

Stanford Interactive Workspaces: A Framework for Physical and Graphical User Interface Prototyping

Jan Borchers, Meredith Ringel, Joshua Tyler, and Armando Fox
Computer Science Department, Stanford University

Overview

Most Smart Homes are created evolutionarily by adding more and more technologies to an existing home, rather than being developed on a single occasion by building a new home from scratch. This incremental addition of technology requires a highly flexible infrastructure to accommodate both future extensions and legacy systems without requiring extensive rewiring of hardware or extensive reconfiguration on the software level. Stanford's *iStuff* ("Interactive Stuff") provides an example of a hardware interface abstraction technique that enables quick customization and reconfiguration of Smart Home solutions. *iStuff* gains its power from its combination with the Stanford Interactive Room Operating System (iROS), which creates a flexible and robust software framework that allows custom and legacy applications to communicate with each other and with user interface devices in a dynamically configurable way.

The Stanford Interactive Room ("iRoom", see Figure 1), while not a residential environment, has many characteristics of a Smart Home: a wide array of advanced user interface technologies, abundant computation power, and infrastructure with which to coordinate the use of these resources (for more information on the iRoom or the Interactive Workspaces project, please visit <http://iwork.stanford.edu>). As a result, many aspects of the iRoom environment have strong implications for, and can be intuitively translated to, Smart Homes. In particular, the rapid and fluid development of physical user interfaces using *iStuff* and the iROS, which has been demonstrated in the iRoom, is an equally powerful concept for designing and living in Smart Homes.

Before focusing on the details of *iStuff*, we describe the software infrastructure on which it is based, and the considerations that went into designing that infrastructure.



Figure 1. The Stanford iRoom contains a wireless GyroMouse and keyboard (visible on the table), three touch-sensitive SmartBOARDS and one non-touch-sensitive tabletop display, and a custom-built OpenGL hi-res graphic mural. The room is networked using IEEE 802.11b wireless Ethernet. Except for the hi-res mural and the tabletop, all hardware is off-the-shelf, all operating systems are unmodified Windows (various flavors) or Linux, and all software we have written is user-level.

1 iROS: Application Coordination in Ubiquitous Computing Environments

1.1 Software Requirements for Rapid Integration and Evolution

The ability to continually integrate new technologies and to handle failures in a non-catastrophic manner is essential to Smart Homes and related ubiquitous computing environments. Our experience working in the Stanford iRoom enables us to identify four important requirements for a software infrastructure in a ubiquitous computing environment:

- **Heterogeneity:** The software infrastructure must accommodate a tremendous variety of devices with widely ranging capabilities. This implies that it should be lightweight and make few assumptions about client devices, so that the effort to “port” any necessary software components to new devices will be small.
- **Robustness:** The software system as a whole must be robust against transient or partial failures of particular components. Failures should not cascade, and failure or unexpected behavior of one component should not be able to infect the rest of the working system.
- **Evolvability:** The application program interface (API) provided must be sufficiently flexible to maintain forward and backward compatibility as technology evolves. For example, it should be possible to integrate a new type of pointing device that provides higher resolution

or additional features not found in older devices, without breaking compatibility with those older devices or existing applications.

- **Compatibility:** It should be easy to leverage legacy applications and technologies as building blocks. For example, Web technologies have been used for user interface (UI) prototyping, accessing remote applications, and bringing rich content to small devices; desktop productivity applications such as Microsoft PowerPoint™ contain many elements of a “rich content display server”; and so on. Furthermore, since technology in smart spaces tends to accrete over time, today’s new hardware and software will rapidly become tomorrow’s legacy hardware and software, so this problem will not go away.

Our prototype meta-operating system, iROS (Interactive Room Operating System), meets the above criteria. We call it a meta-OS since it consists entirely of user-level code running on unmodified commodity operating systems, connecting the various iRoom entities into a “system of systems”. We discuss the main principles of iROS here to give the reader an understanding of how it facilitates building new behaviors using iStuff.

1.2 iROS and Application Coordination

We will frame our discussion in the context of the Stanford iRoom, a prototype environment we constructed that we believe is representative of an important class of ubiquitous computing installations. The iRoom is intended to be a dedicated, technology-enhanced space where people come together for collaborative problem solving (meetings, design reviews, brainstorming, etc.), and applications we prototyped and deployed were driven by such scenarios.

The basis of iROS is *application coordination*. In the original formulation of Gelernter and Carriero [2], coordination languages express the interaction between autonomous processes, and computation languages express how calculations of those processes proceed. For example, procedure calls are a special case in which the caller process suspends itself pending a response from the callee. Gelernter and Carriero argue that computation and coordination are orthogonal and that there are benefits to expressing coordination in a separate general-purpose coordination language; our problem constraint of integrating existing diverse components across heterogeneous platforms leads directly to separating computation (the existing applications themselves) from coordination (how their behaviors can be linked).

In iROS, the coordination layer is called the *Event Heap* [6]. The name was chosen to reflect that its functionality could be viewed as analogous to the traditional event queue in single-computer operating systems. The Event Heap is an enhanced version of a *tuplespace*, one of the general-purpose coordination languages identified by Gelernter and Carriero. A tuple is a collection of ordered fields; a tuplespace is a “blackboard” visible to all participants in a particular scope (in our case, all software entities in the iRoom), in which any entity can post a tuple and any entity can retrieve or subscribe for notification of new tuples matching a wildcard-based matching template. We have identified important advantages of this coordination approach over using rendezvous and RMI (as Jini does) or simple client-server techniques (as has been done using HTTP, Tel/Tk [4], and other approaches); these advantages include improved robustness due to decoupling of communicating entities, rapid integration of new platforms due to the extremely lightweight client API (we support all major programming languages, including HTML, for posting and retrieving tuples), and the ability to accommodate legacy applications (simple “hooks” written in Visual Basic or Java can be used to connect existing productivity, Web, and desktop applications to the iRoom). The only criterion for making a new device or application “iRoom-aware” is its ability to post and/or subscribe to tuples in the Event Heap; since we can create Web pages that do this, any device that enters the room running a Web browser is already minimally iRoom-aware.

1.3 On-the-fly User Interface Generation in iROS

The Event Heap is the core of iROS, but we have also built other iROS services that provide higher-level functionality. Most notably, the Interface Crafter (iCrafter) framework [9] can generate UIs dynamically for virtually any iRoom entity, and on virtually any iRoom-aware device. Although it extends previous work in several important ways, including integration of service discovery with robustness and the ability to create UIs ranging from fully-custom to fully-automatic, its main role in the present scenarios is to serve as an abstraction layer between devices and UIs. Briefly, iCrafter is used as follows:

- Controllable entities “beacon” their presence by depositing self-expiring advertisements in the Event Heap. These advertisements contain a description of the service’s controllable behaviors (i.e., methods and their parameters) expressed in SDL, a simple XML-based markup language we developed.
- A device capable of displaying a UI (web browser, handheld, etc.) can make a request for the UI of a specific service, or can query the iCrafter’s *Interface Manager* (which tracks all advertisements) to request a list of available services. This initial request is made via whatever request-response technology is available on the client: visiting a well-known dynamically generated Web page is one possibility.
- The desired UI is created by feeding the SDL contained in a recent service advertisement to one or more interface generators. These may be local or remote (i.e. downloaded on demand over the Web), and may be specialized per-service and/or per-device. The Interface Manager determines the policy for selecting a generator. Part of this process includes integrating contextual information from a separate context database relevant to each workspace, making a “static” UI description portable across installations. For example, in a workspace such as ours with three large wall-mounted displays, it is preferable for a UI to refer to these as “Left, Center, Right” rather than using generic names such as “screen0, screen1, screen2” (see Figure 2).

Note that in the last step the client device and the service do not need to establish a direct connection (client-server style). This makes each robust to the failure of the other. They do not even need to be able to name each other using lower-level names such as network addresses because the tuple-matching mechanism can be based on application-level names or attributes of the service (“retrieve advertisements for all devices of type ‘LightSwitch’”). The same service can be controlled from a variety of different devices without knowing in advance what types of devices are involved, since the same SDL description can be processed into quite different UIs by different interface generators.

The ability to insulate the service and the UI from each other in these ways has been critical to the rapid prototyping of new UI behaviors. iStuff builds on this ability: It uses this indirection to enable rapid prototyping of *physical* UIs as well.

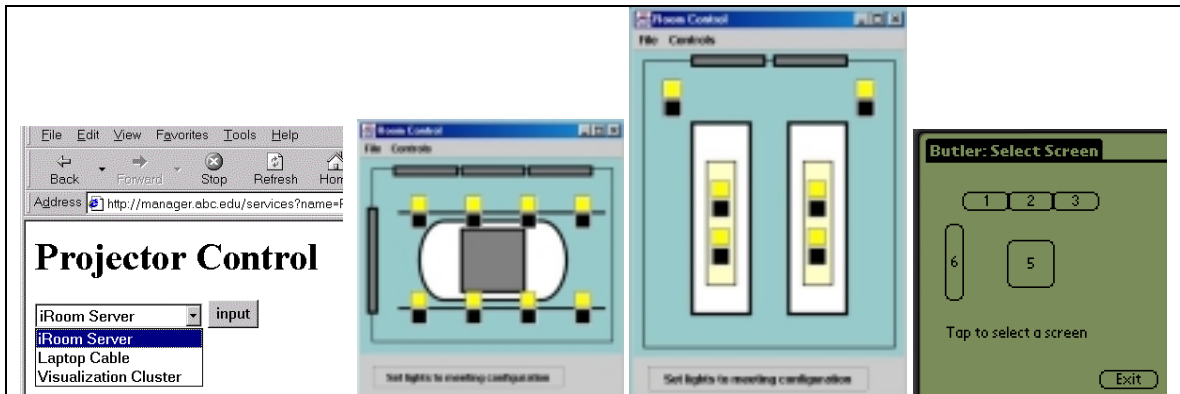


Figure 2. Screen/projector control UIs customized for various devices and incorporating installation-specific information from the context database. (a) a fragment of a projector HTML interface that can be rendered by any Web browser. This UI was generated by a projector-specific HTML generator. The symbolic names “iRoom Server”, etc. are stored in the Context Database and appear in the SDL markup as “machine0”, “machine1”, etc. (b) Room control applet for the same room, generated by a Java Swing-UI generator. The geometry information for drawing the widgets comes from the Context Database, so the generator itself is not installation-specific. Users can drag and drop web pages onto the screen widgets to cause those documents to appear on the corresponding room screens. (c) The same UI using different geometry information (for a different room) from the context database. (d) A Palm UI rendered in MoDAL [8] that lacks the drag-and-drop feature.

2 iStuff: Physical Devices for Ubiquitous Computing

2.1 iStuff Motivation and Definition

iStuff is a toolbox of wireless, platform-independent, physical user interface components designed to leverage the iROS infrastructure (which allows our custom-designed physical devices to send and receive arbitrary commands to and from other devices, machines, and applications in the iRoom). The capability to connect a physical actuator to a software service “on the fly” has appeal for the users of a ubiquitous computing environment such as a Smart Home. Residents would have the ability to flexibly set up sensors and actuators in their home, and designers of such homes would be able to prototype and test various configurations of their technology before a final installation.

There are several characteristics that are crucial for our iStuff:

- Completely autonomous packaging, wireless connection to the rest of the room, and battery-powered operation;
- Seamless integration of the devices with iROS as an existing, available, cross-platform ubiquitous computing environment to let devices, machines, and services talk to each other and pass information and control around;
- Easy configuration of mappings between devices and their application functionality, by customizing application source code, or even just updating function mappings using a web interface;
- Simple and affordable circuitry.

Various other research projects have looked at physical devices in the past; Ishii’s Tangible Bits project [5] introduced the notion of bridging the world between bits and atoms in user interfaces, and more recently, Greenberg’s Phidgets [3] represent an advanced and novel project in physical

widgets. Phidgets, however, are designed for use in isolation with a single computer, are tethered, and do not work across multiple platforms.

2.2 Device Classification and Implementation

The range of potentially useful user interface components is almost unlimited, and the really useful devices to go in a standard toolbox will only be identified over time. Ideas for such devices can be categorized according to whether they are input or output devices, and according to the amount of information they handle, as in the following examples:

- One-bit input devices, such as push buttons and toggle buttons, or binary sensors such as light gates;
- Multiple-bit discrete input devices, such as rotary switches or digital joysticks as well as packaged complex readers that deliver identification data as a result (for example, barcode readers or fingerprint scanners);
- Near-continuous input devices, such as sliders, potentiometers, analog joysticks or trackballs, and various sensors (light, heat, force, location, motion, acceleration...)
- Streaming input devices, such as microphones and small cameras;
- One-bit output devices, such as a control or status light, beepers/buzzers, solenoids and power switchers;
- Multiple-bit discrete output devices, such as LED arrays or alphanumeric LCD displays;
- Near-continuous output devices, such as servos, motors, dials, and dimmers;
- Streaming output devices, such as speakers and small screens.

Thus far, students in our lab have designed and built five types of prototype iStuff devices spanning four of the above categories: iButtons, iSliders, iBuzzers, iLEDs, and iSpeakers (Figure 3). While our hardware designs have proven surprisingly powerful and proofs-of-concept, they are simple enough to be reproduced easily (see sidebar).

Implementation Details [SIDEBAR]

Transmitting devices (buttons, sliders) contain a Ming TX-99 V3.0 300 MHz FM radio-frequency (RF) transmitter and a Holtek HT-640 encoder to send 8 bits of data to a receiver board, which contains a Ming RE-99 V3.0 RF receiver and a Holtek HT-648L decoder. The receiver board sends its data to a PC using either the parallel or USB port, and a listener program running on the PC then posts an appropriate tuple (based on the ID received) to the iRoom's Event Heap. Receiving devices (buzzers, LEDs) work in the opposite manner—a listener program receives an event intended for the iStuff and sends the target device ID through either the parallel or USB port to an RF transmitter. This data is then received wirelessly by an RF receiver in the device, resulting in the desired behavior. The iSpeaker has a different architecture, as the RF technology we employed is not sufficient for handling streaming media. Instead, a listener program on the PC waits for speaker-targeted events, and in response streams sound files over an FM transmitter, which our iSpeaker (a small, portable FM radio) then broadcasts.



Figure 3. The various types of iStuff created so far—buttons, potentiometers, speakers, and buzzers.

We have developed several successful setups using iStuff in the Stanford iRoom:

- New users coming into our iRoom are not familiar with the environment, and need an “entry point” to learn about the room and its features. Using our iStuff configuration web interface, we programmed one iButton to
 - send events that turn on all the lights in the room
 - switch on all SMARTBoards (large touch-sensitive displays) and our interactive table display
 - bring up a web-based introduction to the room on one SMARTBoard
 - show an overview of document directories for the various user groups on a second SMARTBoard, and
 - open up our brainstorming and sketching application on the third SMARTBoard.
 It is worth nothing that setting up this “iRoom Starter” took less than fifteen minutes of configuration using the web interface.
- SMARTBoards provide only a rather inconvenient way to issue right-clicks when using the touch-sensitive board for input—users have to press a right-click “mode key” on the tray in front of the board to have their next touch at the board be interpreted as a right-click. To study whether having the right-click modifier closer to the actual input location

at the board would make this interaction more fluid, we built a specialized iButton that was shaped to fit inside a hollow pen prototype made from RenShape plastic by a Product Design student in our model shop. When the button is pressed it sends an event to the Event Heap that is then received by a listener application running on the computer associated with the SMARTBoard. The listener then tells the SMARTBoard driver to interpret the next tap as a right-click. Users can now simply press the button on the pen and then tap to issue a right-click.

- We found our iSlider could conveniently control the paddles for our multi-screen version of the classic video game Pong, as described below.
- The iSpeaker has been extended to provide verbal feedback for user actions (e.g., “Pong game started”) by means of a text-to-speech program—applications simply send a “SpeakText” event to the iSpeaker containing the ASCII text to be spoken.
- We are experimenting with our iLEDs and iBuzzers to provide feedback about the status of devices in the room.

As discussed before, the Event Heap is a core component of the iRoom Operating System that makes it possible to decouple sender and receiver from the contents of the message itself (Figure 4). This architecture allows great flexibility for the prototyping of interfaces—for instance, an application can be controlled by either a traditional mouse, a graphical slider widget, or an iSlider so long as each of those devices sends the event type (perhaps an event containing a new Y-coordinate) for which the application is listening.

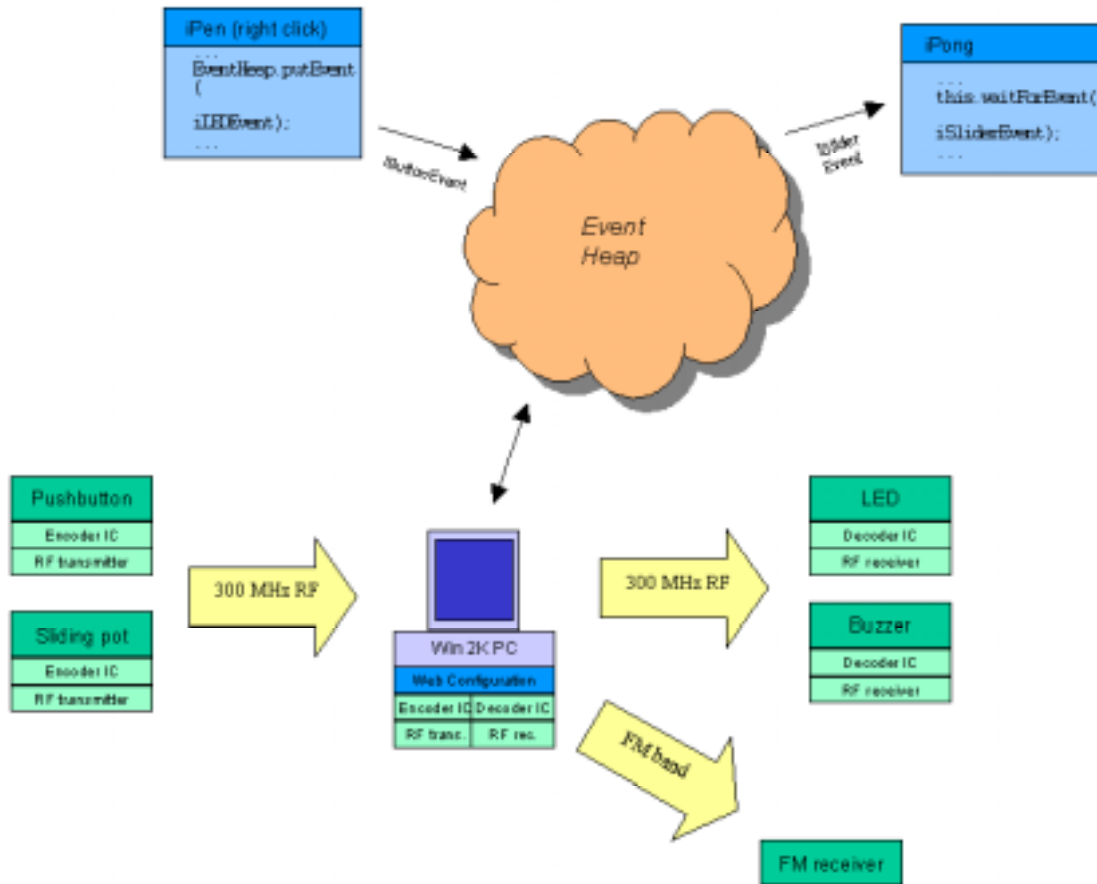


Figure 4. The overall system architecture for iStuff. In green are the actual physical devices, and in blue are a couple of examples of applications using iStuff and the iROS Event Heap. An iStuff server translates the wireless device transmissions into software events, which can be consumed by interested applications. For example, when the potentiometer of the iSlider is moved, it sends a radio signal, which is received by the server and turned into a SliderEvent. The event is posted to the Event Heap, and subsequently received by the iPong application, which is listening for SliderEvents.

In our iRoom we have demonstrated the utility of the combination of Event Heap software with iStuff hardware by developing iPong, a multi-machine version of the video game where players control the vertical position of virtual paddles to make contact with a virtual bouncing ball. The game was written to listen for Paddle Events, which contained information about the new position of the target paddle. Any input method that can generate a Paddle Event can control the paddle position. We have mapped the standard mouse, touch panel input, and an iSlider (a sliding-potentiometer iStuff widget) to drive the paddle. To the application, the physical source of the events is irrelevant. In this way, we have decoupled the link between hardware and software components in a physical user interface.

Our iButtons are already reconfigurable dynamically via a web interface that lets users enter arbitrary events to send when a specific button is pressed. We intend to provide this flexible interactive mechanism for mapping applications and events for all iStuff, using the “on-the-fly” service discovery tools of iROS (described in section 1.3). The result will be a general virtual “Patch Panel” that allows even end-users to map events to services and to map conversions between related event types. Thus, iStuff makers can send and receive their own types of events (such as button events or slider events) without concern for the exact names of events desired by

end user applications, and application developers can send and receive their own types of events (such as Paddle event) without prior knowledge of every possible type of device the user might choose to interface with their application.

The iStuff/Event Heap combination has direct applications to the Smart Home that incrementally acquires new technologies. When residents acquire a new device, or wish to reconfigure existing devices, they can simply use a utility such as our “Patch Panel” to map the event type sent by the new device to the event type expected by the target application.

2.3 Smart Home Applications

While our iStuff was originally designed with our iRoom (a space used for meetings and brainstorming/design sessions) in mind, our technology and infrastructure could be useful in a Smart Home environment. In particular, the ability to create *task-oriented user interfaces*—interfaces reflecting the user’s task as opposed to the technical features of an appliance—makes iStuff particularly compelling for Smart Home applications:

- **Dynamic, task-based remote controls:** Currently, when a user wants to watch a movie on a DVD, they need several remote controls—one to control the DVD player, another to control their home’s surround-sound system, and a third to control the television set (and then the user has to get up to dim the lights!). Today’s remotes are device-based, but because the Event Heap architecture allows for the decoupling of devices from messages we are able to use iStuff to construct task-based remote controls. By gathering appropriate iStuff components and using the Patch Panel application to ensure that the appropriate iStuff events are converted to the events appropriate for the target devices (DVD player, speakers, TV set, lights), the user can construct a task-oriented controller—one device which controls all appliances relevant to viewing a DVD movie, regardless of their physical connectivity. iCrafter could be used in an analogous manner to dynamically create GUI controllers for household appliances, thus transforming a PDA into a task-based universal remote control.
- **Monitoring house state:** A user is on her way out the door of her Smart Home, about to head off to work. The display near her door shows her the status of several devices in her home that have been instrumented with iSensors—did she leave the stove on? The lights in her bedroom? Is the thermostat too high? Is the burglar alarm on?
- **Setting house state:** A user can create an iButton or similar device to set the house’s “state” as she leaves for work everyday, and mount this button by her door. She might configure it to lower her thermostat, switch off all lights, and activate her security system, for example. This type of button is analogous to our “Start iRoom” button mentioned earlier.
- **Smart Home Design:** Architects and interior designers could use iStuff to fine-tune the placement of controls, speakers, and other interactive elements of a Smart Home. Researchers and technology developers could use iStuff to quickly prototype and test their products before putting them on the market for addition to Smart Homes.

3 Discussion and Summary

Technology advancements have made much of the original vision of ubiquitous computing feasible. A software framework, however, that integrates those heterogeneous technologies in a dynamic, robust, and legacy-aware fashion and provides a seamless user experience has been missing. We have created the Stanford iRoom, a physical prototypical space for ubiquitous computing scenarios that has been in constant use for almost two years now, to address this need. iROS, our iRoom Operating System, runs as a meta-OS to coordinate the various applications in

the room. iROS is based on a tuplespace model, which leads to the aforementioned desired characteristics. Its failure robustness has been better than average both for induced and real faults. Its ability to leverage and extend existing applications has been critical for rapid prototyping in our research.

The iStuff project builds on iROS, and tackles the problem that customizing or prototyping physical user interfaces for ubiquitous computing scenarios (such as Smart Homes) is still a very arduous process. It offers a toolbox of wireless, physical user interface components that can be combined to quickly create nonstandard user interfaces for experimentation. The iROS infrastructure has proven invaluable in making the software integration of these custom devices very straightforward. The flexibility of the technology we developed for Stanford's iRoom has potential benefits in a Smart Home scenario—for example, by enabling users to quickly create a customized, task-based interface to a system in their home.

In all, we hope that our approach to building a software and hardware framework for ubiquitous computing environments, and the various building blocks we have implemented and deployed, are general and useful enough so that others will find them of value. For more information on our Stanford Interactive Workspaces project, to access iStuff documentation, or to download iROS software, please visit our project homepage at <http://iwork.stanford.edu/>.

Acknowledgements

The authors would like to thank Maureen Stone, Michael Champlin, Hans Anderson and Jeff Raymakers for their contributions to this work, as well as the Wallenberg Foundation (<http://www.wgln.org>) for its financial support.

References

- [1] W. Keith Edwards and Rebecca E. Grinter. *At Home with Ubiquitous Computing: Seven Challenges*. In Proceedings of UbiComp 01, Atlanta, GA, pp. 256-272.
- [2] Gelernter, D., and Carriero, N., *Coordination Languages and their Significance*, Communications of the ACM, Vol. 32, Number 2, February, 1992.
- [3] Saul Greenberg and Chester Fitchett. *Phidgets: Easy development of physical interfaces through physical widgets*. In Proceedings of the UIST 2001 14th Annual ACM Symposium on User Interface Software and Technology (Orlando, FL, USA, November 11–14, 2001), pages 209–218, ACM, New York, 2001.
- [4] Todd D. Hodes, Randy H. Katz, Edouard Servan-Schreiber, and Larry Rowe. *Composable Ad-Hoc Mobile Services for Universal Interaction*. In Proceedings of Third Int'l. Symposium on Mobile Computing and Communication (ACM MobiCom 97), Budapest, Hungary, Sept. 1997.
- [5] Hiroshi Ishii and Brygg Ullmer. *Tangible bits: Towards seamless interfaces between people, bits and atoms*. In Proceedings of the CHI 97 Conference on Human Factors in Computing Systems (Atlanta, GA, USA, March 22–27, 1997), pages 234–241, ACM, New York, 1997.
- [6] Brad Johanson and Armando Fox. *The Event Heap: A Coordination Infrastructure for Interactive Workspaces*. To appear in Proceedings of the 4th IEEE Workshop on Mobile Computer Systems and Applications (WMCSA-2002), Callicoon, New York, June 2002.
- [7] Emre Kiciman and Armando Fox. *Using Dynamic Mediation to Integrate COTS Entities in a Ubiquitous Computing Environment*. In Proceedings of the Second International Symposium

on Handheld and Ubiquitous Computing 2000 (HUC2K) (Lecture Notes in Computer Science, Springer Verlag).

[8] Toby Lehman et al. MoDAL (Mobile Document Application Language). See <http://www.almaden.ibm.com/cs/TSpaces/MoDAL>

[9] Shankar R. Ponnekanti, Brian Lee, Armando Fox, Pat Hanrahan, Terry Winograd. *ICrafter: A Service Framework for Ubiquitous Computing Environments*. In Proc. UbiComp 01, Atlanta, GA.

About the authors




Jan Borchers is an Acting Assistant Professor of Computer Science at Stanford University. He works on Human-Computer Interaction in the Stanford Interactivity Lab, where he studies post-desktop user interfaces, HCI design patterns, and new interaction metaphors for music and other types of multimedia. He holds a Ph.D. in computer science from Darmstadt University, and has been known to turn his research into public interactive exhibits. He can be reached at borchers@cs.stanford.edu.



Meredith Ringel is a first-year Ph.D. student in Computer Science at Stanford, with a focus on Human-Computer Interaction. Merrie received her BS in Computer Science from Brown University. She can be reached at merrie@cs.stanford.edu.



Joshua Tyler is a second-year Masters student in Computer Science at Stanford with a specialization in Human-Computer Interaction. He received a BS in Computer Science from Washington University. He can be reached at jtyler@cs.stanford.edu.



Armando Fox joined the Stanford faculty in January 1999. His research interests include the design of robust Internet-scale software infrastructure, particularly as it relates to the support of mobile and ubiquitous computing, and user interface issues related to mobile and ubiquitous computing. Armando received a BSEE from M.I.T., an MSEE from the University of Illinois,

and a Ph.D. from UC Berkeley. He is a founder of ProxiNet, Inc. (now a division of PumaTech), which commercialized thin client mobile computing technology developed at UC Berkeley. He can be reached at fox@cs.stanford.edu.