of the major challenges of current computer systems is to provide users with suitable means to plan, model, and control complex operations which consist of many inter-related processes and dependencies. Multimodal, multiuser interaction schemes are needed to provide adequate control metaphors. It is our hypothesis that tangible user interfaces provide particularly intuitive means for controlling complex systems.

⁹<http://www.ar-tracking.de>




a) View through the HMD while picking up a virtual sheep.



b) Attracting virtual sheep with a tangible sheep.



c) A laptop as windows to the 3D world.



d) Scooping virtual sheep with an iPAQ



e) Demonstration of SHEEP at ISMAR2002

Figure 9. Subfigures a) to e) give an overview over the various interactions realized in SHEEP . Additionally infrared-reflective markers (yellow circles) and the camera that is detecting them are highlighted (red circle in figure e).

To demonstrate the potential of tangible user interfaces to dynamically visualize, manipulate and control inter-related processes, we have built SHEEP. SHEEP is a multiplayer game centered around a physical table with a pastoral landscape that contains a herd of virtual and real sheep. The landscape and virtual sheep are projected from a ceiling-mounted projector. Players can assume one of several roles. According to their different roles, players use different input devices and interaction technologies to interact with the game.

Within this system, every sheep is an independent process, communicating with other processes to attract or repel each other. Sheep processes can be created or deleted. They can be visualized and manipulated using various modalities that are packaged into tangible units. The system was first demonstrated at ISMAR2002 [43] and formally presented at ISMAR2003 [25].

Interactions

The SHEEP game contains numerous interactions which are explained below:

Coloring sheep Figure 9a shows the view through an HMD onto the table. A player wearing an HMD can pick up sheep with his tracked hand (note the marker that is attached to the player's hand) and color them by moving them into color bars shown inside the HMD. After that, he can drop the sheep back onto the table again.

Attracting sheep The scene also allows for real, tangible sheep. Once such object, a small LEGO toy is shown in figure 9b. Since all sheep are programmed to stay together, the entire herd can be controlled by moving the tangible sheep – thereby making it to the leader of the herd. Moving the sheep around constitutes a tangible interaction which was very popular among users, because of its immediate comprehensibility.

Exploring the 3D world A separate laptop can be used to view the scene on the table in three dimensions from arbitrary vantage points (figure 9c). This constitutes a tangible interaction, with the metaphor of moving a window about in the 3D world similar to the *active lens* of the *metaDESK* [50].

Creating and removing sheep By putting on a headset with a microphone and grabbing a tracked magic wand, a player can create new sheep and remove sheep from the table. This is done by multimodal *point-and-speak input*. The technical realisation and visual sensation for the users has been shown in section 2.3 (figure 3).

Scooping sheep Players equipped with a tracked iPAQ (figure 9d) can use it to scoop sheep up from the table. Scooped sheep can be dropped back somewhere else on the table. During the scooping operation, the scooped sheep is displayed on the palm-sized computer. The entire interaction is illustrated in figure 10. This interaction is similar to Pick-and Drop [42], however we use a PDA to pick up virtual objects instead of a stylus,

Discussion

SHEEP is a stationary system which extensively uses tangible interactions. Synergistic multimodal input was used to improve usability of the system. Coordinated multimedia output was deployed to enhance the immersivity of the user's experience. Audio feedback was used whenever possible and a large number of displays was used: tracked laptop, iPAQ, projected landscape on table and a HMD.

This system was the first test of the multichannel abilities of the DWARF user interface framework. On a technical level, we were successful in validating our claims about the benefits of our framework. The 3D Viewer had to be reimplemented, as we had several problems with its implementation for SHEEP (details can be found in [25]). The improved version of the 3D Viewer was successfully deployed in ARCHIE, which is described in the next section.




Figure 10. Scooping a virtual sheep with an iPAQ that acts as tangible display

3.4 ARCHIE

System Overview

ARCHIE¹⁰ was developed in a team effort by eight master students within half a year. The goal of the project was to develop a system to support architects in constructing buildings.

Usually a number of people with different interests are involved in the development process of a building. The potential buyer has to mandate an architectural office to initiate the building process because the process is too complex to handle for himself. A mediator is responsible to represent the interest of the later building owners towards the architect. The architect assigns work to specialized persons, for example to technical engineers for designing plans of the wiring system. Although the later building owner is the contact person for the architects office, he is only one of the many stakeholders interested in the building. Furthermore landscape architects have to approve the geographic placement of the new building in the landscape. Last but not least a building company has to be involved as well. End users should be given the option to give feedback during the design phase too. So, the architects office receives feedback from many participants about their plans.

The system we developed to support the collaboration between these groups of users is described below: An architect who is equipped with a wearable computer enters the office of the future. He can now select which of the two applications to start: modelling or presentation. Modelling lets the architects work collaboratively on a virtual model that is placed on a table. They have various tangible tools to modify the model. Presentation lets the architect present the modelled building to customers.

There have already been a variety of similar systems [38, 52, 1, 7]. With ARCHIE we had a different focus compared to those systems. Instead of focussing on the overall usability by the end-users, we were more trying to test drive our framework and evaluate how it can deal with such a complex scenario.


¹⁰<http://www1.in.tum.de/DWARF/ProjectArchie>

Interactions


Mobile interactions The architect enters the office with his wearable (see figure 11a). The wearable is equipped with an HMD and a custom input device that was developed by Columbia University's Computer Graphics Lab [6]. Basically it is a touch pad fixed inside a glove, as can be seen in figure 11b. When the user enters the room, the involved components of the DWARF framework detect that they are ready to start up the modelling or the presentation application. An appropriate GUI is displayed to the user. Depending on the actual configuration of the wearable, there are two possible GUIs and dependent interactions. Either the GUI is displayed on an iPAQ (see figure 11c) and the user can select the desired application by touching the respective menu entry. Or when the architect wears an HMD, the GUI is displayed inside the HMD and the selection is made via the TouchGlove.



(a) The wearable system



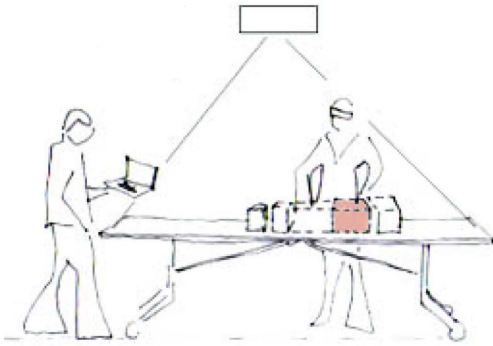
(b) The TouchGlove input device




(c) Selecting functions with the iPAQ.

Figure 11. Mobile parts of ARCHIE

Modelling The setup for this part, which is similar to SHEEP can be seen in figure 12a. An architect can take a tangible wall, move it somewhere and create a new virtual wall object by pressing the create button on his input device or using speech input. Another architect can help by moving around a tangible sun to simulate different lighting conditions, which is very important for architectural design. Both continue their work and build an approximate outer shape of a new building (see figure 12b).



(a) Conceptual drawing of the setup (courtesy of Manja Kurzak)




(b) Tangible objects for modelling and visualization: sun and wall

Figure 12. Collaborative modelling of a building


Presentation When the architect selects the presentation application, a video view is started which is projected on a wall – instead of using an HMD. Instead of the previous used HMD, now a projected view is started, providing scenes as seen from a tangible camera object, as seen in figure 13c. The audience can now get a spatial understanding of the proposed building which is displayed on the projection screen. The model to be presented is rendered stereoscopically in anaglyphic red-cyan (figure 13b) For a realistic 3D view the visitors need to wear the corresponding red-cyan glasses (figure 13a).



(a) Audience wearing red-cyan glasses



(b) View rendered in red-cyan



(c) The tangible camera controlling the viewpoint for the presentation

Figure 13. Presentation of a planned building to an audience

Discussion

ARCHIE is one of the most complex applications we have built so far. Its mobile part shows the abilities of DWARF for the combination towards wearable and ubiquitous computing. Multi-channel and tangible interactions were implemented similarly as in SHEEP and once again validated the claims regarding the expressiveness of our framework.

ARCHIE was presented to architects who liked it a lot. They suggested that an integration with already existing tools used in architecture (e.g. 3D Studio Max) would increase acceptance among architects. The modelling within ARCHIE is still very simple and will be enhanced in future versions by features like selection of different materials or group selections and manipulations of virtual objects.

Most importantly to us, ARCHIE provided new tools towards automatically setting up and testing new user interfaces for architects. While such customers were inawarely discussing topics related to the architectural design of a building at hand, we were able to also collect spontaneously provided usability data, regarding the preference of users to use various interaction tools, as well as the time they needed to use such tools to respond to system requests [24]. This intermediate level of system use, as well as system evaluation is what we are striving for. We are gathering increasing amounts of evidence that DWARF is a very suitable toolkit for this purpose.

4 Summary and Future Work

In this paper, we have postulated a holistic approach to user interface development. Based on the observation that a confluence between multimedia based, multimodal, wearable, ubiquitous, tangible and AR-based user interfaces exists, we have set out to describe the overall technical requirements for an overarching architecture. Our Distributed Wearable Augmented Reality Framework (DWARF) is able to fulfill the central requirements and has recently¹¹ been extended with interaction elements from Zoomable and Attentive User Interfaces. On its basis, we have proposed a user interface architecture for UAR. A number of prototype demonstration systems were shown to fit well into this framework.

One of the major goals of our research is to provide a rapid prototyping environment within which new user interface ideas can be prototyped and tested easily. To this end, the current setup has already proven to be suitable for joint, online development, testing and enhancement of both individual interaction facilities and multimodal, ubiquitous combinations thereof. The SHEEP game was partly developed in such joint sessions, which we called *Jam sessions* [25].

Furthermore, Kulas demonstrated within the ARCHIE system that usability evaluations can be seamlessly integrated into the live development and testing process [24]. To this end, user evaluation processes can be created to automatically inspect and evaluate the data streams flowing between the individual interaction devices, tangible objects, users, displays, etc.

As a next step, the user interface architecture (Figure 2) can serve as the basis to dynamically integrate further, in-depth enhancements to the analysis and interpretation of user input. By including tools to track a user's gestures or mimics (e.g., by eye tracking), a cognitive model of the user can be accumulated. Combined with further sources of environmental context data, the User Interface Controller can be extended to react adaptively to changing situations. By extending the information presentation primitives of the 3D Viewer, new presentation metaphors, as well as context-dependant layouts of information within an augmented environment can be provided. Note that all of these enhancements can be included, tested, modified and/or rejected online and non-exclusively.

The result is a live, dynamically changeable and also dynamically adaptive development environment for UAR user interfaces. The environment can thereby provide us with the opportunity to easily explore, combine and test different concepts that are currently emerging while also providing increasing degrees of automatic adaptation by the tools themselves.

We expect this flexible, dynamically changeable setup to be able to provide us with the tools to generate and test new concepts much more rapidly and flexibly.

Acknowledgements

Special thanks go to our students Daniel Pustka, Franz Strasser, Gerrit Hillebrand, Marco Feuerstein, Ming-Ju Lee, Otmar Hilliges, Chris Kulas, Manja Kurzak, Felix Loew, Marcus Tönnis, Johannes Wöhler and Bernhard Zaun for developing many useful components and tools, without these this work would not have been possible. The Ph.D. students who designed and developed most parts of DWARF are: Thomas Reicher, Martin Bauer, Martin Wagner and Asa MacWilliams. The Columbia University's Computer Graphics and User Interfaces Lab provided us with two prototypes of the TouchGlove; we would like to thank Steven Feiner and Gabor Blasko for their support. The

¹¹<http://www1.in.tum.de/DWARF/CarMovie>

tracking system used for SHEEP was partially on loan from BMW (TI-360). This work was partially supported by the High-Tech-Offensive of the Bayerische Staatskanzlei (Prof. Brügge).

References

- [1] D. ANDERSON, J. MARKS, A. AGARWALA, P. BEARDSLEY, D. LEIGH, E. SULLIVAN, J. YEDIDIA, J. FRANKEL, J. K. HODGINS, and K. RYALL, *Tangible Interaction and Graphical Interpretation: A New Approach to Intuitive 3D Modeling*, in Proceedings of SIGGRAPH 2000, 2000.
- [2] R. AZUMA, *A Survey of Augmented Reality*, Presence, Special Issue on Augmented Reality, 6 (1997), pp. 355–385.
- [3] M. BAUER, B. BRUEGGE, G. KLINKER, A. MACWILLIAMS, T. REICHER, S. RISS, C. SANDOR, and M. WAGNER, *Design of a Component-Based Augmented Reality Framework*, in Proc. IEEE International Symposium on Augmented Reality (ISAR'01), New York, 2001, pp. 45–53.
- [4] M. BAUER, B. BRUEGGE, G. KLINKER, A. MACWILLIAMS, T. REICHER, C. SANDOR, and M. WAGNER, *An Architecture Concept for Ubiquitous Computing Aware Wearable Computers*, in Proceedings of the 2nd International Workshop on Smart Appliances and Wearable Computing (IWSAWC 2002), Vienna, Austria, July 2002.
- [5] M. BAUER, O. HILLIGES, A. MACWILLIAMS, C. SANDOR, M. WAGNER, J. NEWMAN, G. REITMAYR, T. FAHMY, G. KLINKER, T. PINTARIC, and D. SCHMALSTIEG, *Integrating Studierstube and DWARF*, in International Workshop on Software Technology for Augmented Reality Systems (STARS 2003), 2003.
- [6] G. BLASKO and S. FEINER, *A Menu Interface for Wearable Computing*, 6th International Symposium on Wearable Computers (ISWC 2002), (2002), pp. 164–165.
- [7] W. BROLL, E. MEIER, and T. SCHARDT, *The virtual round table - a collaborative augmented multi-user environment*, in Proceedings of the third international conference on Collaborative virtual environments, ACM Press, 2000, pp. 39–45.
- [8] W. BUXTON, *The three mirrors of interaction: a holistic approach to user interfaces*, in Friend21 International Symposium on Next Generation Human Interface, Tokyo, Japan, 1991.
- [9] J. M. CAROLL, *Human-Computer Interaction in the New Millenium*, Addison-Wesley Publishing, 2002.
- [10] J. M. DAVIES, *An Ambient Computing System*, Master's thesis, Department of Electrical Engineering and Computer Science at the University of Kansas, Kansas, United States of America, 1998.
- [11] F. ECHTLER, F. STURM, K. KINDERMANN, G. KLINKER, J. STILLA, J. TRILK, and H. NAJAFI, *The Intelligent Welding Gun: Augmented Reality for Experimental Vehicle Construction*, in Virtual and Augmented Reality Applications in Manufacturing, Chapter 17, S. Ong and A. Nee, eds., Springer Verlag, 2003.
- [12] S. FEINER, B. MACINTYRE, M. HAUPT, and E. SOLOMON, *Windows on the world: 2D windows for 3D augmented reality*, in Proc. UIST'93, Atlanta, GA, November 3-5 1993, pp. 145–155.
- [13] E. GAMMA, R. HELM, R. JOHNSON, and J. VLISSIDES, *Design Patterns: Elements of Reusable Object-Oriented Software*, Addison-Wesley, Reading, MA, 1995.
- [14] D. GARLAN, D. SIEWIOREK, A. S MAILAGIC, and P. STEENKISTE, *Project Aura: Toward Distraction-Free Pervasive Computing*, IEEE Pervasive Computing, 1 (2002).
- [15] R. J. K. JACOB, L. DELIGIANNIDIS, and S. MORRISON, *A Software Model and Specification Language for Non-WIMP User Interfaces*, ACM Transactions on Computer-Human Interaction, 6 (1999), pp. 1–46.
- [16] K. JENSEN and G. ROZENBERG, *High-level Petri Nets. Theory and Applications*, Springer, 1991. ISBN 3-540-54125 X.

- [17] B. JOHANSON and A. FOX, *Extending tuplespaces for coordination in interactive workspaces*, J. Syst. Softw., 69 (2004), pp. 243–266.
- [18] B. JOHANSON, A. FOX, and T. WINOGRAD, *The Interactive Workspaces Project: Experiences with Ubiquitous Computing Rooms*, IEEE Pervasive Computing, 1 (2002), pp. 67–74.
- [19] E. KAISER, A. OLWAL, D. MCGEE, H. BENKO, A. CORRADINI, X. LI, S. FEINER, and P. COHEN, *Mutual Dissambiguation of 3D Multimodal Interaction in Augmented and Virtual Reality*, in Proceedings of The Fifth International Conference on Multimodal Interfaces (ICMI 2003)., 2003, pp. 12–19.
- [20] H. KATO, M. BILLINGHURST, I. POUPYREV, and N. TETSUTANI, *Tangible augmented reality for human computer interaction*, in Proceedings Nicograph, 2001.
- [21] G. KLINKER, T. REICHER, and B. BRUEGGE, *Distributed Tracking Concepts for Augmented Reality Applications*, in Proc. IEEE International Symposium on Augmented Reality (ISAR'00), Munich, 2000, IEEE.
- [22] G. KLINKER, D. STRICKER, and D. REINERS, *An Optically Based Direct Manipulation Interface for Human-Computer Interaction in an Augmented World*, in Proc. Eurographics Workshop on Virtual Environments (EGVE), Vienna, Austria, June 1999.
- [23] G. KORTUEM and J. SCHNEIDER, *An Application Platform for Mobile Ad-hoc Networks*, in Workshop on Application Models and Programming Tools for Ubiquitous Computing, Atlanta, USA, 2001.
- [24] C. KULAS, *Usability Engineering for Ubiquitous Computing*, Master's thesis, Technische Universität München, München, Deutschland, 2003.
- [25] A. MACWILLIAMS, C. SANDOR, M. WAGNER, M. BAUER, G. KLINKER, and B. BRUEGGE, *Herdinng SHEEP: Live Development of a Distributed Augmented Reality System*, The Second International Symposium on Mixed and Augmented Reality (ISMAR 2003), (2003).
- [26] F. MATTERN, *Pervasive/Ubiquitous Computing*, Informatik-Spektrum, 24 (2001), pp. 145–147.
- [27] M. MAYBURY and W. WHALSTER, *Readings in Intelligent User Interfaces*, 1998.
- [28] P. MILGRAM and F. KISHINO, *A taxonomy of mixed reality visual displays*, IEICE Transactions on Information Systems, E77-D (1994).
- [29] L. NIGAY and J. COUTAZ, *A Design Space for Multimodal Systems: Concurrent Processing and Data Fusion*, in Proceedings of INTERCHI '93, ACM Press: New York, 1993, pp. 172–178.
- [30] V. NOVAK, *Attentive User Interfaces for DWARF*, Master's thesis, TU München, 2004.
- [31] D. OLSEN, *User Interface Management Systems: Models and Algorithms*, Morgan Kaufmann, 1992.
- [32] A. OLWAL, *Unit - A Modular Framework for Interaction Technique Design, Development and Implementation*, Master's thesis, Royal Institute of Technology (KTH), Stockholm, Sweden, 2002.
- [33] S. OVIATT, *Ten myths of multimodal interaction*, Communications of the ACM, 42 (1999), pp. 74–81.
- [34] S. L. OVIATT, *Mutual Disambiguation of Recognition Errors in a Multimodel Architecture*, in CHI, 1999, pp. 576–583.
- [35] S. L. OVIATT, *Multimodal Interface Research: A science without borders*, Proc. of the 6th Int. Conf. on Spoken Language Processing (ICSLP), (2000).
- [36] K. PERLIN and D. FOX, *Pad: An Alternative Approach to the Computer Interface*, Computer Graphics, 27 (1993), pp. 57–72.

- [37] I. POUPYREV, M. BILLINGHURST, S. WEGHORST, and T. ICHIKAWA, *The Go-Go Interaction Technique: Non-Linear Mapping for Direct Manipulation in VR*, in ACM Symposium on User Interface Software and Technology, 1996, pp. 79–80.
- [38] M. RAUTERBERG, M. FJELD, H. KRUEGER, M. BICHEL, U. LEONHARDT, and M. MEIER, *BUILD-IT: A Planning Tool for Construction and Design*, in Proceedings of CHI 1998, 1998, pp. 177–178.
- [39] H. REGENBRECHT and M. WAGNER, *Interaction in a Collaborative Augmented Reality Environment*, in Proceedings of CHI 2002, Minneapolis, USA, 2002.
- [40] T. REICHER, A. MACWILLIAMS, B. BRUEGGE, and G. KLINKER, *Results of a study on Software Architectures for Augmented Reality Systems*, The Second International Symposium on Mixed and Augmented Reality (ISMAR 2003), (2003).
- [41] G. REITMAYR and D. SCHMALSTIEG, *OpenTracker-An Open Software Architecture for Reconfigurable Tracking Based on XML*, in VR, 2001, pp. 285–286.
- [42] J. REKIMOTO, *Pick-and-Drop: A Direct Manipulation Technique for Multiple Computer Environments*, in Proceedings of UIST, 1997, pp. 31–39.
- [43] C. SANDOR, A. MACWILLIAMS, M. WAGNER, M. BAUER, and G. KLINKER, *SHEEP: The Shared Environment Entertainment Pasture*, in Demonstration at the IEEE and ACM International Symposium on Mixed and Augmented Reality – ISMAR 2002, Darmstadt, Germany, 2002.
- [44] A. SCHMIDT, H.-W. GELLERSEN, M. BEIGL, and O. THATE, *Developing User Interfaces for Wearable Computers - Don't Stop to Point and Click*, in International Workshop on Interactive Applications of Mobile Computing (IMC2000), 2000.
- [45] B. SHNEIDERMAN, *Designing the User Interface*, Addison-Wesley Publishing, 1997.
- [46] D. SONG and M. NORMAN, *Nonlinear Interactive Motion Control Techniques for Virtual Space Navigation*, in Proceedings of IEEE Virtual Reality Annual International Symposium, 1993, pp. 111–117.
- [47] P. STRAUSS and R. CAREY, *An Object-Oriented 3D Graphics Toolkit*, Computer Graphics (SIGGRAPH 1992 Proceedings), 26(2) (1992), pp. 341–349.
- [48] D. STRICKER, G. KLINKER, and D. REINERS, *A Fast and Robust Line-based Optical Tracker for Augmented Reality Applications*, in 1. International Workshop on Augmented Reality (IWAR'98), San Francisco, November 1998, AK Peters, pp. 129–145.
- [49] M. TURK and G. ROBERTSON, *Perceptual user interfaces (introduction)*, Communications of the ACM, 43 (2000), pp. 32–34.
- [50] B. ULLMER and H. ISHII, *The metaDESK: Models and Prototypes for Tangible User Interfaces*, in Proc. UIST'97, Banff, Alberta, Canada, Oct 14-17 1997, ACM Press, pp. 223–232.
- [51] B. ULLMER and H. ISHII, *Emerging Frameworks for Tangible User Interfaces*, IBM Syst. J., 39 (2000).
- [52] J. UNDERKOFFLER and H. ISHII, *Urp: A Luminous-Tangible Workbench for Urban Planning and Design*, in Proceedings of CHI 1999, 1999, pp. 386–393.
- [53] R. VERTEGAAL, *Attentive User Interfaces*, Communications of ACM, Special Issue on Attentive User Interfaces, 46 (2003).
- [54] A. WAIBEL, M. T. VO, P. DUCHNOWSKI, and S. MANKE, *Multimodal Interfaces*, Artificial Intelligence Review, Vol.10 (1995), pp. 299–319.
- [55] M. WEISER, *Hot Topics: Ubiquitous Computing*, IEEE Computer, (1993).