The BlueWand as Interface for Ubiquitous and Wearable Computing Environments

Thomas Fuhrmann, Markus Klein, and Manuel Odendahl
Institut für Telematik
Universität Karlsruhe (TH), Germany

Abstract

The achievements of modern electronics more and more enmesh our daily life with various wirelessly communicating gadgets. Some of them are carried around e.g. cell-phones, personal digital assistents, or portable music players. Others are absorbed into the environment, e.g. access control systems, vehicle electronics, remotely controlled home appliances and consumer electronics. Typically, each such device is handled by a dedicated keypad. And sometimes, miniaturization has pushed the keys' size already beyond the limit of easy usage.

This paper presents the BlueWand as a control means for such ubiquitous and wearable computing environments. The BlueWand is a small pen-like device that can be used to control other Bluetooth enabled devices by hand-movements. Based on a 6-axis accelerometer and gyroscope system it is able to detect its orientation and movement in space and to transmit this data via Bluetooth to any device that can interpret these movements and execute associated commands. The BlueWand is especially suited for scenarios that require only a limited set of commands or where the controlled devices themselves are an unnatural place for the human-computer interface. Thus the BlueWand facilitates a true wireless experience.

Keywords: Ubiquitous Computing, Wearable Computing, Bluetooth, Human-Computer Interface

1 Introduction

One of the most amazing trends in both computing in general and wireless communication in particular is the perpetual miniaturization. Due to this trend the number of wirelessly networked computing devices in our environment constantly increases. At the same time, conventional interfaces like keypads, push-buttons, and turning knobs become more and more ill-suited for such tiny devices that are about to be absorbed into our homes, furniture, cars or even clothes.

Additionally, usage scenarios that involve two or more such wirelessly connected devices, can prove either of these an unnatural place for the human-computer interface. Consider e.g. a person wearing a wireless ear-phone. While walking around, this ear-phone could be connected to the cell-phone or a portable music player. From time to time, it may additionally connect to other devices that need an audio connection to that person, like the car navigation system, the intercom or a home appliance. In case a user interaction is required, none of these devices might be readily at hand because they could be stowed away in a bag or be mounted at a place the user is not yet aware of. Instead of integrating the input device into the ear-phone or forcing the user to find and reach for the appropriate correspondent, we propose to split the human-computer interface by supplementing the ear-phone with a small, lightweight and easy-to-use input device, the BlueWand.

Being pen-like in its physical form, the BlueWand follows a well established usage paradigm. It is very natural to grasp a pen to point at something, to enhance a gesture, or to use a pen to demonstrate a movement. Accordingly, the BlueWand can be used to operate the on-screen-display of a TV-set, a DVD-player, or a videocassette recorder. It can help to control the lights or blinds at home, the car-navigation system, and public info terminals or vending machines without forcing the user to touch unhygienic buttons or screens. Devices like cell-phones or portable music players that are wirelessly connected to an ear-phone can read out commands or play short jingles while the user runs the BlueWand over a virtual menu. Simple commands (e.g. accept or dismiss a call, start or stop a machine, etc.) can also be issued by simple gestures alone.

Extending this vision to a world where users not only wear wireless ear-phones but also have wireless head-up displays attached to their glasses, one can imagine an even wider range of applications for the BlueWand. Besides the already mentioned control of visual menus for video applications it could e.g. also act as pointer for an augmented reality scenario or serve as joystick for gaming on the move.

In addition to these point and control applications

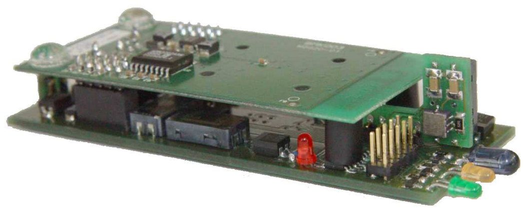

Figure 1: First prototype board of the BlueWand with orthogonal sensor fin and Ericsson Bluetooth module on top.

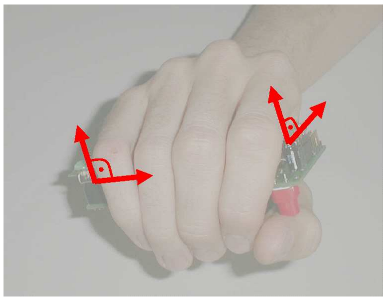

Figure 2: Bluewand held in the human hand, arrows show sensor axes

the BlueWand can also be used to store personal profiles, cryptographic keys or serve as an electronic wallet. A personal gesture like a signature instead of a PIN code can be used to express the owner's intention to open her wallet, apply her profile to a given device or use the key to open a lock. Thus the BlueWand has the potential to become a central part for a human-computer interface in the world of ubiquitous and wearable computing, replacing smart cards and PIN codes. Even more, signature gestures are for many people more easy to remember than passwords and PIN codes but hard to imitate without extensive training.

In this paper we describe the BlueWand, its principal hardware architecture (section 2), the sensor system (section 3) and the key points of the implementation of the Bluetooth stack (section 4). Section 5 gives an overview over some related work and section 6 concludes with an outlook to future work.

2 Hardware Architecture

The BlueWand architecture has to accomplish three goals: Detection of device's complete motion in space (six degrees of freedom), full Bluetooth compliance with sufficient computing resources to handle the communication with various devices that are to be controlled by the BlueWand, and a size, shape, and weight that facilitates everyday use in mass-market mobile environments. Thus the hardware design is governed by the need for a small, lightweight and cost-efficient device that has the potential to be shrunken to pen-size. Due to the projects origin as contribution to the *IEEE International Design Competition* some trade-offs had to be made for the first prototype (see figure 1): It had to house the Ericsson Bluetooth module that was provided for the competition. The second prototype currently built does not have this restriction and will thus be less than third the size shown in the photograph. Expected weight is about 50 gram (including battery).

The Bluewand is based on the Atmel AVR ATmega103 microcontroller, a high-performance low-power RISC architecture with 4 kB of internal SRAM and 128 kB in-system programmable flash. Its I/O-capabilities facilitate connection both, to the sensors and the Bluetooth module. In the BlueWand, the AVR is operated with 7.3728 MHz, providing ample resources for sensor data preprocessing and the communication protocols. If this computing power is not needed, the power saving functions of the AVR can be used to increase battery lifetime.

Rotations of the Bluewand are detected by three orthogonally mounted ENC03J piezoelectric vibrating gyroscopes from Murata. Lateral movements of the Bluewand are tracked by two ADXL202 accelerometers from Analog Devices. Compared to the gyroscopes the accelerometers provide access to absolute tilt values (i.e. no drift in measurements) and the full lateral movement. As one sensor can track two orthogonal axes, the Bluewand has to use two orthogonal sensors to track the full 3D orientation. Since the casing of the Bluewand fits into the human hand in a defined orientation only, the fourth (i.e. duplicate) axis is used to additionally detect wrist rotation independently of the gyroscopes (see figure 2). To enable the required sensor orientations one gyroscope and both accelerometers are placed together with their external resistors and capacitors on two mini-boards ("fins") mounted upright on the board.

For additional user input (e.g. selection of a virtual menu entry or activation for a gesture) the BlueWand is equipped with a button. A vibrator provides haptic feedback besides the acoustic or visual feedback via a Bluetooth ear-phone or an on-screen-display. A built-in IR LED can transmit codes for pairing the BlueWand with an accordingly equipped controlled device. The IR range can be limited such that only that device is activated that the BlueWand directly points at. Pairing is required only once, e.g. after purchase of a new BlueWand controllable device.

3 Sensor System

Sensor data is read at a rate of approximately 180 data points per second. Both, rotation and acceleration data values can be used directly, e.g. to yield an unscaled angle in the BlueWand reference frame or a relative tilt signal. Alternatively, they can be scaled to physical values using the gauge of the earth's gravitational field. Currently, this gauge has to be measured once after fabrication. A continuous adaption that combines absolute tilt values from the accelerometers with the integrated relative values from the gyroscopes can further improve the accuracy but has not fully been implemented in the prototype yet.

Preprocessing of the raw sensor data separates tilt and motion-induced acceleration by combination of rotation and acceleration data in a 6-axis motion model. Through this, the BlueWand can not only map gesture patterns to pre-recorded patterns or detect relative motions but actually reconstruct its complete 3-D movement. Drift due to inevitable sensor impreciseness can be avoided by recurrence to the gauge of the earth's gravitational field. Solely rotations within the horizontal plane (where gravity does not determine an absolute reference frame) and translatory drift has to to be accounted for by other means.

To achieve this and support the 6-axis motion model we additionally separate tilt and acceleration due to their different timescales: Tilt can be assumed to be either only slowly varying or – if changing more quickly – to change steadily to a new value without oscillation. Thus all oscillations occurring on short timescales (less than about 300 ms) can be assumed to be caused by movements, i.e., acceleration followed by retardation. This separation is achieved by a least-square fit procedure with a window of 64 sample points. The fit is performed both forwards and backwards. Afterwards both fits are combined by a weighted selection such that the resulting fit closely follows even sharp bents in the tilt curve instead of erroneously smoothing them out (cf. figure 3). The inevitable latency that is introduced by the 300 ms fit window has shown to be well tolerable in practice[1]. As an additional advantage the computing complexity of the fit algorithm does not depend on the window size, thus yielding a good performance.

4 Bluetooth Stack

A microcontroller-based Bluetooth stack has to follow design goals that completely differ from typical PC-based implementations: It does not need to support various applications through a multipurpose API, but has to spare resources as much as possible. Since none of the free Bluetooth implementations sofar available follows these goals sufficiently, we decided to build a customized stack from scratch [11]. It combines a small RAM usage (only for basic state information) with small code size (about 30 kB) and quick execution on an 8-bit microcontroller. Its key idea is to avoid buffers entirely and to generate all Bluetooth commands on-the-fly directly from the data of the preprocessing unit. Similarly, Bluetooth events are immediately processed from the byte-stream of the Bluetooth host controller interface. Thus, commands and events can be processed within a few machine cycles only, making the stack suitable for high-throughput applications too. (See [11] for details and further sample applications besides the BlueWand.)

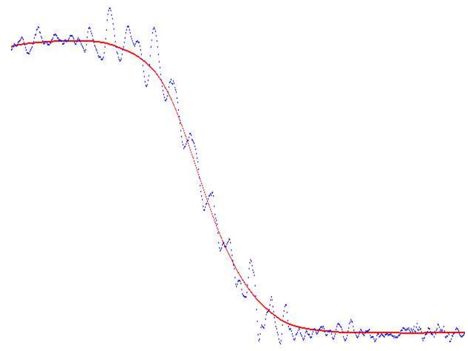

Figure 3: Single axis showing tilt-like hand-movement. Raw data shows saccadic movement that is typical for human motion. Smoothed data contains undistorted tilt.

Our Bluewand Bluetooth stack is designed to support two different application scenarios (see figure 4). In the first scenario the controlled device (e.g. a coffee machine or TV-set) contacts the Bluewand and requests the transmission of sensor or motion data or gesture commands. To this end, the controlled device makes a Bluetooth inquiry and pages devices matching the BlueWand class of device. After an authenticated connection is established, the BlueWand starts sending the requested data. Additionally, the controlled device can send configuration packets to the BlueWand in order to configure various parameters, such as sensor calibration, requested amount of preprocessing, etc.

In the second application scenario the Bluewand device periodically makes a Bluetooth inquiry, in order to find devices that offer to be controlled by the BlueWand (e.g., an MP3-player or a cell-phone maybe carried in a bag). If those devices are able to generate audio feedback, devices which can playback audio feedback are searched (e.g. a Bluetooth earphone). The BlueWand then makes an authenticated connection to the audio feedback device, sending the information required to connect to the other device (the Bluetooth device address, as well as the clock offset difference between its own internal clock and the remote clock, see [16]). The controlled device can then efficiently connect to the feedback device without having to make an inquiry. Via this connection it sends audio feedback of commands triggered by the BlueWand. The same mechanism will be used to efficiently connect to a Bluewand that is moving, e.g. in a building.

[1] This result might be surprising if one knows that the human auditive cortex is sensitive to latencies as small as 100ms. However, haptical feedback seems to be less rigorous.

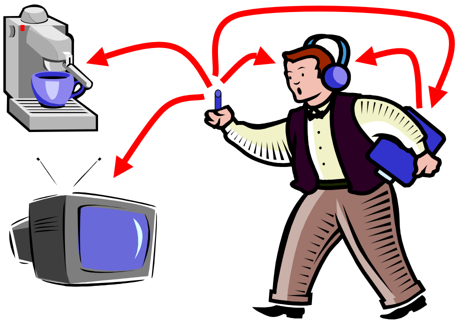

Figure 4: Typical usage scenario for the BlueWand

5 Related Work

The BlueWand pertains to both, research on human-computer interfaces for wireless and mobile environments in particular and on ubiquitous and wearable computing in general. We hence briefly discuss aspects from both areas:

3-D input is one of the most natural interfaces for human-computer interaction [7], especially for virtual environments [2] and augmented reality scenarios [18]. Most position and motion tracking techniques employ computer vision; others like the well-known data gloves rely on active sensing through mechanic, electromagnetic, acoustic or fiber optic devices [10], or also use accelerometers to track the movements of the hand and individual fingers [13]. Contrary to the Bluewand which focuses on pointing and controlling things by simple movements, these techniques can capture a much richer set of commands and gestures. But although data gloves are, in principle, very powerful, we think that permanently wearing a glove might prove too hindering in the daily practice of wearable and ubiquitous computing. Therefore people are likely to refrain from using data gloves widely. Video analysis techniques, on the other hand, require an outside view onto the person exercising the control, i.e. they can not be carried around easily by that person and are thus inappropriate for the mobile scenarios considered here.

Tangible user interfaces [19] integrate a physical data-representation with an interface to control that data or associated functionality. [1] have constructed an inertial measurement component that is very similar to the one used in BlueWand. It can be embedded into various artefacts that can then be used to control certain data representations, e.g. an avatar in a virtual environment. Gyration [6] commercially offers a gyroscope-based pointing device for beamer presentations. [15, 12] propose to use accelerometers and other gestural sensors for musical interaction and performance; and [17, 4, 5] use these techniques to create spatially aware displays for personal digital assistents. While those displays are very useful for direct manipulation of the graphical representation of a virtual environment, we believe that generally the input part of the interface should be separated from the output part, as is the case with BlueWand. Otherwise the gesture input would unnecessarily disturb the readability of a graphical output.

Concerning ubiquitous computing in general, many recent approaches still rely on improving and employing the traditional keypad scenario: [9] presents a method to increase the key density on a keypad of a mobile phone, and [3] employ cellphones and personal digital assistents for accessing ubiquitous communication walls. Recently, IBM [14] has shown a prototype of a wristwatch that contains a 2-axis accelerometer. Since a watch is worn permanently without blocking one's hands this might prove it a preferable place for a 3-D input device. However, being situated behind the wrist, it is only capable of detecting arm movements, not the much richer gestures from hand movements.

Besides the aforementioned uses of 3-D input and tangible user interfaces such motion aware devices can be of even further use in mobile environments. [20] describes a small device that combines motion patterns with map knowledge to achieve support for in-building navigation. [8] investigates the combination of gesture and speech input to resolve the problem of context awareness in speech recognition. If a short spoken command is supplemented by a gesture, it can be better separated from ambient conversation. Altogether, this could bring us another step nearer to a natural human-computer interface for ubiquitous and wearable computing scenarios.

6 Conclusions and Outlook

In this paper we have presented the BlueWand as input device for wirelessly connected ubiquitous and wearable computing environments. Based on a 6-axis sensor system and according preprocessing algorithms the BlueWand can continuously detect and report its position, orientation, and motion in 3-D space. Bluetooth enabled devices that offer to be controlled by the BlueWand do thus not require an input device of their own and can hence be fully integrated into various appliances, machines, or other artifacts. This holds even more, if the controlled device offers a natural output device (e.g. the on-screen-display of a TV-set or videocassette recorder) or e.g. a Bluetooth ear-phone can give acoustic feedback.

We believe that the BlueWand demonstrates a new human-computer interface paradigm for ubiquitous and wearable computing environments that will prove a key-element for future wireless scenarios. Our future work will thus focus on the development of sample applications for the BlueWand that demonstrate its use in such scenarios.

References

[1] Ari Y. Benbasat and Joseph A. Paradiso. An inertial measurement framework for gesture recognition and applications. In *Proccedings of the International Gesture Workshop*, pages 9–20, London, UK, April 2001.

[2] Doug A. Bowman, Ernst Kruijff, Joseph J. LaViola, and Ivan Poupyrev. An introduction to 3d user interface design. *Presence: Teleoperators and Virtual Environments*, 10(1):96–108, 2001.

[3] Alois Ferscha and Simon Vogl. Pervasive web access via public communication walls. In *Proceedings of the First International Conference, Pervasive 2002*, pages 84–97, Zürich, Switzerland, August 2002.

[4] George W. Fitzmaurice and William Buxton. An empirical evaluation of graspable user interfaces: towards specialized, space-multiplexed input. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pages 43–50, Atlanta, Georgia, USA, March 1997.

[5] Ken Hinckley, Jeffrey S. Pierce, Mike Sinclair, and Eric Horvitz. Sensing techniques for mobile interaction. In *Proceedings of the 13th Annual ACM Symposium on User Interface Software and Technology*, pages 91–100, New York, NY, 2000.

[6] Gyration Inc. The GyroMouse. http:// www.gyration.com.

[7] Robert J. K. Jacob. Computers in human-computer interaction. In Julie A. Jacko and Andrew Sears, editors, *Handbook for Human-Computer Interaction*. Lawrence Erlbaum, Hillsdale, N.J., 2002.

[8] Joseph J. LaViola. Whole-hand and speech input in virtual environments. Master's thesis, Department of Computer Science, Brown University, 1999. CS-99-15.

[9] David Levy. The fastap keypad and pervasive computing. In *Proceedings of the First International Conference, Pervasive 2002*, pages 58–68, Zürich, Switzerland, August 2002.

[10] Thomas B. Moeslund. Interacting with a virtual world through motion capture. In Lars Qvortrup, editor, *Interaction in Virtual Inhabited 3D Worlds*, chapter 11. Springer-Verlag, 2000.

[11] Manuel Odendahl and Thomas Fuhrmann. An optimized bluetooth stack for microcontrollers. to be published.

[12] Joseph A. Paradiso. The brain opera technology: New instruments and gestural sensors for musical interaction and performance. *Journal of New Music Research*, 28(2):130–149, 1999.

[13] John Kangchun Perng, Brian Fisher, Seth Hollar, and Kristofer S. J. Pister. Acceleration sensing glove (ASG). In *Third International Symposium on Wearable Computers (ISWC'99), San Francisco*, pages 178–180, 1999.

[14] IBM Research. Linux wristwatch. http:// www.research.ibm.com/ WearableComputing.

[15] Hideyuki Sawada and Shuji Hashimoto. Gesture recognition using an accelerometer sensor and its application to musical performance control. *Electronics and Communications in Japan, Part 3*, 80(5):9–17, 1997.

[16] Frank Siegemund and Michael Rohs. Rendezvous layer protocols for bluetooth-enabled smart devices. In *International Conference on Architecture of Computing Systems (ARCS'02), Karlsruhe*, pages 256–273, 2002.

[17] David Small and Hiroshi Ishii. Design of spatially aware graspable displays. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pages 367–368, Atlanta, Georgia, USA, March 1997.

[18] Thad Starner, Steve Mann, Bradley Rhodes, Jeffery Levine, Jennifer Healey, Dana Kirsch Rosalind W. Picard, and Alex Pentland. Augmented reality through wearable computing. *Presence, Special Issue on Augmented Reality*, 6(4):386–398, August 1997.

[19] Brygg Ullmer and Hiroshi Ishii. Emerging frameworks for tangible user interfaces. In John M. Caroll, editor, *Human-Computer Interaction in the New Millennium*, pages 579–601. Addison-Wesley, August 2001.

[20] Elena Vildjiounaite, Esko-Juhani Malm, Jouni Kaartinen, and Petteri Alahuhta. Location estimation indoors by means of small computing power devices, accelerometers, magnetic sensors, and map knowledge. In *Proceedings of the First International Conference, Pervasive 2002*, pages 211–224, Zürich, Switzerland, August 2002.