Metaphor and Manifestation—Cross-Reality with Ubiquitous Sensor/Actuator Networks

MIT Media Lab's Responsive Environments Group is exploring ways to bridge networked electronic sensors and human perception through "cross reality" implementations that render and manifest phenomena between real world and virtual environments via densely embedded sensor and actuator networks.

Joshua Lifton, Mathew Laibowitz, Drew Harry, Nan-Wei Gong, Manas Mittal, and Joseph A. Paradiso
MIT Media Laboratory

The world is evolving its own electronic nervous system as sensor networks begin to cover the planet, and a rich set of research opportunities and challenges are generated where these cybersenses are projected onto our physical affordances. Much of this will play out where real meets virtual. Real sensed phenomena can freely manifest in virtual realms, where unconstrained by physics, users can adroitly browse and engage them. Similarly, interactions in virtual worlds can incarnate into reality via ubiquitously distributed displays and actuation. Accordingly, we can leverage virtual environments to extend our awareness and participation beyond the clutches of here and now. These environments can serve as a fluid conduit to interface our perception into the fast-evolving electronic realm of ubiquitous sensing and media, leading perhaps toward something of the "digital omniscience" envisioned by some of today's leading speculative-fiction authors.[1,2]

Researchers and practitioners have been working with intermediate blends of the real and virtual for decades (see the sidebar "Related Work with Online Virtual Worlds"). Classically subsumed under the heading of mixed reality,[3] common implementations range from installations where entire surfaces of rooms or objects are virtual (such as the partially built houses made complete with projection walls for use in military or situational training exercises) to augmented reality environments, which can be thought of as an "information prosthetic" that overlays normally invisible data onto real objects, often using mobile or head-worn devices.

We see cross-reality precipitating when diverse and ubiquitous sensor and actuator networks meet pervasively shared online virtual worlds, where phenomena freely tunnel between real and contrived continua at a multitude of "wormholes" opened by densely deployed networked devices, seamlessly adapting the level of immersion to match a variable ecology of available interfaces and user context or preference.

This article overviews several recent and ongoing projects in the MIT Media Laboratory's Responsive Environments Group that are aimed at interfacing humans with ubiquitous sensor/actuator networks. We describe several exam-

Published by the IEEE CS ■ 1536-1268/09/$25.00

Related Work with Online Virtual Worlds

Although shared graphical online virtual worlds have been around for over two decades (for example, LucasFilm's 1986 Habitat), they have evolved enormously in recent years, as computational resources, graphics performance, and high-bandwidth network penetration have skyrocketed. The most popular worlds are specific to online gaming (for example, Blizzard Entertainment's World of Warcraft), but general-purpose environments that support user-generated content have become popular and have begun to thrive with diverse applications and rich content. For example, Second Life, launched by Linden Lab in 2003, now boasts over 15 million subscribers, and hosts experiences ranging from nightclubs to a virtual Paris. With its easy-to-use graphics scripting environment, wide user base, and simple navigation interface, we've chosen Second Life as the virtual platform for our cross-reality experiments.

Some research groups, especially those specializing in environmental sensing, have explored superimposing real-time or cached sensor information over 2D maps, such as the Google Earth visualizations of the James Reserve done by the Center for Embedded Network Sensing (CENS).[1] Similarly, commercial traffic-reporting Web sites (for example, traffic.com or Google Maps) graphically append near-real-time vehicular-congestion information to maps. Diverse sensor-derived information has also been rendered atop Web-accessed map content via Microsoft's SensorMap project.[2]

There's considerable discussion in the visualization literature as to whether information is best presented in 2D or 3D—this depends on the particular data, the desired interaction modality, the user's level of experience, and how information is portrayed.[3] The advantages of 3D representation are best demonstrated when users can dynamically change their viewpoint and when information can be readily instantiated via recognizable 3D realizations that leverage the natural spatial intuitions humans have evolved in the real world. Our research focuses on browsing and interacting with real-time sensors and actuators installed in inhabited structures (typical of ubiquitous computing environments), so, Second Life's architectural bias naturally lends itself to representing such data through location-specific 3D animated constructs.

The convergence of shared 3D virtual worlds with popular Web-based data sources to form a "Second Earth" has been broadly predicted.[4] Uses of such a "hyper reality" include navigating cached and interpreted real-world data, as in the Economic Weather Map project.[5] Commercial implementations of cross-reality include IBM's visualization of data center operation[6] and VRcontext's ProcessLife technology (www.vrcontext.com), which uses high-fidelity 3D virtual replicas of real plants or factories to remotely browse and influence industrial processes in real time. The IBM project represents real data center operation through sophisticated animation but has been implemented mainly as a browser lacking virtual-to-real manifestation, while ProcessLife works in both directions. The sensor animations shown with ProcessLife, however, tend to be only simple text tags and alerts printed over a static CAD model. In contrast, the VRcontext team has blended 2D and 3D visualizations, letting users easily transition from one representation to another.

As the technology involved becomes more established, ideas relating to cross reality have been increasingly appearing in interactive art and online gaming. One of the many examples is Drew Harry's "Stiff People's League" installation at the 2008 Homo Ludens Ludens Exhibition (see http://labcast.media.mit.edu/?p=26), where virtual foosball players participating from Second Life are projected onto a real foosball table to compete with real-world players manipulating instrumented foosball rods—both look to score goals by kicking the same virtual ball.

REFERENCES

1. S. Askay, "New Visualization Tools for Environmental Sensor Networks: Using Google Earth as an Interface to Micro-Climate and Multimedia Datasets," master's thesis, Dept. Environmental Sciences, Univ. of California, Riverside, 2006.
2. S. Nath, J. Liu, and F. Zhao, "SensorMap for Wide-Area Sensor Webs," *Computer*, vol. 40, no. 7, 2007, pp. 90–93.
3. C. Ware, *Information Visualization*, Morgan Kaufmann, 2004.
4. W. Roush, "Second Earth," *Technology Rev.*, vol. 110, no. 4, 2007, pp. 38–48.
5. G. Boone, "Reality Mining: Browsing Reality with Sensor Networks," *Sensors*, vol. 21, no. 9, 2004, pp. 14–19.
6. "IBM Launches 3-D Data Center in OpenSim," *Virtual Worlds News*, 21 February 2008; www.virtualworldsnews.com/2008/02/ibm-launches-3.html.

ples of cross-reality environments that drive virtual phenomena in Second Life through real-world sensor data and tunnel virtual phenomena into the real world through distributed displays and actuators. We also introduce mobile devices for in-situ browsing, interaction, and context-scripting of sensor networks. As the level of pervasive media capture and embedded sensing that we are building into our physical environment can be perceived as being invasive, we also describe ongoing work that uses a wearable token to manage dynamic privacy.

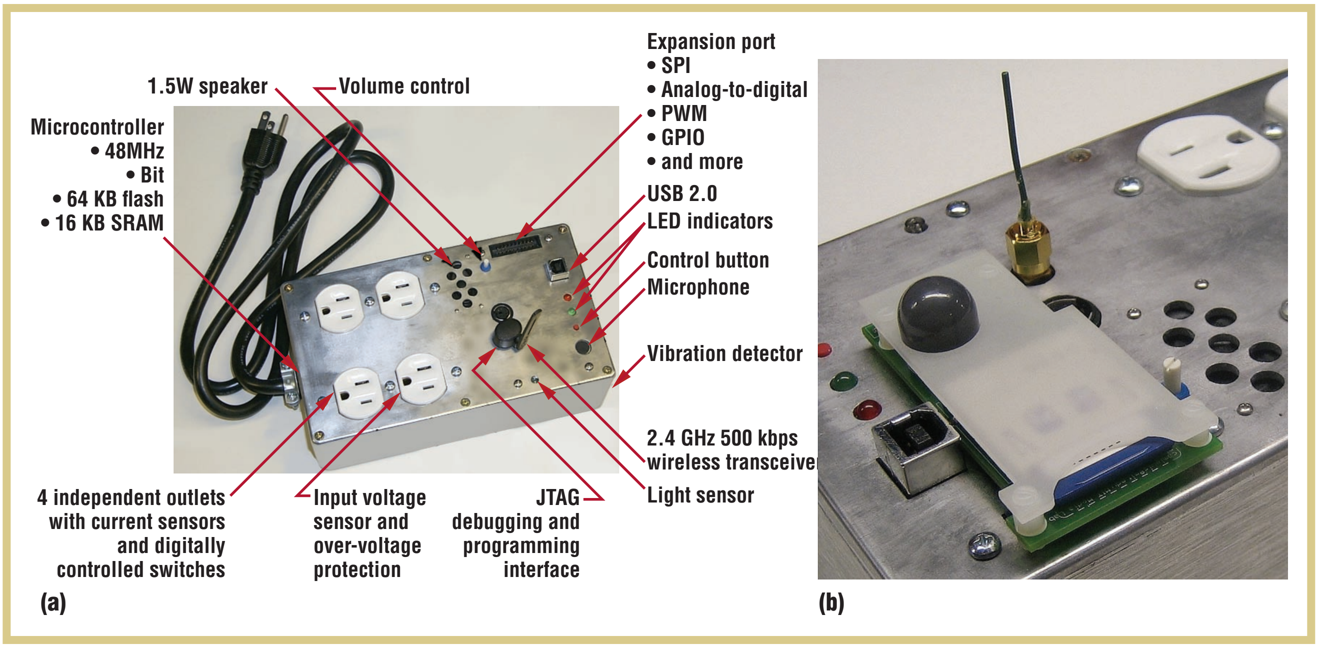

Figure 1. A sensor/actuator node as embedded in a power strip PLUG: (a) top detail and (b) a close-up showing PLUG with expansion board installed. Forty-five of these devices were deployed in our building and used in many applications, including cross-reality applications.

ShadowLab

Our initial forays into cross-reality resulted in an environment we called ShadowLab, a Second Life map of the Media Lab's third floor animated by data collected from a network of 35 smart, sensor-laden power strips (termed PLUGs) that we designed and built.[4] We chose power strips because they're already ubiquitous in offices and homes. They're ideal candidates to accrete auxiliary low-cost sensor capability because their AC connection supplies power and a potential network connection. Indeed, commercial strips are already evolving tangential functionality—for example, many include accessory surge protectors for nearby data cables. Our prototypes (see Figure 1) sense light, vibration, sound, and dynamic AC current draw from each outlet, and they host a low-power radio (based on TI/Chipcon's CC2500 transceiver chip) that can communicate with wireless and wearable devices that we've also constructed and network with other PLUGs. An expansion board on each device also supports a motion sensor, temperature sensor, and removable Secure Digital (SD) memory card for local data logging. The PLUGs can also actuate through an in-module speaker and dimming the voltage on each outlet.

Figure 2 shows two versions of ShadowLab visualizing real-time data from the PLUG network deployed throughout our floor. We've chosen simple 3D data representations that intuitively and qualitatively suggest the sensed phenomena. Instead of rendering the data literally (for example, as graphs, dials, and so on), which might not be visually informative when large numbers of data sources are shown, we animated virtual phenomena in ways that naturally suggest the sensor stimuli. For example, in the earlier rendering[4] shown in Figure 2a,each fire column's height corresponds to the amount of electrical current being pulled from a PLUG node at the corresponding map location. The twisting purple ribbon denotes significant motion sensor activity, and the height of translucent "walls" signifies an activity metric derived from nearby sound, motion, and vibration levels.

Figure 2b shows a more evolved version based on metaphors that we termed *DataPonds* animated at points on the map where real PLUG nodes are located.[5] Here, the fronds' height corresponds to the amount of light detected (color indicates their temperature—darker is warmer), the amount they wave in a synthetic breeze corresponds to the amount of locally detected motion, the radius of the blue "ripples" at bottom indicates the local sound level, the core pyramid jitters with detected vibration (meaning that somebody is walking nearby or working on the surface the PLUG is sitting on), and the quantity of orange "smoke" emanating from a node is proportional to the net amount of electrical power that's being drawn from the outlets. Rendering all sensor data from each PLUG into a single corresponding object (the DataPond) generated a visualization that was easier to meaningfully browse than its predecessor, which more freely distributed the animated constructs.

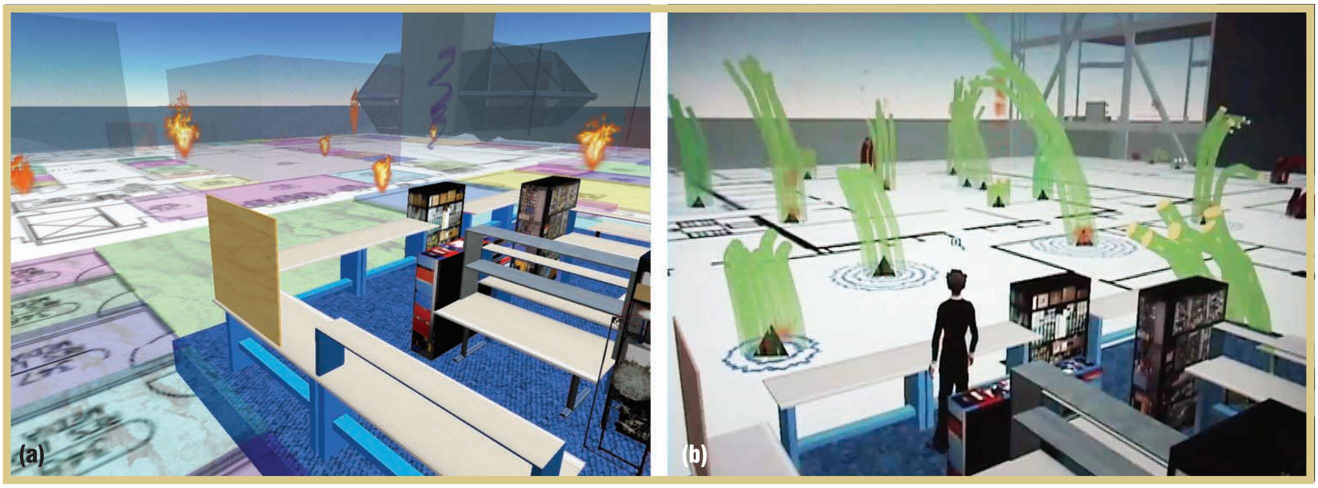

Figure 2. Two versions of ShadowLab—browsing real-world PLUG sensor data in Second Life — (a) using a dispersed visualization and (b) concentrating the visualization in discrete multimodal DataPonds. In both examples, the Responsive Environments Group's lab space is rendered in more detail. We used a simple map for the remainder of the floor.

Because we implemented ShadowLab entirely in Second Life, visiting avatars can freely browse the real environment as they walk, float, or fly about, unconstrained by physical boundaries while they note regions of high activity (from visualized sound and motion), devices that use copious electric power (from smoke and flames), and so on. Avatars can come one-by-one or, as Second Life is a shared online environment, in groups. We can also freely invite users to visit our land via a simple Web address (termed SLurl). We began experimenting with scalable interaction and zooming in these environments (for example, when an avatar approaches or touches a DataPond, audio would stream into Second Life from the corresponding PLUG). We further explored this direction with our subsequent system, however, which we describe in the next section. We also implemented early versions of virtual widgets that communicated with the real world through ShadowLab—for example, objects that would play audio clips through PLUG speakers when explicitly triggered by avatars, as well as physical versions of DataPond fronds (plastic windsocks mounted atop fans driven by PLUG outlets) that would be stimulated by avatar motion in a particular region of ShadowLab (see Figure 3).

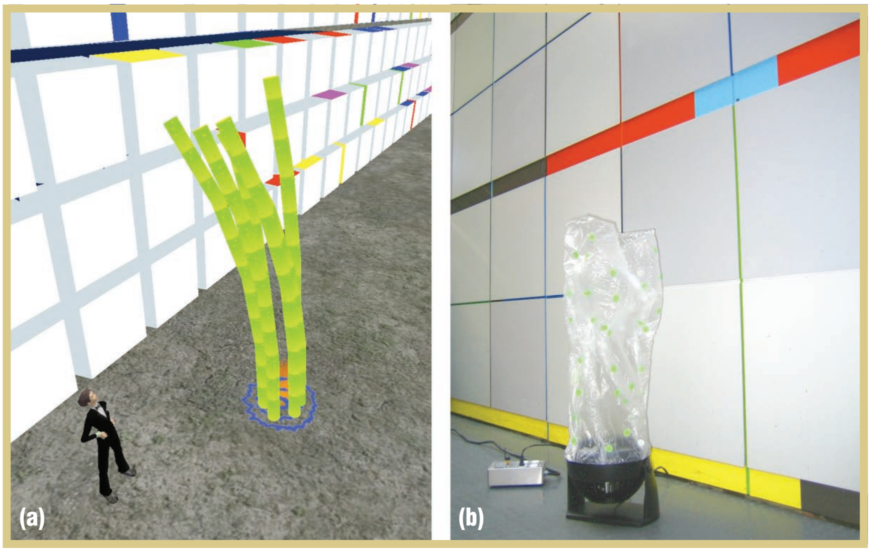

Figure 3. DataPonds: (a) A virtual DataPond in the Virtual Atrium and (b) a real DataPond in the real atrium.

Sensor-driven animation needn't be constrained to inanimate virtual objects—avatars themselves can change their appearance with real-world sensor data. The only unintentional body language exhibited in Second Life is the typing gesture avatars make when the user is typing a chat message, the slumped-over sleeping stance assumed when the user's mouse and keyboard have been inactive for a preset amount of time, automatically turning to look at nearby avatars who have just spoken, and a series of stances randomly triggered when the avatar isn't moving, such as hands on hips and a bored slouch. The user must intentionally choose all other body language and avatar actions. Accordingly, Figure 4 shows an example of what we term

Figure 4. Avatar metamorphosis. The avatar's hair is tied to PLUG sensor activity measured near the avatar's user in the real world. Changes in the avatars' appearance can be made to reflect sensor-derived real-world features, such as aspects of the user's local environment or inferred context.

metamorphosis, where the avatar begins as a typical human and transforms into a Lovecraftian alien as activity levels build outside the corresponding user's office. Although this particular example is outlandish and grotesque, in practice the mapping used in a metamorphosis is arbitrary, which is exactly its appeal as a method of self-expression. Metamorphosis can be mapped to other arbitrary stimuli and unfold in any fashion.

Today's virtual worlds are plagued by the "vacancy problem"—although millions of users have accounts in environments such as Second Life, only a small percentage of users are online at a given time. Because avatars are present only when a user is logged in, this results in a very low average instantaneous population density, producing vast landscapes void of inhabitants. Taking inspiration from the previous example, however, we can envision people continuously straddling the boundary between real and virtual through what we term "scalable virtuality," where they're never truly offline because sensor networks and mobile devices maintain a continuous background interworld connection (Mirco Musolesi and his colleagues gave an early exploration of this idea).[6] This can be tenuous, with a user's avatar, objects on their land, or even a "talisman" that a user can give to another user ambiently and metaphorically animating aspects of the user's real-world location and activity when he or she is offline. This allows a persistent virtual presence, enabling virtual visitors familiar with these mappings (for example, connected to the user's social group) to know something about what the user is doing in the real world. In principle, a touch from an authorized avatar in virtual space can reach out to the corresponding dormant user in the real world (for example, manifesting an alert or informational cue on the user's mobile device or nearby ambient display), perhaps asking for live interaction or to bring the user more fully into the virtual sphere.

Ubiquitous Sensor Portals

The physical platform on which our current cross-reality research is based is the Ubiquitous Sensor Portal (see Figure 5). The portals are I/O devices designed for rich, two-way cross-reality experiences. In addition to their myriad sensors, which we describe later, they provide a small touch-screen display and audio speaker. Information doesn't just stream away from the user's environment—the portals can also manifest virtual and remote phenomena into the user's physical space. The portals host a variety of environmental sensors that measure PIR (passive infrared) motion, light and sound level, vibration, temperature, and humidity. They also feature active IR links that can communicate with various families of badges the Media Lab has developed[7,8] to identify badged individuals facing the portal (these can also serve as reflection proximity sensors to detect the anonymous presence of an unbadged participant). The portals also act as base stations for an 802.15.4 network, enabling wireless communication with (and, via TI's CC2480 protocol, coarse localization of) a variety of wearable sensors that we've recently developed.[8] The portals also capture stereo audio and feature a camera that can snap 3.1 megapixel still images and stream DVD-quality video. An onboard TI DaVinci processor runs Linux in its ARM core, and its onchip DSP provides agility in processing the camera data, performing vision computations, and compressing or decoding video. The portals accommodate slotted flash memory, allowing a large local media and data store in addition to enabling fast streaming over the wired network. The devices also have motorized pan, tilt, and autofocus for automatic shot composition and virtual control of real-world camera gaze.

We've built 45 portals and have installed them throughout the Media Lab as a facility with which to explore applications in pervasive media, ubiquitous presence, and dense, ambient cross-reality integration. Each portal has an extension into Second Life (see Figure 6). This lets people visit our laboratory in virtual space, seeing into and appearing through any portal and float-

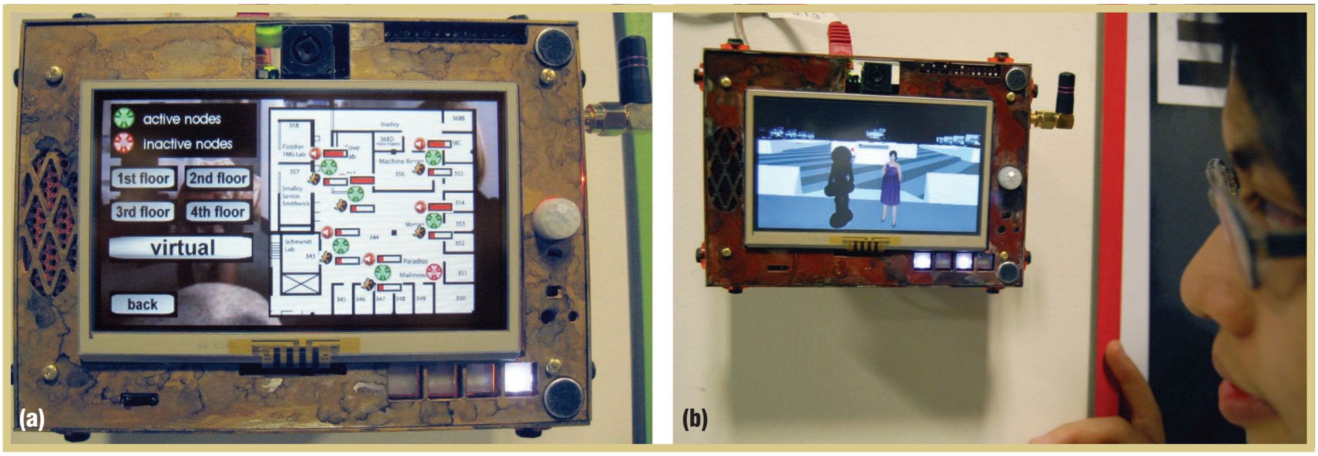

Figure 5. The Ubiquitous Sensor Portal. (a) A portal running an interactive application that lets users browse sensor and video data from other portals. (b) A portal connecting Second Life avatars looking into its virtual end with real-world people.

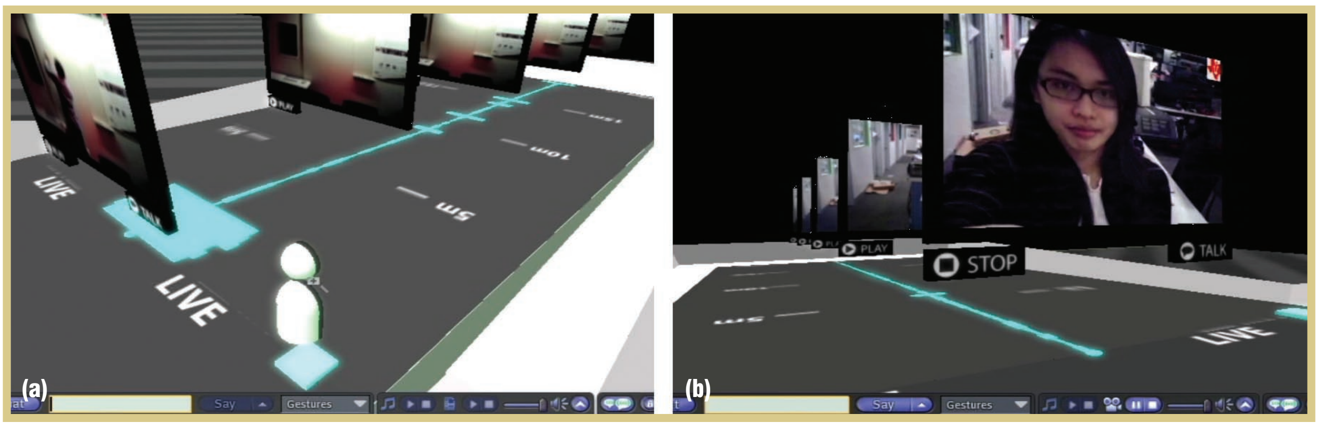

Figure 6. The virtual extension of a Ubiquitous Sensor Portal into Second Life that portrays its sensor data over time, showing current and past images, a trace of local activity level, (a) detected people nearby, and (b) streaming real-time audio and video into Second Life on request (when authorized at the real-world portal). By assembling all virtual portals into a small region of virtual space, our entire laboratory can be easily browsed and visited by avatars who come to our Second Life land.

ing from one to the other without real-world constraints of physics, walls, and so on—in essence, a fluid approach to browsing and interacting with the physical world. In Figure 6, the virtual portals show a recent photo uploaded from the corresponding real-world camera at their front surface, but also extend into the past, offering images and media clips of prior events on request. The signal trace at the bottom of each portal represents an activity metric compositing sound amplitude and detected motion. It extends into the past as the line moves toward the back of the portal, letting a virtual user identify periods of activity in the physical portal's area. A white "ghost" appears in front of the virtual portal when the real portal detects a proximate user—if it identifies the user, their name appears above the ghost. This apparition also leaves time trails, letting a virtual visitor see when people are present.

If a virtual user requests a real-time connection, video (and audio, if authorized) from the portal streams to the screen at the front of the virtual structure, whereas information from Second Life appears in the real-world portal. Originally, we displayed a static avatar icon on the real portal and only supported a text stream from Second Life (texting is still the standard medium of communication in Second Life). Our recent implementations feature full bidirectional video and audio connections between virtual and real portals, achieved by streaming media to the portals scraped from servers running Second Life clients (see Figure 5b).

Because portals can stream intimate data (for example, video, audio, and identity), some people justifiably perceive them as invasive. Because we're starting to live within this network, it's vital that Media Lab building residents establish control over the system's

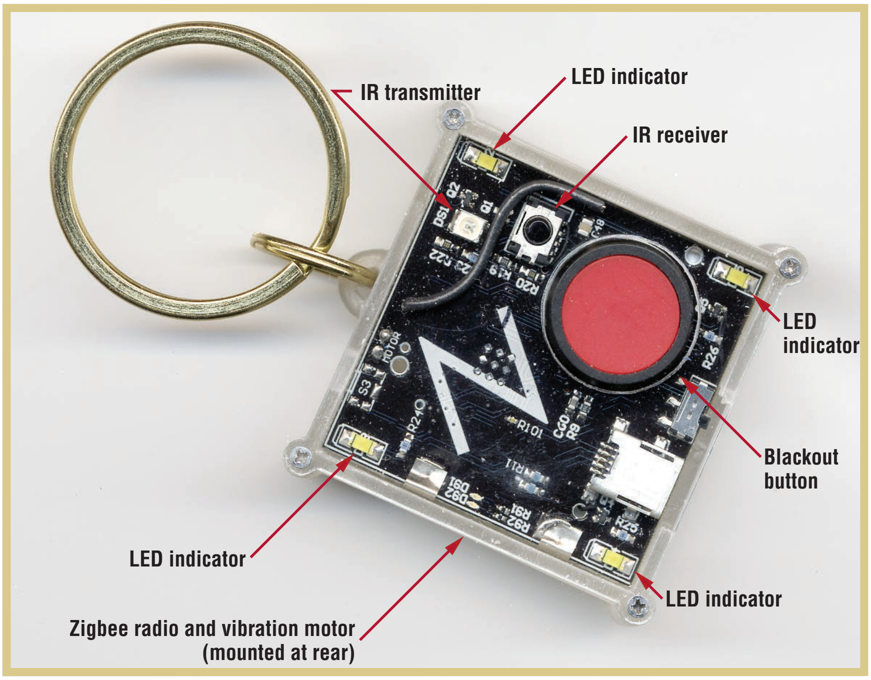

Figure 7. A prototype badge for dynamic privacy control. Each badge transmits an ID and key to the surrounding portals, which consequently stream only features, information, and media that the badge user allows. The large button transmits an immediate blackout command, disabling all audio and video from proximate portals.

boundaries and attain confidence that its capabilities will answer local concerns—a necessity demonstrated in previous ubiquitous media capture installations.[9] We've decided to pursue a solution that works on several levels.[8,10] At the physical level, all portals have an obvious switch on their power line, letting users manually deactivate them (when off, they go dark and the servo motors release, producing an obvious downward nod in the avatar, analogous to sleep).

We're also managing privacy in various ways at a logical level—all these approaches, of course, assume that secure protocols are used and that hardware and code verification and network security are regularly monitored on our devices, just as they are in the PCs and many of the diverse computers that already populate current networks. Our research concentrates more on how people manage privacy rather than how they secure it. In general, the architecture of the portal software determines that any video streaming or media recording will produce obvious visuals on portal displays (the display will also indicate when a portal is logically blindfolded). In the case of end-to-end streaming, all video is reciprocal—you see any entity on the portal screen that is watching remotely.

With 45 portals already distributed through portions of our laboratory, video and audio capture is dense enough to make it difficult to ensure that all portals in range are manually deactivated when privacy is needed. This will only become worse as ubiquitous computing truly arises and potentially invasive media capture becomes an intrinsic property of devices scattered all over our environments. We've addressed this with badge systems, which periodically beacon a unique ID, to wirelessly mediate privacy (Figure 7 shows a prototype privacy badge). Using received signal strength or Chipcon's localization engine, the portals know which badges are potentially in sensor capture range and can passively and dynamically control data access according to the badge users' preferences. When a user presses the red "no" button, however, an immediate opt-out signal is transmitted to block any sensors in range—an important option if users initiate a sensitive conversation. The current protocol answers to the most restrictive privacy setting the portal has received. If we have any indication that the wireless network is being jammed or spoofed, the portals will revert to a conservative privacy level.

Browsing and Interfacing with Mobile Devices

Although designers have created thin Second Life clients to run on mobile devices,[11] full-up virtual worlds might not always be the most appropriate way to represent sensor data on a small display when engaging with sensor networks in which the user is physically immersed. Accordingly, we have developed two devices to explore mobile browsing of and interaction with local sensor nets using 2D displays that feature coarse localization and augmented-reality-style device pointing.

Figure 8 shows our Tricorder,[12] inspired by the handheld device from *Star Trek* that enabled planetary explorers to access relevant information about their extended environment. Doing this with all sensors contained in the Tricorder is still a technical challenge—our approach, however, assumes that a network of localized sensors is already distributed throughout the user's region. The data is polled by the handheld and geographically represented on the device's display. We implemented our Tricorder on a Nokia 770 Linux PDA, augmented by a three-axis tilt-compensated magnetometer that estimates the device's pointing angle and a CC2500 transceiver that enables it to communicate with the PLUG sensor network we described previously.

We render multimodal PLUG data

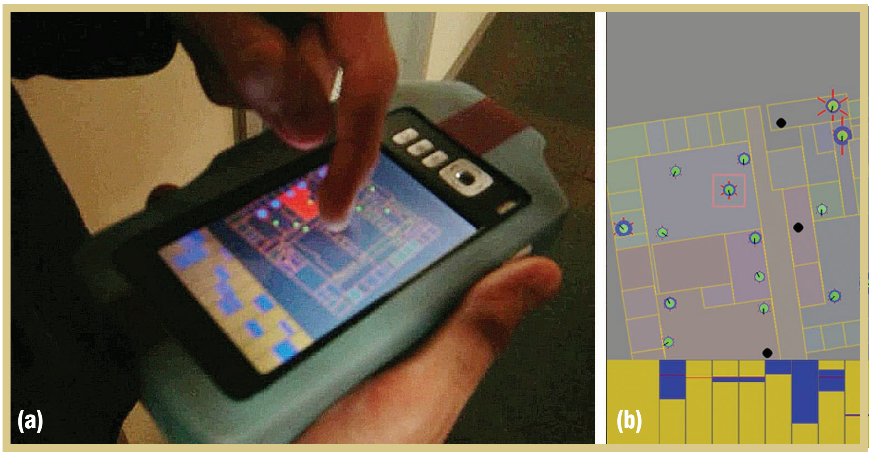

Figure 8. Tricorders. (a) The Nokia 770-based Tricorder in action and (b) a typical Tricorder display. The icons are plotted atop a situated map at the assumed PLUG locations and are animated by data from the corresponding sensors. The bar graph shows sensor data from the selected node.

Figure 9. The Ubicorder. The electronics package at the tablet's upper right contains the compass and accelerometers used for pointing, together with the 802.15.4 radio and IR link for communicating with the portal sensor nodes.

as simple animated dots drawn atop a floor map, as shown in Figure 5 and detailed in our previous work.[12] Localized PLUGs within range coarsely tracked our Tricorder via radio signal strength, which established the plotted map's origin. Rotating the Tricorder correspondingly rotated the map, thus situating the display in real-world coordinates. The user can zoom the view in and out, and plot data from particular nodes by touching corresponding dots.

We designed our subsequent device, termed the Ubicorder[13] (see Figure 9), as a mobile platform that could both browse and interact with sensor networks. Implemented on a Tablet PC supplemented by a compass, IR transceiver, and 802.15.4 radio, the Ubicorder renders data from two families of sensor nodes deployed across our floor—namely, densely deployed ceiling-mounted PIR motion sensors and the multimodal sensor portals we described earlier. Real-time data (here pulled off dedicated servers via the Tablet's WiFi) similarly changes the appearance of dots drawn at sensor locations on a map. The user's estimated location (inferred from the TI's 802.15.4 localization engine) is also plotted as a blob on the Ubicorder's map, with a radius corresponding to the expected location accuracy and an arrow pointing in the direction along which the Ubicorder is oriented. Because the location and orientation resolution were too coarse for reliable pointing at nearby nodes, we provided the Ubicorder with a line-of-sight IR transceiver that could communicate with any portal nodes that it was directed toward. Evolving approaches using coarse radio or sporadic GPS location to constrain a computer vision system hint at a tractable, robust indoor location capability that could improve the close-range performance of such systems.[14] The emergence of 3D augmented reality on smart phones[15] that can query sensor networks and connect with shared online worlds promises to gestate immersive mobile cross-reality.

We also designed the Ubicorder as a programmable interface to sensor networks. Users can select nodes of interest on the Ubicorder's display and interactively condition and combine corresponding sensor signals by an adjustable set of parameters and rules. Because the Ubicorder is portable, the user can bring it into areas where a particular set of activities must be flagged and recognized. The user can then progressively build and tweak the algorithm to approximate the desired result while watching or performing the activity. Although such operation might be beyond the ken of average end users, devices like the Ubicorder will be essential for future ubiquitous computing "utility workers" who need to tweak, adjust, customize, and maintain pervasive sensor systems and their associated context engines. More detail on the Ubicorder's sensor programming interface appears elsewhere.[13]

We've built two cross-reality implementations that differ somewhat in approach—the first used coarse, noninvasive sensors (PLUGs), a standard building map, and heavily leveraged metaphorical animation (ShadowLab) to facilitate browsing of real-world phenomena in a virtual space, while the second (portals) also featured full audio and video streaming but used a more literal and less abstract virtual representation with a geometric layout that reflected intellectual affiliation as opposed to real-world locations. Our experience with these environments indicated a need for both levels of presentation—for example, coarse and grossly obvious representation when quickly browsing, and a more literal and detailed mapping when interacting in particular regions. The benefits of both the metaphorical and direct representations can be realized together through a zooming dynamic. If the user's avatar is far from the visualization site, high-level metaphoric animation can dominate, but as it approaches or touches an apparition of a real-world sensor node, more detail can be revealed and streaming media accessed. Although the Second Life environment has many advantages (for example, simplicity, standardization, and ubiquity), our future research suggests customizations not supported by Linden's environment, so we're exploring other less restrictive platforms (for example, Sun's Wonderland).

We've run metaphorical PLUG-based ShadowLabs many times, but thus far have only tested preliminary cross-reality environments on the portals. We're now keeping the portals activated, and plan to host several "cross-reality" days in the next months, where our research partners, collaborators, and prospective students can virtually visit our laboratory, peering through and appearing on unblocked portals to engage with real-world Media Lab researchers. We plan to soon install many more portals that feature a large display to enable a more engaging virtual penetration into the real world (by analogy, opening windows rather than portholes). We'll also use our privacy badges to explore several approaches to passive throttling of the data provided by the portals, depending on the dynamic privacy preferences of proximate individuals and the social distance of the user launching the query. Handhelds and mobile devices will play an important role in cross-reality applications, both as a source of data to animate their users' environments and avatars and as augmented reality terminals through which local sensor networks can be explored and programmed, as touched on by our Tricorder and Ubicorder initiatives. We expect that this will expand quickly once augmented reality becomes better established on smart phones.

Handhelds and mobile devices will play an important role in cross-reality applications, both as a source of data to animate their users' environments and avatars and as augmented reality terminals

Immersive virtual worlds could some day act as a unifying metalayer for pervasive computing, where sensed phenomena from the real world are extrapolated, processed, interpreted, represented, and interacted with before being projected back down into reality through appropriate incarnations. Cross-reality looks to build an accessible medium in which real phenomena can fluidly meet information constructs—a crossroad where collective human perception can be naturally plugged into the exponentially expanding reach of sensor and actuator networks.

Video clips and other information relating to the projects described here are at www.media.mit.edu/resenv/projects.html.

ACKNOWLEDGEMENTS

We thank our colleagues in the Responsive Environments Group who helped out with these systems, especially Michael Lapinski, Bo Morgan, Alex Reben, Matt Aldrich, and Mark Feldmeier for hardware and software development. We also appreciate the help of Kent Quirk and colleagues at Linden Lab for technical support. This work was supported by the Things That Think Consortium and the other research sponsors of the MIT Media Laboratory

REFERENCES

1. V. Vinge, *Rainbows End*, Tor Books, 2006.

2. C. Stross, *Halting State*, Ace Books, 2007.

3. P. Milgram and F. Kishino, "A Taxonomy of Mixed Reality Visual Displays," *IEICE Trans. Information and Systems*, vol. E77-D, no. 12, 1994, pp. 1321–1329.

4. J. Lifton et al., "A Platform for Ubiquitous Sensor Deployment in Occupational

and Domestic Environments," *Proc. Int'l Conf. Information Processing in Sensor Networks* (IPSN 07), ACM Press, 2007, pp. 119–127.

5. J. Lifton and J.A. Paradiso, "Dual Reality—Merging the Real and Virtual," to be presented at the 1st Int'l Conf. Facets of Virtual Environments (FAVE 09), Springer LNICST, 27–29 July 2009.
6. M. Musolesi et al., "The Second Life of a Sensor: Integrating Real-World Experience in Virtual Worlds Using Mobile Phones," *Proc. 5th ACM Workshop Embedded Networked Sensors* (HotEmNets 08), ACM Press, 2008, pp. 12–16.
7. J.A. Paradiso et al., "Identifying and Facilitating Social Interaction with a Wearable Wireless Sensor Network," to be published in the *J. Pervasive and Ubiquitous Computing*, Springer, 2009.
8. M. Laibowitz, N.-W. Gong, and J.A. Paradiso, "Wearable Sensing for Dynamic Management of Dense Ubiquitous Media," *Proc. 6th Int'l Workshop on Wearable and Implantable Body Sensor Networks* (BSN 09), IEEE CS Press, 2009, pp. 3-8.
9. V. Bellotti and A. Sellen, "Design for Privacy in Ubiquitous Computing Environments," *Proc. 3rd European Conf. Computer-Supported Cooperative Work* (ECSCW 93), Kluwer Academic, 1993, pp. 77–92.
10. N.-W. Gong, "Configurable Dynamic Privacy for Pervasive Sensor Networks," master's thesis, MIT Media Lab, 2009.
11. R. Roush, "New Portal to Second Life: Your Phone," *Technology Rev. Online*, 16 Feb. 2007; www.technologyreview.com/Infotech/18195.
12. J. Lifton et al., "Tricorder: A Mobile Sensor Network Browser," *Proc. ACM Computer-Human Interaction Conf.—Mobile Spatial Interaction Workshop* (CHI 07), 2007; http://msi.ftw.at/program.
13. M. Mittal, "Ubicorder: A Mobile Interface to Sensor Networks," master's thesis, MIT Media Lab, 2008.
14. H. Hile and G. Borriello, "Information Overlay for Camera Phones in Indoor Environments," *Proc. 3rd Int'l Symp. Location- and Context-Awareness* (LoCA 07), LNCS 4718, Springer, 2007, pp. 68–84.
15. D. Wagner and D. Schmalstieg, "Making Augmented Reality Practical on Mobile Phones, Part 1," *IEEE Computer Graphics and Applications*, vol. 29, no. 3, 2009, pp. 12–15.

the AUTHORS

Joshua Lifton works as a technologist and developer for the Electric Sheep Company, the Emmy Award-winning virtual-worlds shop. His research interests include merging the real and virtual worlds in the context of a mainstream commercial environment. Lifton completed his PhD at the Massachusetts Institute of Technology's Media Laboratory. Contact him at lifton@electricsheepcompany.com.

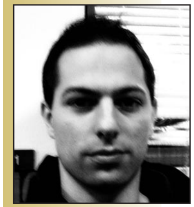

Mathew Laibowitz is a PhD candidate in the Massachusetts Institute of Technology Media Laboratory's Responsive Environments Group. His research seeks to combine his interests in narrative storytelling with pervasive human-centric sensor and camera networks to develop applications for a world of ubiquitous media and new methods of content creation. Laibowitz has a BS in Electrical Engineering and Computer Science from Columbia University. Contact him at mat@media.mit.edu

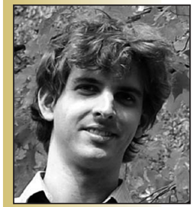

Drew Harry is a first-year PhD student studying the design of mediated social spaces in the Massachusetts Institute of Technology Media Laboratory's Speech + Mobility Group. Most recently, he has designed virtual spaces for meeting, presentation, and play that expand our ideas about how we can use the properties of virtual environments to enable new kinds of interaction. Harry has an MS in media arts and sciences from MIT as a member of the Sociable Media Group. Contact him at dharry@media.mit.edu.

Nan-Wei Gong is a graduate student in the Massachusetts Institute of Technology Media Laboratory's Responsive Environments Group. Gong has an MS in materials science and engineering from National Tsing Hua University. Contact her at nanwei@media.mit.edu.

Manas Mittal is a software engineer at Apple. His research interests include sensor networks, ubiquitous computing, and human-computer interaction. Mittal has an MS from the Massachusetts Institute of Technology Media Laboratory. Contact him at manas@media.mit.edu.

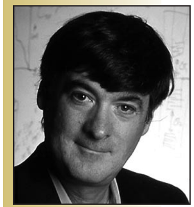

Joseph A. Paradiso is an associate professor of Media Arts and Sciences at the Massachusetts Institute of Technology Media Laboratory, where he directs the Responsive Environments Group, which explores how sensor networks augment and mediate human experience, interaction, and perception. In addition, he codirects the Things That Think Consortium, a group of industrial partners and Media Lab researchers who explore the extreme fringe of embedded computation, communication, and sensing. Paradiso has a PhD in physics from MIT's Lab for Nuclear Science. He's a senior member of the IEEE and the American Institute of Aeronautics and Astronautics, and a member of the APS, the ACM, and Sigma Xi. Contact him at joep@media.mit.edu.

For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.