A Feasible Low-Power Augmented-Reality Terminal

Johan Pouwelse* Koen Langendoen Henk Sips
Delft University of Technology
The Netherlands
{pouwelse,koen,sips}@ubicom.tudelft.nl

Abstract

This paper studies the requirements for a truly wearable augmented-reality (AR) terminal. The requirements translate into a generic hardware architecture consisting of programmable modules communicating through a central interconnect. Careful selection of low-power components shows that it is feasible to construct an AR terminal that weighs about 2 kg and roughly dissipates 26 W. With state-of-the-art batteries and a 50% average resource utilization, the terminal can operate for about 10 hours.

1. Introduction

The goal of *ubiquitous computing* is to have computers act as "human assistants" that support us instantly. Computers should move out of our awareness instead of being at the center of our attention [14]. For ubiquitous computing to become reality we need two important technologies to mature: wireless communication (wearability) and augmented reality (user-interface). Wireless communication is obviously required to obtain services provided by an arbitrary computer regardless the location of the user. State of the art in wireless communications, however, does not provide the bandwidth and/or coverage needed by most (networked) applications. For example, web surfing on a wireless computer is limited to text pages, since graphical images take forever to load. Fortunately, commercial interest in wireless communications advances the bandwidth available to individual users rather quickly.

Augmented reality (AR) is the other technology needed for ubiquitous computing, since it provides a user interface that truly hides computers from human awareness. Applications that use augmented reality are capable of enhancing the real world with additional information at any convenient location in the view of the user. This information can be as simple as plain text, and as complex as synthetic 3D objects displayed in stereo. Augmented reality obviously requires the exact location and orientation of the user to render the visual information correctly. Applications can take advantage by becoming location sensitive, for example, interactive travel guides may provide information on the spot about the spot.

The computational intensity of high-quality augmented reality is rather high and unsuitable for general-purpose wearable computers. Tasks like 3D rendering, position/orientation tracking, and video compression require special purpose components to achieve high performance and low energy consumption. Energy consumption is very important, since batteries take up a significant fraction of the weight of wireless terminals. A general-purpose wearable computer does not provide the raw processing power required by augmented-reality applications.

In the Ubiquitous Communications (UbiCom) project [3] we therefore develop a next generation of wireless augmented-reality terminals that enable ubiquitous communications. This paper proposes a terminal that meets challenging requirements in terms of augmented-reality performance, energy consumption, wearability, and flexibility. The terminal under study is based on a modular design that eases the exploitation of -down modes provided by the individual low-power components. The AR terminal is feasible and could be build today since the hardware consists of standard integrated circuits that are commercially available. Integrating all functionality, however, will be a major undertaking.

2. Requirements

The requirements of a next generation augmented-reality terminal are difficult to specify, since AR technology is rapidly advancing and true experience is limited. Nevertheless we can derive some high-level requirements from the user's perspective that a terminal should be wearable, display augmented information timely and accurately, and be flexible enough to support various applications.

---

*Supported by the Dutch organization for Applied Scientific Research (TNO), Physics and Electronics Laboratory.

2.1. AR performance

We assume that AR applications will superimpose scenes of synthetic objects on the real-world view of the user. A stereo display is required to augment the environment with true 3D objects. Objects may be composed out of detailed textures and could be dynamically requested from servers in a backbone network. This requires a wireless connection with sufficient capacity. We anticipate that some applications want to display (or capture) full motion video, so the throughput of the wireless connection should be in the order of 10 Mbps per user. The end-to-end latency over the wireless link, which is probably most stringent for interactive audio applications, is required to be less than 25 ms.

The 3D objects should, of course, be positioned accurately in the user's view. A rotational accuracy of 0.25° is a reasonable requirement [9]; it corresponds roughly to a quarter of a thumb an arm's length away. To maintain this accuracy when the user moves his head, say with a moderate 50°/s, the end-to-end latency of the AR display subsystem is required to be less than 5 ms. This is a very short time to 1) detect the user's movement, 2) compute the augmented view based on the new head position, and 3) update the display. With (perfect) head-motion prediction, however, the 5 ms latency requirement only constrains the rendering, which should therefore exceed 200 Hz.

The rotational accuracy has also consequences for the resolution of the display. The size of a pixel should be less than 0.25°, so for a field of view of 160° a minimal resolution of 640 pixels is required. In addition, the positioning system should be reasonably accurate. If we want to place virtual objects within one meter from the user, its location must be known accurately within 1 cm.

2.2. Wearability

Ideally, an AR terminal should be as wearable as clothing. At least, its weight should not hinder the users movement or balance. Head-worn components like see-through displays, trackers, and cameras should weigh less than 400 g. Components meant to be stored away in a vest pocket are limited to about the weight and size of a PDA like the palmpilot, which weighs about 170 gram. Belt-worn devices may weigh more. Properly shaped components that match the form of the human body may weigh up to 2 kg without hindering the user [4].

An important consequence of the wearability requirement is that the weight of the batteries that can be part of the AR terminal is limited to about 1 kg. With state of the art battery technology (re-chargeable Lithium cells), this means that an AR terminal can carry about 120 Wh of energy. Consequently, low-power design is a major issue when developing an AR terminal.

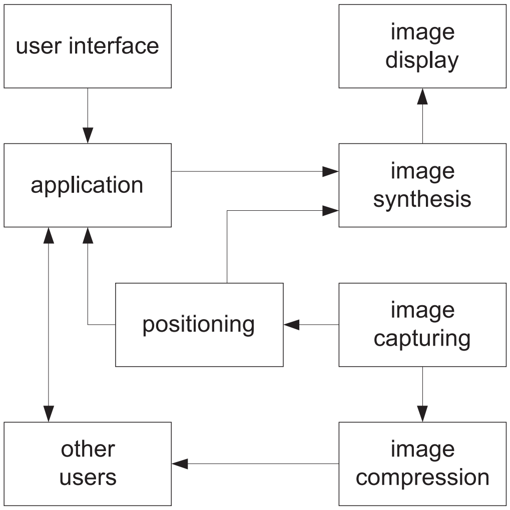

Figure 1. System structure.

2.3. Flexibility

The field of augmented reality is rapidly evolving. By requiring the mobile terminal to be flexible, the user is more likely to profit from future developments. The terminal should be designed such that it can handle advances in hardware as well as developments in software. At the hardware level, for example, the user may start off with a rather cheap and low performing head-mounted display and upgrade to an expensive alternative when becoming convinced of AR's potentials. At the software level, it is important to offer support for general purpose processing at the terminal since applications still have to be developed.

3. System structure

The requirements from the previous section are an external specification of how an AR terminal should perform. Some of the requirements concern the terminal as a whole (e.g., the total weight should be less than 2 kg), while others concern only a single unit (e.g., communication throughput of 10 Mbps). To understand the impact of global requirements it is important to identify how the units interact and depend on each other.

Figure 1 shows the functional system structure of the AR terminal envisioned in the UbiCom project. It consists of eight units. The *application* is crucial and has most connections to other units. It drives the *image synthesis* by instructing what to display, and responds to input from the *user interface* and location changes reported by the *positioning* unit. The application is also responsible for handling contact with other users reachable through the wireless network.

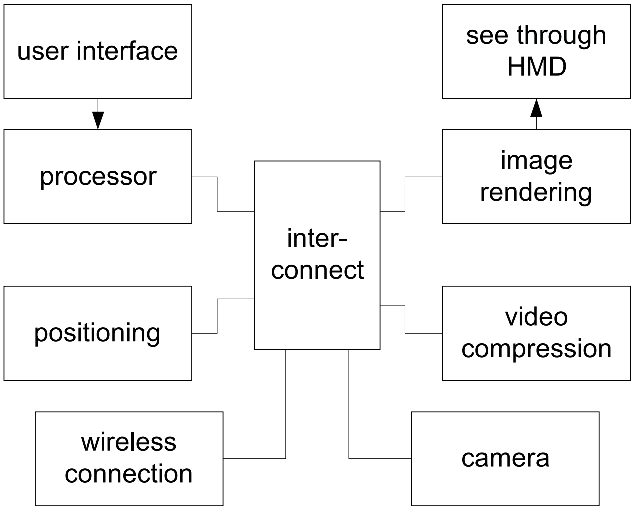

Figure 2. Hardware architecture.

Note that the image synthesis is not solely controlled by the application, but also takes information from the positioning unit. This is necessary to be able to meet the 5 ms requirement on end-to-end latency between user movement and display update. Including the application in the response loop would introduce too much latency. The application specifies which objects should appear in the augmented view, while the image synthesis determines where the objects should be displayed in the user's view based upon the exact position and orientation.

Finally, note that the *image capturing* unit provides images not only to the *image compression* unit, but also to the positioning system. This connection is to aid in meeting the 0.25° rotational accuracy requirement. Feature extraction (and tracking) in images of the real world offers an additional means to correct drift and other errors in positioning devices like accelerometers and gyroscopes [11]. Precise knowledge of the position and orientation of the user's viewpoint is vital for synthesizing images that can be accurately superimposed on the real-world view.

4. Hardware architecture

The system structure is a high-level functional description of the AR terminal and must be mapped to specific hardware such that the requirements of Section 2 are met. This is a complex task since many design issues have to be addressed and often one choice also impacts several other issues. The hardware architecture in Figure 2 constrains the seemingly unlimited design alternatives for an AR terminal. This hardware template still offers many options, while enforcing a flexible structure consisting of modules interacting through a central interconnect.

4.1. Flexibility

The modular hardware architecture provides the required flexibility for AR terminals. When an improved product becomes available, say a more accurate positioning device, it can be easily integrated by replacing the original positioning module. New products and gadgets can also be incorporated by simply adding them to the central interconnect. The general purpose processor provides the necessary flexibility for the application software.

4.2. Low power

The modular structure is also suitable for power management, which is important given the limited battery capabilities of an AR terminal. Modules can be switched on or off depending on the needs of the application. For example, when watching video clips while commuting by train the camera and positioning modules are not needed and can be powered down. Preferably a module is equipped with a local processor, so it can behave intelligently and power down its devices autonomously without explicit control from the application.

Another frequently employed method to reduce power requirements is to off load tasks to servers in the backbone network. If real-time constraints (latency and jitter) allow a task to be executed by a server, the cost and complexity of the terminal can be reduced as demonstrated by the InfoPad project [13]. The software, however, becomes more complex because it has to deal with a distributed system and must address issues like data consistency. With our hardware architecture each module is connected to the wireless network through the central interconnect and, thus, can be configured to off-load tasks to the backbone.

4.3. Multiple devices

A third advantage of the modular structure is that each module is free to use multiple devices while providing a single interface to the other modules. For example, the positioning system may combine information from accelerometers, gyroscopes, GPS, and camera images to provide the required sub-cm position accuracy. Since the sensors are part of a single module, combining the inputs occurs *locally* without interrupting other modules. Local processing is also important to avoid unnecessary data copies across the interconnect.

Multiple devices can also be used to offer a wide range of performance, which cannot be obtained with a single device. For example, a wireless connection module equipped with infrared and GSM can offer communications to the backbone network both indoor and outdoor. By selecting autonomously which device to use, the remainder of the terminal is shielded from the hand overs [12].

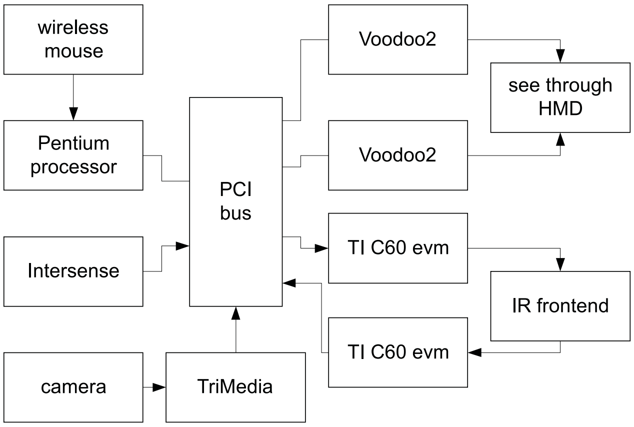

Figure 3. Demonstrator architecture.

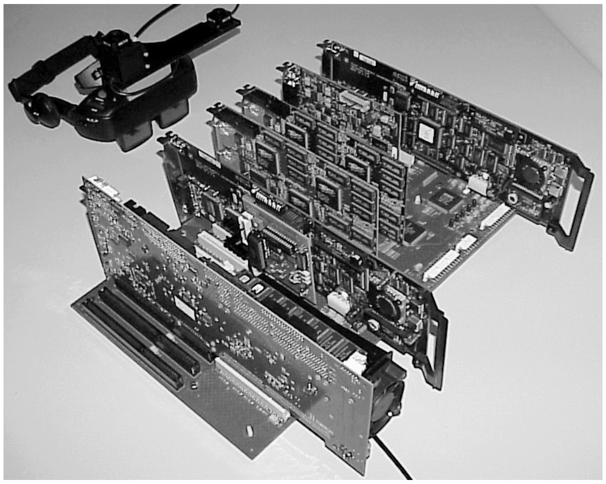

Figure 4. Demonstrator.

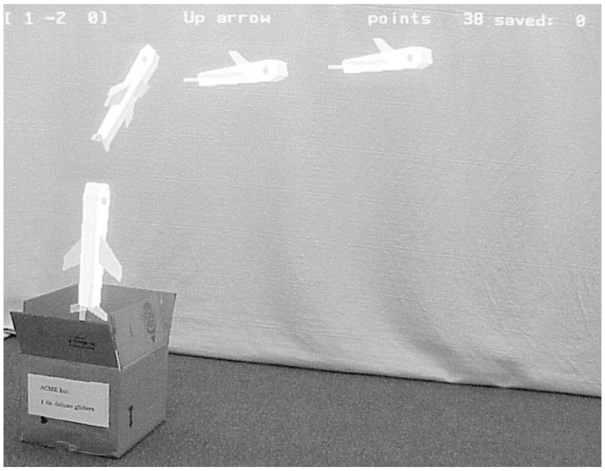

Figure 5. Application with artificial lemming gliders coming out of a real box.

5. UbiCom demonstrator

To validate the system structure and hardware architecture discussed in sections 3 and 4, a demonstrator was assembled out of off-the-shelf components. The emphasis was on developing a first terminal with all required augmented-reality functionality in a short time frame of only six months. The usage of off-the-shelf components conflicts with the wearability requirements, and resulted in a demonstrator weighing about 20 kg and dissipating 200 W. The experience gained with designing, constructing and using the demonstrator, however, provided important input for the design of a truly wearable terminal to be presented in Section 6.

The UbiCom demonstrator consists of a base station and a wireless terminal. The hardware architecture of the terminal is shown in Figure 3. The central interconnect of the terminal is a PCI bus on a passive backplane with 16 slots. PCI has the advantage that standard PC boards can be used, for example, the real-time rendering of 3D images in stereo is performed by two ordinary *Voodoo2* cards. The head mounted display (HMD) is a commercially available display from I-glasses, which provides an LCD-based *see-through stereo display* with a 42° field of view and a resolution of 640x480. It weighs almost 200 g.

The demonstrator architecture (Figure 3) closely matches the proposed hardware architecture from Section 4. The wireless connection is a custom designed infrared link driven by two DSP boards (*TI C60*): one for sending and one for receiving. The *TriMedia* DSP board performs video compression, but it takes its input directly from the camera instead of through the interconnect as proposed in Figure 2. Finally, the positioning module, a magnetic tracker by *InterSense*, uses a number of fixed ultrasonic acoustics sensors, and limits the usage of the demonstrator to an indoor environment. Figure 4 shows the demonstrator hardware without casing, power supply, and ultrasound sensor.

From the software perspective, the demonstrator is an open system based on the Linux kernel. The software is organized according to the system structure in Section 3 except for the image capturing module, which does not need to connect to the positioning system since that uses the ultrasound system for calibration. Unlike the hardware, a large fraction of the software is developed within UbiCom. The hardware products are typically accompanied by vendor software for Windows, which cannot be used on Linux and is impossible to change because of its proprietary nature. We managed to get a simple game, known as lemming gliders, to execute on the UbiCom demonstrator. Figure 5 shows an example view through the HMD with real and artificial objects.

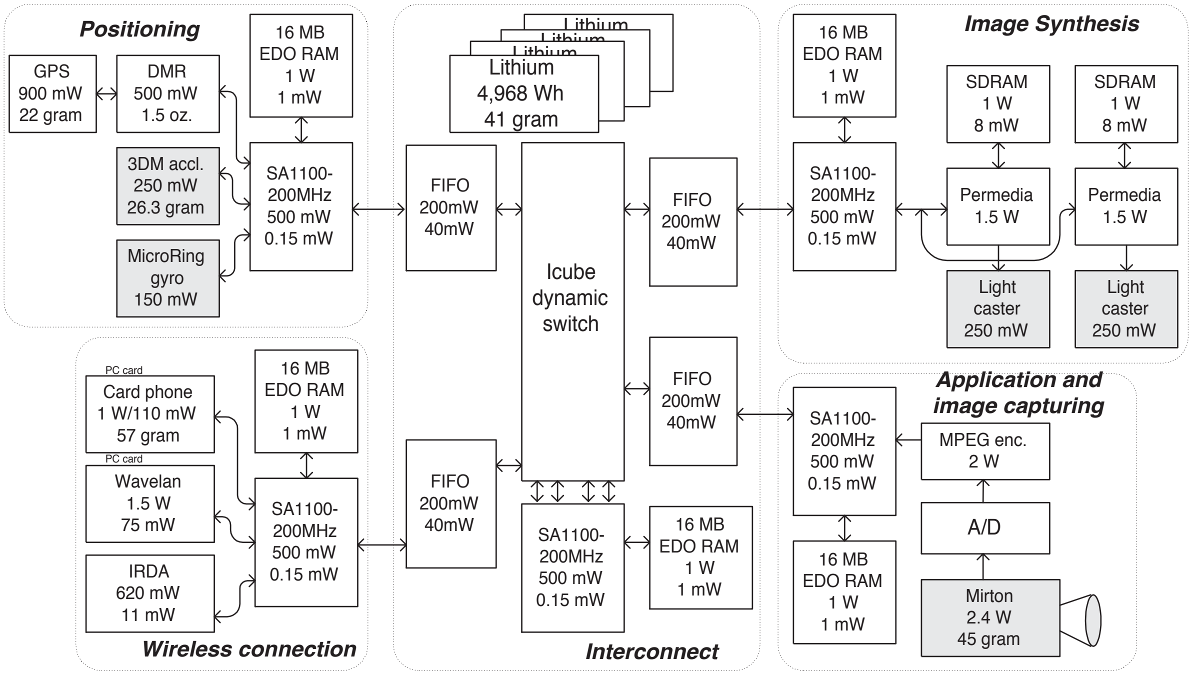

Figure 6. Low-power architecture.

6. Low-power design

This section discusses the next iteration in the design process of a wearable AR terminal and emphasizes the low-power aspect that was neglected in the UbiCom demonstrator. Our limited, but favorable, experience with intelligent modules is reflected in the low-power terminal architecture presented in Figure 6. Each module is equipped with a C-programmable StrongARM CPU (SA 1100), which controls the dedicated hardware and interfaces to the other modules. Using the same processor in all *heterogeneous programmable modules* simplifies the software considerably. Given the modest power demands of the SA 1100 (max 500 mW), our low-power architecture provides a very flexible system without compromising wearability too much. The alternative of designing dedicated hardware to reduce power consumption further may be deployed in the future when the functionality of the modules is much better identified, and thus, flexibility is not as important as today.

The remainder of this section discusses the design considerations for various modules that resulted in the selection of the components present in Figure 6. The goal is to keep overall power dissipation limited to about 20 W, so the AR terminal can operate for at least six hours with 1 kg of batteries. The operation time of the terminal depends on how effective the power-down modes of the hardware components can be deployed when executing a particular mix of AR applications. This issue will be discussed in Section 7.

6.1. Interconnect

The interconnect forms the heart of the AR terminal. It must be capable of transporting data between the various modules at the required speed. The actual throughput requirements depend on the data streams generated by the AR application at hand. We anticipate a maximum throughput with bursts up to roughly 1 Gbps, which is derived from the PCI throughput in the UbiCom demonstrator. Another requirement is that the interconnect must interface with minimal glue logic to the StrongARM processors of our intelligent modules. Two basic interconnect alternatives are available: a shared bus like Ethernet and a packet switch like ATM.

A shared bus is easily extended with new modules and, therefore, provides the flexibility required for an AR terminal. The disadvantages of bus technology are the limited data rates and the power consumption. Power-down modes are difficult to realize since all bus interfaces are required to listen continuously to determine which data is addressed to them. For example, the 400 Mbps Firewire (IEEE 1394) parts by Texas Instruments dissipate 1.4 W when transmitting/receiving at full speed.

Currently, packet switching provides the best performance/power ratio [7], since crossbar switches transport data in parallel and only active modules need to participate. We considered IEEE 1355, Octopus, and Icube as candidates for the AR terminal interconnect. The IEEE 1355 as-

model	manufacturer	resolution	color	brightness [fl]	field of view [°]	refresh [fps]	power [mW]
P5 display	reflectiontech	864x256	1 bit mono	4		60	191
cyberdisplay	kopin color	320x240	24 bit color	22	16	180	56
lightcaster	displaytech	320x240	24 bit color	20		90	100
lightcaster	displaytech	640x480	15 bit color				60
lightcaster	displaytech	1280x1024	24 bit color			60	250
microbright	planar	640x480	8 bit color	10	10	30	1300
microbright	planar	1280x1024	1 bit mono	800	160	60	300
microbright	planar	1280x1024	8 bit mono	75	160	60	1300

Table 1. Characteristics of micro displays.

sociation offers a 32-port crossbar switch with a data rate up to 3.2 Gbps [2]. This meets the throughput requirements, but it dissipates 10 W rendering it unsuitable for our AR terminal. The Octopus switch [5] is specifically designed for low-power. The current FPGA implementation only dissipates only 330 mW and provides 6 ports with an aggregate throughput of 96 Mbps.

The Icube switching fabric provides an intermediate solution between the performance of IEEE 1355 and the low power dissipation of the Octopus switch. It can transfer up to 133 Mbps per 100 mW channel. Unlike the others, it is a passive switch that needs to be driven by a control processor (another SA 1100). The FIFOs needed to buffer the data transported through the switch can easily be interfaced to the StrongARM CPU present in each module. When configured as a 1 Gbps four-port switch (including a control processor) it dissipates about 2.3 W.

6.2. Positioning

The positioning module is driven by a control program on the StrongARM CPU that combines the information generated by multiple sensors. Two sensors are head worn: the MicroRing gyroscope sensor (25 g) and the 3DM orientation sensor (26 g). These sensors track local movements and orientation of the user's head. The global position is tracked with a GPS receiver (Jupiter LP) that dissipates 250 mW excluding the antenna (100 mW). The accuracy of the GPS receiver is enhanced with a dead reckoning navigation system (DMR) using an internal magnetic compass, electronic pedometer, and barometric altimeter. The enhanced GPS information is used to correct the drift and accumulating errors of the head worn sensors.

The disadvantage of using GPS is that it only operates in an outdoor environment. Since we require the AR terminal to function indoor as well, we need an additional sensor to calibrate the head-worn orientation devices when GPS is not functioning. The InterSense system used in the UbiCom demonstrator is unsuitable since it is limited to a single room and dissipates 60 W. Within the UbiCom project we, therefore, research the possibility of using feature extraction and tracking from objects in the user's view. This requires a camera to deliver a stream of images to the positioning system and some sophisticated algorithms to analyze them. The camera images may either be derived from the general video capturing module, or from an additional camera hooked up to the positioning system directly.

6.3. Image rendering and display

The processor controlling the rendering and display system is used to convert high-level descriptions of scenes with 3D objects to render specific input formats (i.e., polygons and textures). It also decompresses audio and video encoded in the MPEG stream obtained from backbone servers over the wireless connection. The Permedia II rendering chips from 3Dlabs render in real time. They each dissipate 1.5 W and require an additional 8 MB of SDRAM (1 W).

The choice of display is quite important since it determines the look-and-feel of the AR terminal to a large extent. Head-mounted displays come in various forms and with different capabilities (see [1]). Of particular interest are the micro-displays, which offer a small form factor in combination with low power requirements. Table 1 lists the characteristics of a few commercially available micro displays. Given our requirements for a high resolution display (see Section 2) the lightcaster display is selected as the basis for the low-power AR terminal.

6.4. Application and image capturing

The main requirement on the camera is a low weight since it is to be mounted on the user's head. The MPEG2 encoder (EMPIRE by Philips) compresses the images, and an audio signal as well, so the camera should preferably output a digital signal. The MTV-5366 camera from Mitron approaches the requirements: it weighs 45 g, dissipates 2.4 W, and provides a 752x582 resolution; the analog output has to be explicitly converted to a digital signal.

The StrongARM CPU is not needed for the image capturing other than that it moves data from the MPEG2 encoder to the interconnect. It can thus also be used for other tasks. In particular, it will be used to execute applications. This saves a general purpose module at virtually no cost.

6.5. Wireless connection

The wireless communication module employs multiple devices to provide the best performance/cost ratio, where cost is defined as power dissipation. The power dissipation of a particular radio typically depends on the range of operation. We use three different radios: IrDA for short range ($\leq$ 5 m), WaveLAN (up to 500 m), and GSM for global coverage. The StrongARM CPU runs the communication software that drives the three radios, and selects the appropriate one for connecting to the backbone network depending on the application's needs (i.e. throughput) and the performance of the radios at the current location. It must provide seamless hand overs between multiple cells, and even between different radios when contact is lost, for example, by going outdoors.

IrDA is a widely accepted standard for infrared communications with low power requirements. The HP implementation, for example, dissipates 620 mW when transmitting 4 Mbps across 1.5 m. The short range limits IrDA to indoor environments such as offices and hallways. WaveLAN by Lucent, on the other hand, is a CDMA radio with a range up to 500 m outdoor. The performance depends on the distance and environment; the target of 10 Mbps limits the range to about 120 m outdoor, and 40 m indoor. It is commercially available as a PC card that dissipates 1.5 W. Finally, the GSM radio by Nokia comes as an integrated PC card with GSM phone and antenna. It delivers 9.6 Kbps, dissipates 1 W, and is intended as a backup when moving out of WaveLan coverage.

6.6. Software

The SA 1100 processor, which forms the core of each intelligent module, was primarily selected for its low-power aspects. The second advantage is that the SA 1100 is supported by the Linux operating system. Since all software for the demonstrator system runs on top of Linux, we hope that a considerable fraction can be reused for the low-power terminal. The major obstacles are the low-level device drivers and the distributed computing aspects introduced by the intelligent module concept. For example, the application and rendering software, which share the main CPU in the demonstrator, will run on separate modules in the low-power terminal.

Software is the glue that binds the intelligent modules together. As such it has a large impact on energy consumption. If organized properly, the software in a module can control the low-power modes of its hardware devices to deliver the required performance at the lowest cost. In some cases the operating system can transparently exploit low-power modes, but in other cases the application software must change its requirements or explicitly control the hardware to reduce the energy consumption to acceptable levels. For example, the Linux scheduler can be modified to adjust the processor's clock speed to automatically match the workload. Another example is the application deciding to switch to black-and-white imaging when batteries run low.

Module	power (mW)	weight (g)
positioning	**2750**	**226**
GPS receiver	250	25
GPS antenna	100	50
DMR	500	50
3 DM	250	26
gyroscope	150	25*
SA1100	500	50
16 MB EDORAM	1000	
rendering and display	**7000**	**300**
Permedia II (2)	3000	200*
8 MB SDRAM (2)	2000	
Lightcaster (2)	500	50
SA1100	500	50
16 MB EDORAM	1000	
video and application	**5900**	**145**
camera MTV-5366	2400	45
MPEG2 encoder	2000	50
SA1100	500	50
16 MB EDORAM	1000	
wireless connection	**3000**	**207**
infrared	(620)	20
WaveLAN	1500	80*
GSM	(1000)	57
SA1100	500	50
16 MB EDORAM	1000	
interconnect	**3100**	**135**
Icube	800	85*
FIFO (4)	800	
SA1100	500	50
16 MB EDORAM	1000	
Lithium battery		**1000**
AR terminal	**21750**	**2013**

* estimated value

Table 2. Power and weight breakdown of the AR terminal.

application	positioning (2.75W)	rendering (7W)	video capture (5.9W)	wireless link (3.1W)
entertainment	+	++	−	+
video communication	+/−	+	++	++
guidance	++	+	−−	+
information retrieval	+	+	−	++

Table 3. Application classification.

7. Feasibility

The integration of all components presented in the previous section will require a major engineering effort. Boards must be designed and built to host the intelligent modules, and software must be ported to or developed for them. This effort is only justified when the low-power AR terminal meets the initial requirements stated in Section 2. In particular, the terminal should be wearable. Table 2 summarizes the data previously presented and gives a breakdown in terms of weight and power dissipation. Note that the power dissipation numbers are reported for peak usage; for most components the dissipation in power-down mode is considerably lower.

The low-power AR terminal weighs about 2 kg (including batteries) and dissipates about 22 W. Since the breakdown only considers the bare components, we must roughly add another 20% for glue logic yielding a total of about 26 W. With 1 kg of batteries (120 Wh) the terminal can operate for about five hours. The power breakdown shows that the rendering module uses the most energy: 7 W, which corresponds to about one third of the total. In case of the wireless connection we have (optimistically) assumed that just one of the three radios is active at any time and, therefore, only accounted the highest contribution by WaveLAN.

7.1. Application scenarios

The operation of the proposed terminal under peak load is just five hours. In practice, however, AR applications rarely use all components and power-down modes can be exploited to extend battery life. To identify the opportunities of power-down modes we have looked at a number of application scenarios. The following scenarios are considered:

entertainment AR provides the potentials for interactive games in a real-world environment. This type of application stresses the rendering module and positioning system.

video communication With a wireless connection that can handle (compressed) real-time full motion video a whole class of applications like remote presence, wearable webcam [8], remote inspection [6] can be made available for AR terminals. These applications stress the wireless connection and video capturing modules.

guidance All kinds of route planning software and interactive travel guides (information on the spot about the spot) are promising applications for AR.

browsing Different types of information (personal notes, yellow pages, multi-media files, etc.) should be accessible from the AR terminal. Intelligence can be added at the user's convenience and results in applications like the remembrance agent [10].

Table 3 classifies the scenarios by listing how they utilize the modules in the low-power terminal. It is difficult to quantify the effective usage of all components, but applications will probably rarely use more than 50% of the terminal resources. This observation extends the operation time of the terminal to about 10 hours.

8. Conclusions

We have identified the requirements for an augmented reality terminal. Based on initial experience with a demonstrator we have developed a system structure and hardware architecture for such a terminal. The flexible architecture is built out of a central interconnect and intelligent modules for positioning, image synthesis, video compression, and wireless communication. Each programmable module is capable of making intelligent decisions about powering down devices under its control. The selection of components to use in each module was driven by the need for a low-power implementation with a reasonable weight.

We propose a low-power AR terminal based on state-of-the-art components. A power breakdown of this terminal shows that it is a feasible proposal with a weight of about 2 kg and a worst-case power dissipation of 26 W. With current battery technology this means that the terminal can operate for five hours, but considering a realistic application load and efficient exploitation of power-down modes we estimate that the terminal can operate more than twice as long.

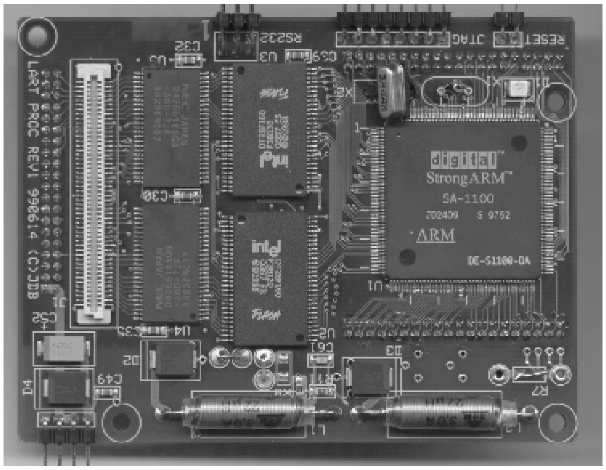

Figure 7. Basic StrongARM board.

Status and future work

The UbiCom members are determined to realize a low-power terminal based on the proposal in this paper. We plan to built upon technology under development by the related Mobile Multimedia Communication (MMC) project at Delft University. A basic mother-board with a StrongARM processor (SA 1100), 32 MB of memory, 1 MB of non-volatile memory, and a high-speed bus connector (see Figure 7) has been built and is under test now (august 1999). The board is small (75 by 100 mm) and must be equipped with extension boards to interact with specific hardware such as PCMCIA devices and the high-speed Icube interconnect. The combination of a simple processor board and device-specific extension boards has the advantage of flexibility: new hardware can be incorporated easily through designing a new extension board. The authors will concentrate on low-power aspects in all software layers of the complete AR terminal: driver-specific power-down policies, operating system support, application awareness, etc.

Acknowledgements

We thank Jan-Derk Bakker, Hylke van Dijk, Dick Epema, Wouter Pasman, Stelian Persa, Kees van Reeuwijk, Henk Rombouts, and the anonymous reviewers for their valuable feed back on the many design issues addressed in this paper.

References

[1] Head-mounted display survey . http://www.emory.edu/BUSINESS/et/vrd.

[2] IEEE 1355 standard. http://www.1355- association.org.

[3] E. Deprettere. Ubiquitous communication. Program proposal, available at http://ubicom.twi.tudelft.nl/, Delft University of Technology, Jan. 1997.

[4] F. Gemperle. Design for wearability. In *2nd ISWC*, Pittsburgh, Pennsylvania, Oct. 1998.

[5] P. Havinga and G. Smit. Octopus: embracing the energy efficiency of handheld multimedia computers. In To appear in *ACM/IEEE Mobicom*, Seattle, Washington, 1999.

[6] G. Kortuem. Some issues in the design of user-interfaces for collaborative wearable computers. *IEEE Virtual Reality Annual International Symposium*, 1998.

[7] P. Lettieri and M. Shrivastava. Advances in wireless terminals. *IEEE Personal Communications*, Feb. 1999.

[8] S. Mann. Smart clothing: The wearable computer and wearcam. *Personal Technologies*, 1(1), Mar. 1997.

[9] W. Pasman, A. van der Schaaf, R. Lagendijk, and F. Jansen. Accurate overlaying for mobile augmented reality. To appear in *Computers & Graphics*, 23(5), 1999.

[10] B. Rhodes. The wearable remembrance agent: A system for augmented memory. *Personal Technologies, Special Issue on Wearable Computing*, 1997. http://www.media.mit.edu/.

[11] A. State, G. Hirota, D. Chen, W. Garrett, and M. Livingston. Superior augmented-reality registration by integrating landmark tracking and magnetic tracking. In *SIGGRAPH 96*, pages 429–438, New Orleans, LA, Aug. 1998.

[12] M. Stemm and R. Katz. Vertical handoffs in wireless overlay networks. *ACM Mobile Networking (MONET), Special Issue on Mobile Networking in the Internet*, Summer 1998.

[13] T. Truman, T. Pering, R. Doering, and R. Brodersen. The infopad multimedia terminal: A portable device for wireless information access. *IEEE Transactions on Computers*, 47(10), Oct. 1998.

[14] M. Weiser. Some computer science issues in ubiquitous computing. *Communications of the ACM*, 36(7):74–83, July 1993.