Dual Reality: An Emerging Medium

by

Joshua Harlan Lifton

S.M., Massachusetts Institute of Technology (2002)
B.A. Swarthmore College (1999)

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning,
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Media Arts and Sciences

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2007

Author ____________________
Program in Media Arts and Sciences
September 6, 2007

Certified by ____________________
Joseph A. Paradiso
Associate Professor of Media Arts and Sciences
Program in Media Arts and Sciences
Thesis Supervisor

Accepted by ____________________
Deb Roy
Chairperson, Departmental Committee on Graduate Students
Program in Media Arts and Sciences

Dual Reality: An Emerging Medium

by

Joshua Harlan Lifton

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning,
on September 6, 2007, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Media Arts and Sciences

Abstract

The commoditization of low-power radios, a rich set of sensors, longer-lasting batteries, and feature-rich microcontrollers has prompted significant research efforts to imbue physical environments with the responsiveness and awareness afforded by ubiquitous, unobtrusive, low-maintenance sensor networks. However, despite these technical advances, there has been relatively little progress toward finding compelling applications enabled by such sensor networks. What few applications have been demonstrated generally use sensor networks to passively monitor environments either inaccessible or uninteresting to people, such as remote wilderness, factory floors, and health care scenarios. Yet, by definition, any "killer application" of sensor networks must be both popular and widespread.

At the same time, online virtual worlds promising complete freedom of creation and interaction are quickly becoming economically, socially, and technically feasible and are making inroads into the mass media market. Yet, despite their popularity, or maybe even because of it, today's online virtual worlds are marred by a stagnation and emptiness inherent in environments so disconnected from the physical world. Furthermore, the demand for richer modes of self-expression in virtual worlds remains unmet.

This dissertation proposes the convergence of sensor networks and virtual worlds not only as a possible solution to their respective limitations, but also as the beginning of a new creative medium. In the "dual reality" resulting from this convergence, both the real and virtual worlds are complete unto themselves, but also enhanced by the ability to mutually reflect, influence, and merge into each other by means of sensor/actuator networks deeply embedded in everyday environments. As a medium, dual reality has the potential to elevate mass creation of media to the same heights television elevated the mass consumption of media and the Internet elevated the mass communication of media.

This dissertation describes a full implementation of a dual reality system using a popular online virtual world and a human-centric sensor network designed around a common electrical power strip. Example applications, interaction techniques, and design strategies for the dual reality domain are demonstrated and discussed.

Thesis Supervisor: Joseph A. Paradiso
Title: Associate Professor of Media Arts and Sciences, Program in Media Arts and Sciences

Dual Reality: An Emerging Medium

by

Joshua Harlan Lifton

The following people served as readers for this thesis:

Thesis Reader____

Deb Roy
Associate Professor of Media Arts and Sciences
Program in Media Arts and Sciences

Thesis Reader____

Samuel Madden
Associate Professor of Electrical Engineering and Computer Science
Department of Electrical Engineering and Computer Science

Acknowledgments

Thanks are due to many people, for many reasons.

To my committee members, Deb Roy and Sam Madden, for their patience, guidance, insight, criticism, and encouragement. They made sure I was where I should have been and helped me get there when I wasn't.

To Brian, Deb, Mike, and Peggy, for the conversations and laughs that kept me sane through all the late nights.

To Lisa Lieberson, NecSys, Linda Peterson, Pat Solakoff, and Gigi Shafer, for holding together the metaphorical ship always on the verge of flying apart, and making it look easy.

To Greg Tucker, Kevin Davis, Pei Wang, and Cornell King, for holding together the literal ship always on the verge of flying apart, and making it look easy.

To Mark Feldmeier, Yasuhiro Ono, Cameron Lewis, Mat Laibowitz, Michael Lapinski, and Manas Mittal, for the ideas, sweat, and tears they poured into the Plugs.

To the participants of the Developing Applications for Sensor Networks class, for the ideas, sweat, and tears they wrung from the Plugs.

To Pattie Maes, John Lester, and all the participants of the IAP Second Life Workshop, for their creativity and curiosity.

To Frank Moss, Judith Donath, Will Glesnes, Jon Ferguson, Henry Holtzman, Paula Aguilera, Jonathan Williams, Becky Bermont, Drew Harry, Jun Ki Lee, Barbara Barry and her Cafe Trio, and all the participants in the IBM Virtual Worlds Conference, for proudly, efficiently, and with good humor flying the demo or die colors.

To Grace Colby, Henry Jenkins, Beth Coleman, and Sibley Verbeck, for letting me pick their large brains.

To Senior Haus, for supplying fervor when in demand, and to the MIT Tae Kwon Do Club, for draining fervor when in excess.

To the Responsive Environments Group, Tangible Media Group, Physical Language Workshop, and Robotic Life Group for defining my view of the Media Lab.

To Hayes, Guy, Alyssa, and Ela for playing it straight and evening my keel.

To Tam and Chris, for more encouragement than is healthy.

To Mom and Dad, for never batting an eye, and to Zach, for still having better kung fu.

Finally, to Joe Paradiso, for simply being the best advisor anyone could ever hope for.

In the time it has taken to complete this dissertation, I've met many wonderful people and lost others. This work is dedicated to the memories of the latter: Big Jimmy, Debbie, Joanne, and Sergio.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

> *"The most merciful thing in the world, I think, is the inability of the human mind to correlate all its contents. We live on a placid island of ignorance in the midst of black seas of infinity, and it was not meant that we voyage far. The sciences, each straining in its own direction, have hitherto harmed us little; but some day the piecing together of disassociated knowledge will open up such terrifying vistas of reality, and of our frightful position therein, that we shall either go mad from the revelation or flee from the light into the peace and safety of a new dark age."*
>
> *H.P. Lovecraft*

Examining today's popular media and the demographics they play to, it would not be far-fetched to deduce that people are happier to consume media than they are to create media. However, it is equally clear that media creation is on the rise among more and more people, being fed and feeding on the rise of media creation technologies, such as wikis, digital cameras, and audio editing software. Even more prevalent is the surge in the last ten years of people sharing media with each other. Given that the act of sharing media is at a basic level an act of self-expression and that many of the newer forms of media are as participatory as they are consumptive, the distinctions between consuming, creating, and communicating media begin to blur. Indeed, mass media might finally earn its name as the masses become involved in its creation and communication as well as its consumption.

The equalization among creation, communication, and consumption of media depends in large part on the available media technologies. In broad terms, if the printing press, radio, and television were the media engines of the past, and the Internet is the media engine of the present, then what will be the media engine of the future? There are two points to consider before answering this question. First, these different media engines complement, feed into, and influence each other rather than immediately supplant one another; the printing press, radio, and television allow for centralized distribution to a wide audience and the Internet allows for infinite recombination and decentralized distribution. Second, the current media engines and their associated technologies provide ample means of both consuming and communicating media, but lack a truly accessible means of creating media. The media engine of the future, then, is one which simplifies media creation to the point of making it accessible to nearly everyone. To put creation of media on the same footing as consumption and communication of media, the act of creation must be as trivial as watching television or turning on a peer-to-peer file sharing client.

1.1 Dual Reality

At the heart of this dissertation is the concept of "dual reality", which is defined as follows:

Dual Reality: An environment resulting from the interplay between the real world and the virtual world, as mediated by networks of sensors and actuators. While both worlds are complete unto themselves, they are also enriched by their ability to mutually reflect, influence, and merge into one another.

The dual reality concept, in turn, incorporates two key ideas – that data streams from sensor networks are the raw materials that will make media creation as trivial as media consumption and media communication, and that online 3D virtual worlds are an ideal venue for the manifestation of the content generated from such sensor data streams.

In particular, this dissertation proposes that data from networks of intelligent sensors embedded throughout our everyday environments will fuel the next media engine, one which allows people to create content simply by going about their lives as usual and then sifting, filtering, and processing, automatically or manually, the resulting sensor data streams. Such sensor networks have the potential to free people from the task of intentionally staging and recording content by relying instead on people's natural interactions with the physical world to generate digital content automatically.

However, obtaining the sensor data from which to generate content is only a part of the process of media creation – the actual embodiment of the sensor data-derived content and the mapping between raw data streams and the final embodiment are equally critical. This dissertation proposes online 3D virtual worlds as a stage suited to both embodying and mapping sensor network data streams. As social spaces, virtual worlds provide an incentive for people to shape sensor data from their real lives into a form of self-expression in the virtual world. As environments free of many real-world constraints, virtual worlds provide the flexibility to do so in many ways.

In essence, sensor networks will turn the physical world into a palette, virtual worlds will provide the canvas on which the palette is used, and the mappings between the two are what will make their combination, dual reality, an art rather than an exact science. Of course, dual reality media will in no way replace other forms of media, but rather complement them. Indeed, the end product, that which can be consumed and shared, is unlikely to outwardly resemble current forms of media, even if it is just as varied.

Finally, a complete consideration of dual reality must also include the possibility of "sensor" data from the virtual world embodied in the real world. Insofar as technically feasible, dual reality is bi-directional.

1.2 Document Overview

The main thrust of this work is to develop and demonstrate the technologies, design principles, and usage scenarios necessary to begin exploring the speculations portrayed above, and perhaps make them slightly less speculative in the process. This dissertation aims to concretely demonstrate the dual reality concept. Specifically, sensor networks are applied as the glue between the real and virtual worlds in order to enrich them both. To this end, a sensor network, in the form of 35 heavily augmented electrical power strips distributed throughout an approximately 2500-m^2 research lab space, connects the lab space to a virtual version of the space, among other things, located within a popular commercial online virtual world. Data from the real world are collected by the sensor network and used to influence the virtual world and vice versa. Example applications are explored and interaction and mapping techniques discussed.

The remainder of this chapter describes the contributions of this thesis and the research domains in which this work is situated. The next chapter goes into example scenarios of how this work might be used. Chapter 3 expands on the motivation driving this dissertation and puts it in the context of related work. Chapter 4 details the hardware and software systems implemented in support of the contributions of this dissertation. Chapter 5 evaluates and discusses the systems of the preceding chapter. The final chapter concludes with lessons learned and future directions.

1.3 Contributions

This dissertation makes the following contributions:

- The framing of dual reality as a new application domain for sensor networks, virtual worlds, and media creation.
- State-of-the-art hardware and software tools for realizing dual reality interactions.

- Example dual reality environments, interactions, and scenarios.
- A new human-centric approach to sensor network design that addresses practical issues of pushing sensor networks to ubiquity, as needed to realize dual reality environments.
- Evaluation and discussion of various aspects of such dual reality systems.

1.4 Research Domains

At a high level, the work and contributions of this dissertation span the following well-established research domains. A more careful overview and analysis of these domains is presented in Chapter 3.

1.4.1 Sensor Networks

A sensor network can be loosely defined as a collection of electronic nodes that cooperate in order to observe various aspects of the physical world. At a practical level, a sensor network node must have a power source, a means of communicating with other nodes, a set of sensors with which to observe the world, and the computational resources to process those observations as necessary. By this broad definition, many systems can be considered sensor networks: cell phone networks, radio telescopes, and earthquake detection monitors all fit this description. That said, most contemporary research focuses on designing and deploying inexpensive, low-power, easily embedded and maintained, wireless sensor networks, within which the examples just given are typically not included. Power-efficient sensing and network routing protocols, micro- and nano-power analog and digital electronics, and distributed data processing techniques all fall under the purview of contemporary sensor network research. A logical extension of sensor networks, however broadly or narrowly defined, is the ability not only to observe the physical world, but also to act upon it, thus drawing on the lessons of distributed control systems.

1.4.2 Ubiquitous Computing

This work approaches sensor networks from the perspective of ubiquitous computing and considers all the above definitions of "sensor network" as needed, using the term as a catchall. Ubiquitous computing can be loosely defined as the seamless integration of computation into the everyday human environment (as opposed to remote, inhospitable, or automated environments devoid of people, where sensor networks are also deployed). Given these definitions, sensor networks can be considered the sense organs of ubiquitous computing, providing an awareness of the environment with which it has merged.

1.4.3 Online Virtual Worlds

There is no definition of online virtual worlds that is both agreed upon and useful. Everything from mailing lists, wikis, and chat rooms, to massively multiplayer immersive 3D tactical and strategy games can be considered online virtual worlds. This dissertation could apply to all these definitions, but primarily focuses on the conecpt of virtual world as introduced in science fiction works by authors such as William Gibson [1] and Neil Stephenson [2]. This type of online virtual world is characterized by an immersive 3D environment, fluid interactions among inhabitants, and some level of ability for inhabitants to shape their environment.

Chapter 2

Extended Example

"When the going gets weird, the weird turn pro."

Hunter S. Thompson

Among the prominent features of Boston's skyline is the John Hancock Berkeley Building, colloquially known as the Old John Hancock Building [3]. Although not particularly tall, it frequently draws Bostonians' gazes because of its famous top-mounted weather forecasting lights, the behavior of which is summarized by this well known rhyme:

Steady blue, clear view.
Flashing blue, clouds due.
Steady red, rain ahead.
Flashing red, snow instead.

The Old John Hancock Building's lighted weather beacon illustrates how a simple sensor data representation, if well-crafted, can serve to enliven a space. Displaying the weather forecast on a large alphanumeric display would certainly not have the same effect; the feeling of local custom and insider information engendered by the beacon is palpable.

Applying the concepts underlying the weather beacon to real world sensor data representations in virtual worlds promises to be even more compelling for two reasons. First, the cost and complexity of designing, implementing, and modifying virtual-world representations is typically orders of magnitude less than doing so for comparable real-world representations, allowing for a much greater number and variety of representations. Second, the need to enliven virtual worlds is much greater than the need to enliven the real world. Information flow from the world to the observer is infinitely more rich and abundant in the real world than the virtual world. Siphoning some of this real-world information into the virtual world by means of sensor networks not only remedies the virtual world's relative stagnation, but also increases its utility by tying it more closely with real-world concerns.

Of course, as the utility of the virtual world increases, reversing the siphon of information to feed data about the virtual world into the real world becomes more compelling. Devices designed to display stock market or other information streams in an ambient fashion, as the Old John Hancock Building's beacon does with weather, are precursors to this idea [4].

The remainder of this chapter illustrates how the work described in this dissertation can be applied. A fictional description of a near-future scenario hints at the possibilities of a world in which sensor networks and virtual worlds are ubiquitous and accessible. The scenario makes only modest extrapolations of what is currently possible based on the work of this dissertation.

2.1 Near-future Usage Scenario

Vlad logs into VirtuaPlanet and goes directly to Club Zed to meet up with some of his friends. He's just back from a vacation without a net connection and hasn't seen them for a while. He wades through the usual nonstop crowd of dancing avatars to his usual table and finds that he is the first to arrive. He sits down and flips through the messages people have left for him while he was away. In the real world, he has just gotten comfortable on the commuter train on his way to work. Vlad is friends with Min, the owner of Club Zed, and understands the subtle upbeat shift in lighting and music on the Club Zed dance floor to mean Min is now near a terminal somewhere, but not logged on. Vlad and Min have never met in real life and Min didn't bother to tell him that this cue indicates that she is near her computer in her bedroom at home.

In fact, she is in the same timezone as Vlad, but only now waking up and getting out of bed to start her work day from the comfort of her home as the manager of Club Zed and several other VirtuaPlanet properties. Vlad knows to keep an eye on Club Zed's ceiling. Sure enough, as Min logs in, Club Zed's usual lights and music skip a split second and the room is lit up with a shower of grey coming down from the sky light above as a confused mosaic of a dull sun, clouds, traffic, a mess of newspapers, spilled coffee, and a living room set appears, swirls, shifts, and quickly fades. The mosaic is composed of low resolution images from the hundreds of sensor nodes Min has scattered about her high-rise apartment. Her sensor network's privacy settings are such that images of people are never used. The mosaic flashes up every day when she logs in – it's her way of connecting the two spaces, one real and one virtual, where she spends most of her time. It's fast enough that only those looking for it would be able to make anything out. The regular dance lights and music flare up to fill the void as Club Zed frenetically continues it's 24/7 party.

Min and Vlad exchange greetings. Shortly thereafter, David, TJ, and Katrine join them. Each of their avatars takes a seat on the plush sofa surrounding Vlad's table. The poses the avatars affect as they sit reflect some aspect of their owners' real world surroundings. Vlad's on-body sensor network recognizes the characteristic hum and motion of the daily commute and tells his laptop that Vlad is on a train. His laptop uses this information to trigger his avatar's "commuting by train" sitting pose. Vlad's avatar leans back, brings his hands in front of him, and impatiently drums his fingertips together as his hair is blown back by a virtual wind only he can feel. Aside from her daily mosaic, Min likes to keep her worlds separate and chooses not to stream any other sensor data into VirtuaPlanet; her avatar maintains the default sitting pose. In contrast, David prefers using a dense on-body sensor network for fine-grained realistic control of his avatar. The others take this as a sign of his relative inexperience with virtual worlds and they expect him to grow out of it soon after realizing the benefit of a virtual world lies largely in the fact that it is not a copy of the real world. Katrine's avatar's collar, cuffs, and fingernails all are all quickly changing color and almost imperceptibly changing size, indicating she is most likely hooked into her car's vehicular sensor network and is zipping down a highway in light afternoon traffic. TJ's avatar is its usual mess of pulsating neon spines and bubbles, the meaning of which no one but he knows.

After everyone is seated, the conversation turns to questions about Vlad's trip. Expecting as much, Vlad placates his friends by giving them each a small box tied shut with a ribbon. Each box contains a sensate short of Vlad's trip. Vlad composed the short from data collected from his trip's various personal area sensor networks, which consisted of his on-body network and any nearby networks open to him, such as those in his hotel room, the city park, or the amusement park's roller coaster. When making sensate shorts, Vlad usually accepts the default short produced by his editing software, but this time he spent the several hours during his flight back home customizing the shorts. The sensate short will start when an avatar opens the bow-tied box. Once started, the sensate short takes the avatar on a fast-forwarded, highlights-only, virtual version of Vlad's trip that will last about 5 minutes, during which time the avatar will "experience" multiple modalities of recorded sensor data, such as movement, sound, images, wind, and temperature. Vlad's friends pocket their sensate shorts and save them for later. In the real world, Vlad's train enters a tunnel and begins to slow down. His on-body sensor network infers that he will be getting off the train soon. His avatar automatically reacts to this trigger by standing up and morphing into a different set of clothes. The friends say goodbye and go their own ways. Vlad logs out of VirtuaPlanet as the train comes to a stop.

Once at work, Vlad goes straight into a meeting with research sponsors. He is scheduled to lead a group of 150 sponsors through a broad introduction to ten of the lab's research groups, all within half an hour. Between the usual inhabitants of the lab and the 50 sponsors who have already showed up, the lab is crowded. The remaining 100 sponsors are logged into the lab's VirtuaPlanet reflection, consisting of a to-scale floorplan with ethereal walls for full visibility and easy access to any part of the lab. They too can see the bustle of the lab – the ephemeral whirlwinds, spurts of flame, sounds of droning machinery, and sudden orbs of ball lightning are all part of the Post Apocolyptica Theme (also available as a desktop theme and cell phone ringtone) dictating the transformation of raw sensor data collected in the real lab into thematically appropriate metaphorical representations of varying degrees of subtlety. In the physical lab, Vlad puts his cell phone into "lab tour" mode, in the process automatically restricting virtual lab access to sponsors and lab affiliates. As he gives the real-world tour through the lab, his avatar mirrors his movements in the virtual lab and live audio and video from his cell phone give the virtual tour members a window into the real lab.

Chapter 3

Theory and Rationale

> *"There is a theory which states that if ever for any reason anyone discovers what exactly the Universe is for and why it is here it will instantly disappear and be replaced by something even more bizarre and inexplicable. There is another that states that this has already happened."*
>
> *Douglas Adams*

This chapter lays out the theoretical motivations behind this work, outlines the main issues that arise given these motivations, and describes the approaches taken to overcome these issues. Along the way, the pertinent related work is overviewed.

3.1 Sensor Networks

By their nature, sensor networks augment our ability to understand the physical world in ways beyond our innate capabilities. With sensor networks and a record of the data they generate, our senses are expanded in space, time, and modality. As with previous expansions of our ability to perceive the world, some of the first and perhaps in the long run most important upshots will be the creation of new forms of self-expression; take for example

Doc Edgerton's pioneering high-speed photographs [5, 6]. Of course, it is impossible to perfectly predict at the outset how a technology will be used – the initially limited quality of a technology or the particular circumstances of its development and introduction can restrict its actual or imagined potential. For example, Thomas Edison famously thought the phonograph would be used primarily for dictation rather than music, hence the term "Dictaphone" [7]. However, even the Dictaphone example strengthens the argument that self-expression often trumps other uses of a given technology.

Most, if not all, sensor network platforms in use today are characterized by an emphasis on a low-power, unobtrusive, versatile design and an understanding, implicit or otherwise, that a network's sensor nodes are to be handled only briefly, if at all, by expert researchers between long periods of unattended operation. Although this paradigm has generally served the research community well and fits many application scenarios, it precludes a full exploration of the sensor network application space. In particular, this paradigm is not fully appropriate for ubiquitous computing settings. Indeed, the commonly cited vision of living in a truly aware environment, one which senses and responds to our every action, does not necessarily imply that the sensor nodes upon which this vision is built need be low-power, unobtrusive, or versatile.

This thesis introduces the "Plug", a sensor node modeled on a common electrical power outlet strip and designed specifically for ubiquitous computing environments. See Figure 3-1. As with most sensor nodes, each Plug has its own microcontroller for tending to a host of sensors, actuators, and wireless and wired communication interfaces. In addition, a Plug node can serve as a normal power strip, providing four standard three-prong US electrical outlets. As such, a Plug node must be plugged into a power outlet to operate, making the issue of extreme energy conservation, such as needed for long-term battery-powered deployments, nearly irrelevant. Furthermore, considering a Plug node's comparatively large size (approximately 20cm×7cm×12cm) and weight (approximately 1kg), it's difficult to argue that a Plug is unobtrusive based on its physical specifications alone. Finally, a Plug node's versatility is limited to that of a regular power strip – it is not mobile, wearable, embeddable, or otherwise easily reconfigurable to be anything but a power strip.

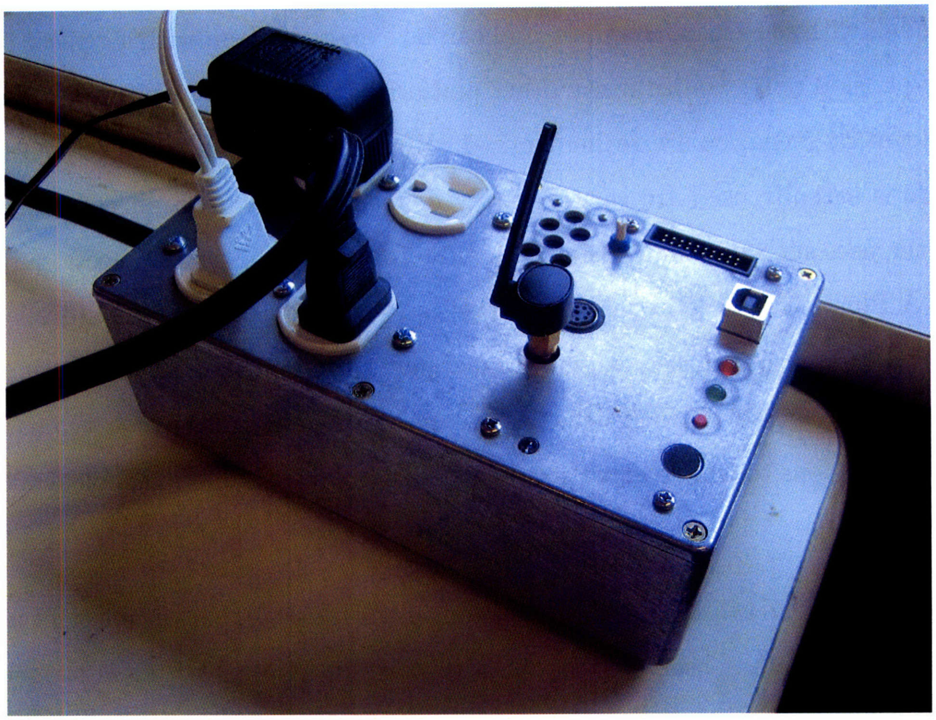

Figure 3-1: A typical usage scenario for a Plug sensor node.

Nonetheless, within the context of ubiquitous computing, a network of Plug nodes is ideally suited for sensor network research and applications. By their nature, ubiquitous computing scenarios take place in environments normally inhabited by people, of which the home and the workplace are the dominant examples. Both these settings are infused with ample electrical power, typically in the form of wall sockets spaced every two or three meters. Thus, the need for exceptionally low-power sensor nodes is mitigated so long as the nodes need not be mobile. Similarly, what is considered unobtrusive depends on the setting. Power strips are common in nearly every home and workplace setting. A cursory examination of a typical 14-square meter office used by three graduate students revealed no less than 10 power strips, not including wall outlets. Despite this pervasiveness, most of the time, power strips go nearly unnoticed. (The exception being when they go missing). A true metric of a sensor node's obtrusiveness must take into account how well it blends with its environment, not just its physical size and weight. Finally, the versatility of a Plug node is somewhat two-sided. For example, like other versatile sensor network platforms, Plug nodes are richly multi-modal, with ten sensor channels, and can be expanded upon with a generic digital and analog expansion port. Unlike many other platforms, the Plug node's physical form factor and deployment scenarios are rather limited. However, although the mechanics of how a Plug node is actually used are very narrowly defined (i.e., you can plug electrical devices into it), the uses such mechanics afford are only limited by the uses of electrically powered devices. A Plug is versatile in the same sense modern electrical infrastructure is versatile. Moreover, a Plug node's power switching capabilities and built-in speaker give it significant actuation advantages over most other sensor network platforms, which typically require hardware extensions to enable actuation. Chapter 4 gives a complete technical specification of the Plug.

The crux of the Plug's utility as a sensor network platform lies in part in the tight integration of the observed and the observer. That is, a primary purpose of a Plug is to measure how it is being used or to ascertain context relating to its immediate neighborhood. The fact Plug nodes have a well-defined use at all is unusual in itself and contrasts sharply with most sensor network nodes, which are largely designed to be hidden throughout an environment and not interacted with directly. The Plug platform may be the first to attract the very phenomena it is meant to sense. This principle of designing sensor networks as integral parts of their environments, as opposed to additions layered on top thereof, is central to the Plug platform and will likely play a major role in bringing sensor networks out of the research lab into the real world. Intelligently augmenting commonplace devices already used for dedicated applications is a clear path toward ubiquitous sensing – the cost of adding sensing, networking, and computing capabilities to individual devices is relatively low and even a single device has utility, allowing the cost of the entire network to be spread over time.

3.1.1 Related Work

It is well documented that the term "sensor network" encompasses many different instantiations. The current formulation of a sensor network is less than 10 years old [8], whereas the term "sensor network" has been in use for at least 30 years (see Figure 3-2) [9], and has come to refer to everything from hopping land mines [10] to artificial sensate skins [11]. The Plug platform expands this list to include electrical power infrastructure.

Studying the power consumption of various electrical devices has a rich history. Such information can be used to identify classes of devices [12, 13] or even individual devices [14], detect and predict electrical and mechanical faults in motors [15], monitor energy costs and consumption [16], and as a form of surveillance [17].

The SeeGreen system uses power line communication to monitor and control metering devices attached to electrical appliances, but does not extend to other sensing modalities or communication channels [18]. The "Kill A Watt" is a commercially available surrogate electrical outlet for home energy consumption monitoring, displaying volts, amps, watts, Hz, and VA for a single electrical outlet [19]. The Home Joule is a night light that displays real-time energy consumption and electricity prices [20]. A Spy Labs product makes evident the privacy concerns related to embedding sensing capabilities into commonplace objects – the AGS-01 is a power strip with built-in GSM cell phone transmitter which can be used to monitor surrounding audio from anywhere in the world by simply phoning the number of the inserted SIM card [21]. At another extreme, Chip PC Technologies' Jack PC product is a fully functional thin client computer designed to fit into a standard LAN wall socket with a monitor, mouse, and keyboard plugging directly into the wall [22]. Power strips themselves are evolving in form and function – it's common now to see them augmented with surge protectors, noise filters, and pass-through connectors for data, cable TV, and phone lines, and designers are looking at radically new packaging to improve usability, such making the physical form factor reconfigurable and the plugged-in power cords more easily differentiable [23].

Intel Research and USC/ISI built and deployed a conference room monitoring and reservation system using a sensor network [24]. This system is notable because it involved a real-world sensor network application within a workplace environment and it demonstrated how existing infrastructure, in this case motion detectors for turning on and off lights, can be leveraged by the sensor network. More recently, the OpenSpace project equipped public areas of the third floor of the MIT Media Lab with 150 motions sensors as an attempt to

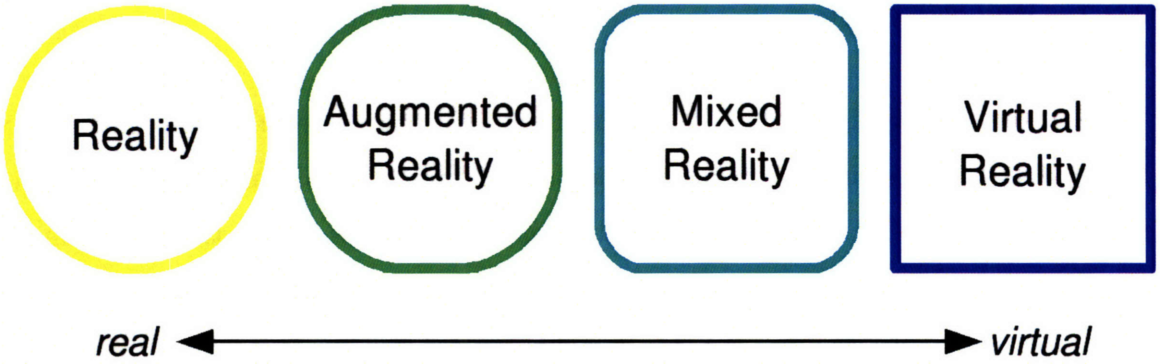

Figure 3-3: A virtual worlds taxonomy as viewed on the real-virtual axis.

take notice, even if they don't yet know how exactly the technology will be used or what it is good for. There are, of course, important differences as well. Whereas the Web to some extent mimics the printed word and image in its presentation, most instances of virtual worlds mimic the spatial and social dimensions of the real world. One incentive for doing so is to take advantage of the highly developed and specialized social mechanisms and biological morphology evolved by humans to interact and navigate within the natural world. The goal may not be, and probably should not be, to replicate all aspects of the real world, but rather only those that facilitate the interaction in a virtual environment. In light of this, imbuing virtual worlds with the ability to sense aspects of the real world is a technique with significant potential.

Of the many axes along which various virtual worlds can be compared, the most relevant for this work is the real-virtual axis, which indicates how much of the constructed world is real and how much virtual. See Figure 3-3. A rough taxonomy can further compartmentalize the real-virtual axis into *reality*, *augmented reality*, *mixed reality*, and *virtual reality*. For the purposes of this work, some of these categories deviate from their commonly accepted definitions and are instead defined as follows:

Reality is simply life in the absence of virtual representations of the world.

Augmented reality has all aspects of reality, as well as an "information prosthetic" which overlays normally invisible information onto real objects. The canonical example introduced in Steven Feiner's augmented reality manifesto involves overlaying information about operincrease social awareness in the work place, while still maintaining a sense of privacy, by visualizing the sensors' activation across time and space [25].

Several recent research efforts have followed a human-centric approach to sensor network design and deployment. The CarTel project is a distributed, mobile sensor network and telematics system that includes approximately 30 automobiles, each equipped with a small embedded computer that collects, processes, and forwards sensor data [26]. A collaboration between the University of Washington and Intel Research focuses on recognizing human activity using sensors embedded in common wearable devices, such as wristwatches and cell phones [27]. A commercial collaboration between NIKE and Apple resulted in the "Nike + iPod" product consisting of a pair of running shoes with an embedded motion sensor that wirelessly communicates data to a portable music player, thereby influencing the music being played [28].

3.2 Ubiquitous Computing

Collections of cooperative, wirelessly networked, intelligent sensors embedded in our surroundings are poised to become integral parts of our everyday work and domestic environments. Although significant challenges must yet be overcome before ubiquitous deployment is realized, such wireless sensor networks are steadily creeping into our lives - cell phones, network-aware digital cameras, and RFID tags are already accepted as the norm. With them come challenging questions of usability and human-network interfacing. Should people be aware of the sensor nodes surrounding them? If so, how? How will people access, search, and sort the copious real-time and stored data available from wireless sensor networks? Will networks be open or closed, centralized or decentralized? Regardless of how these and similar questions are ultimately answered, their consideration will benefit from real-world examples of human interfaces to sensor networks.

A large branch of the human-computer interaction community subscribes to the notion that ubiquitous computing is in our future, that computers will be embedded throughout our environments and we will interact with them seamlessly [29]. To make this vision a reality, the emphasis should not be on computing being embedded in our environments, but rather sensing being embedded in our environments. After all, any seamless interface between people and invisible computers must rely on sensing. As an example of this skewed emphasis, Mark Weiser's seminal ubiquitous computing manifesto mentions sensing only once, and in passing at that [30]. Sensor networks have the potential to become the technology enabling the human-computer interface necessary to realize ubiquitous computing. This has only recently begun to be accepted as an axiom by the ubiquitous computing community as a whole [31].

Weiser's manifesto goes on to lucidly give the electric motor as an example of how technology can disappear into the background. Aside from being coincidentally topical to the Plug, this example illustrates a likely evolution of sensor networks. Taking his example further, when electricity production first began, the thought that it would be available from holes spaced every couple of meters in every wall in every house was looked upon as absurd and highly impractical. Sensor networks must achieve exactly this scale of infrastructure if they are ever to leave the research lab. Just as electricity, and indeed every major utility, is put to use in ways unforeseen at the time of deployment, so too will sensor networks find application.

Any discussion of ubiquitous sensor networks demands consideration of the accompanying privacy and security issues. Whether or not people are willing to submit to such scrutiny is not a question – they already have at every such opportunity and show no signs deviating from this behavior. The deeper questions are whether or not they *should* submit to such scrutiny and, if so, how that scrutiny should be managed. The former boils down to the endless debate of exactly how much freedom should be traded for a given amount of security (or some other benefit). The latter is convincingly portrayed in David Brin's *The Transparent Society* as a binary choice between completely closed networks accessible only by a strict government-like entity and completely open networks accessible to anyone [32]. Clearly, with mass media creation as a focus, this dissertation argues for the completely open option.

3.2.1 Related Work

Much of the work bridging ubiquitous computing and sensor networks takes the form of handheld devices used in conjunction with a sensor network. The Great Duck Island sensor network project for monitoring the habitat of certain birds was among the first examples to use a personal digital assistant (PDA) as a management tool to fine tune the network [33]. The TASK project used a PDA equipped with a radio to inspect and change the state of sensor nodes deployed within radio range [34]. Going beyond network management to conveying actual sensor data, a sniper localization sensor network system used a PDA to display the estimated location of a gunshot [35]. The Cricket indoor location system uses ultrasound pings between sensor network nodes to provide handhelds with location information [36]. The concept of "participatory sensing" suggests using cell phones carried by people as nodes of a sensor network, such that mobile devices form interactive, participatory sensor networks that enable public and professional users to gather, analyze, and share local knowledge [37]. Along the same vein, room-level localization and mapping have been demonstrated based solely on Bluetooth connectivity (as opposed to triangulation) in Bluetooth-rich environments [38]. Leaning more toward the concept of augmented reality [39] are the Periscope [40], Electronic Lens [41], and Point To Discover [42] projects, all of which attempt to tag the physical world with digital information and read the information back with the aid of a handheld equipped with some combination of GPS receiver, digital camera, and orientation sensor, all backed by a central information server.

3.3 Virtual Worlds

As alluded to in Chapter 1, virtual worlds have long been depicted in science fiction literature. The term itself is vague enough to encompass a full spectrum of technologies, from text-based multiple user domains (MUDs) originating in the late 1970s [43] to visually immersive online 3D games commercially available today [44, 45]. Within a limited context, virtual worlds are today what the World Wide Web (or Web) was in 1993; the first widely popular interfaces to virtual worlds are just appearing and entrepreneurs are beginning to

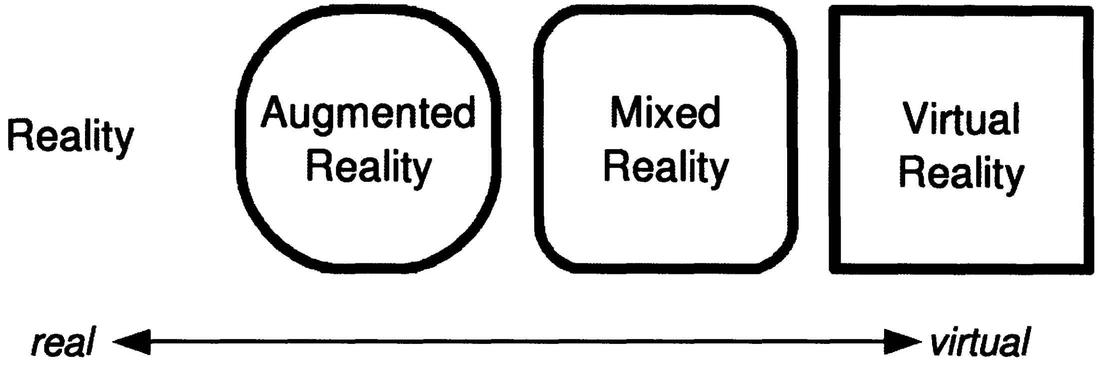

Figure 3-3: A virtual worlds taxonomy as viewed on the real-virtual axis.

take notice, even if they don't yet know how exactly the technology will be used or what it is good for. There are, of course, important differences as well. Whereas the Web to some extent mimics the printed word and image in its presentation, most instances of virtual worlds mimic the spatial and social dimensions of the real world. One incentive for doing so is to take advantage of the highly developed and specialized social mechanisms and biological morphology evolved by humans to interact and navigate within the natural world. The goal may not be, and probably should not be, to replicate all aspects of the real world, but rather only those that facilitate the interaction in a virtual environment. In light of this, imbuing virtual worlds with the ability to sense aspects of the real world is a technique with significant potential.

Of the many axes along which various virtual worlds can be compared, the most relevant for this work is the real-virtual axis, which indicates how much of the constructed world is real and how much virtual. See Figure 3-3. A rough taxonomy can further compartmentalize the real-virtual axis into *reality*, *augmented reality*, *mixed reality*, and *virtual reality*. For the purposes of this work, some of these categories deviate from their commonly accepted definitions and are instead defined as follows:

Reality is simply life in the absence of virtual representations of the world.

Augmented reality has all aspects of reality, as well as an "information prosthetic" which overlays normally invisible information onto real objects. The canonical example introduced in Steven Feiner's augmented reality manifesto involves overlaying information about operating a printer onto the visual field of the user by means of a special head-mounted display [39]. As with Weiser's introduction of ubiquitous computing, Feiner's introduction of augmented reality mentions sensing only peripherally, mainly in the context of on-body sensing and ultrasound transmitters attached to objects. Variations of the augmented reality concept have claimed certain commercial successes, such as Sportvision's "Virtual Yellow 1st and Ten" line used in televised broadcasts of American football games [46].

Mixed reality is one which would be incomplete without both its real and virtual components. For example, the walls and windows of a mixed reality house might be real, but the view out the windows might be virtual, either generated by a projector or as a blue screen effect in a head-mounted display. Without both the real house and the virtual views out the windows, the illusion of a consistent reality is broken. Mixed realities have been studied and implemented primarily as training tools, such as those used by the US military [47].

Virtual reality contains only elements generated by a computer and is typically interfaced to with a computer screen and some user input device, such as a mouse or keyboard. Most computer games, and computer applications in general for that matter, fall into this category. That said, inclusion in this category is only meaningful insofar as the game or application is meant to mimic the real world or aspects of it. Even so, many popular computer games still qualify, such as the SimCity series of games [48].

Each of these four variations on the real-virtual axis represents what is supposed to be a single, complete, and consistent world, regardless of which components are real and which virtual. Although this taxonomy can be successfully applied to most augmented/mixed/virtual reality efforts, it does not address well the concept of dual reality. A **dual reality**, as defined in Section 1.1, comprises a complete reality and a complete virtual reality, both of which are enhanced by their ability to mutually reflect, influence, and merge into each other by means of deeply embedded sensor/actuator networks. See Figure 3-4.

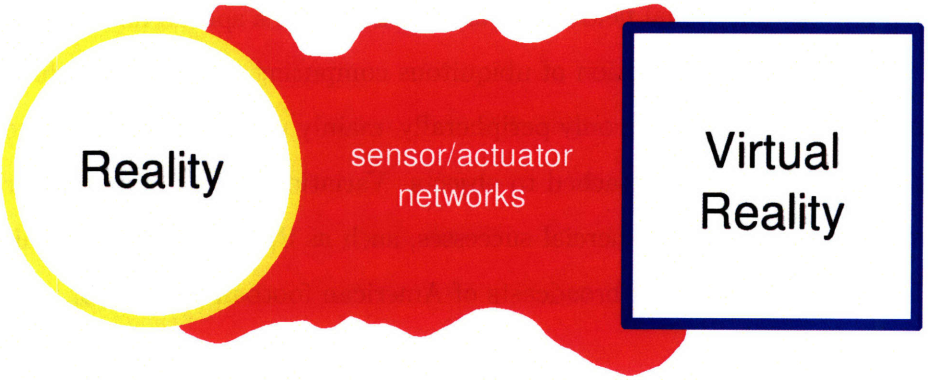

Figure 3-4: Sensor networks merge the real and virtual worlds to form a dual reality.

3.3.1 The *Second Life* Virtual World

As a matter of practicality, the virtual world portions of this work focus exclusively on *Second Life*, an online virtual world launched in 2003 and today still maintained by Linden Lab [49]. The common consensus, shared even by the Linden Lab CTO [50], is that there will be no middle ground for the success of *Second Life*; it will either become wildly popular and change the world, or go the way of the Apple Newton – a good idea whose timing wasn't quite right [51]. The likelihood of *Second Life* succeeding is still an open question. For example, as of this writing, there are approximately nine million registered *Second Life* accounts, but only 50,000 users logged into *Second Life* at any given time [52]. A comprehensive review of all online virtual worlds is beyond the scope of this work and better left to the several websites that specialize in such comparisons [44, 45, 53]. *Second Life* was chosen because of its technical and other advantages in implementing many of the dual reality ideas explored here. In this regard, *Second Life*'s most attractive attributes are:

- **Shared immersive experience**: *Second Life* is an inherently social place with a large and quickly growing population. Unlike browsing the Web, which does not easily permit social interaction among people browsing the same website, *Second Life* "residents" in the same virtual location can freely interact through a chat interface and visuals.

- **Persistent state**: Just as in the real world, any changes made by residents or other entities to the *Second Life* virtual world persist until undone or overwritten. In other words, the world never reboots to a default clean-slate state.
- **Market economy**: Linden Lab actively encourages a market-based economy in which residents can earn "Linden dollars", which can be bought and sold for US dollars, in order to buy virtual or real goods and services. A flexible permissions system enforces intellectual property rights on all transferable goods.
- **Creative medium**: Nearly everything in *Second Life* is created by residents. Even the layout of the land can be modified. A simple graphical construction toolkit allows three-dimensional virtual objects to be created and altered by novice users. A straightforward programming language (Linden Scripting Language, or LSL) can be used to change the behavior of objects, communicate with outside databases and servers, and programmatically manipulate the state of the virtual world within certain security, privacy, and computational resource drain limits.

Everything in *Second Life* exists as some combination of *land*, *avatars*, *objects*, and *scripts*. Land in *Second Life* is mapped directly to Linden Lab server resources, such as computing cycles, memory, and bandwidth. The larger the piece of virtual land, the more of these resources allotted to that land. Each piece of land can support a limited number of avatars, objects, and scripts. As in the real world, land in *Second Life* costs real money and is typically adjacent to land owned by other people. Water, trees, clouds, and celestial bodies are all part of *Second Life* land. Owners of land can freely terraform to suit their needs. Avatars are the virtual manifestation of real people using *Second Life*. The size, shape, appearance, and behavior of an avatar are all highly customizable, although most tend to look vaguely humanoid. All interaction between *Second Life* residents and everything in *Second Life* is carried out through an avatar in one way or another. Avatars can touch, sit, stand, walk, run, fly, speak, etc. Aside from land and avatars, objects compose the remainder of the physical elements of the *Second Life* virtual world. Objects, in turn, are built from one or more primitive three-dimensional solids ("prims"), such as spheres, cubes, tori, and cones. Objects can be designated as physical, phantom, flexible, and the like, depending on their intended use. Objects can be modified manually or programmatically by scripts. A script is a program written in the Linden Scripting Language (LSL). In order to have any effect, a script must be located in a prim's inventory and the prim must be instantiated ("rezzed") in the virtual world. Scripts can modify and move objects, sense and interact with avatars, and respond to and trigger various external events. For a more detailed introduction to *Second Life*, Linden Lab's recently published official guide book and the *Second Life* website are good places to start [54, 49].

3.3.2 Self-expression in Virtual Worlds

The current generation of online virtual worlds, including *Second Life*, is largely social in nature – people enter these worlds in order to meet other people and build connections with them through shared experiences. As in the real world, social interactions in virtual worlds revolve around self-expression. Taking *Second Life* as a representative example of the state-of-the-art in this respect, a resident of *Second Life* can express herself via the appearance and name of her avatar, the information revealed in her avatar's profile (favorite places, preferences, etc.), her avatar's scripted or explicitly triggered actions (dancing, laughing, running, etc.), text chat on public channels (received only by those nearby in the virtual world), text chat on private channels (received by a user-determined list of people regardless of their location in the virtual world), and live voice chat using a headset. A typical encounter when meeting another person for the first time, especially someone new to *Second Life*, revolves around explanations of how names and appearances were chosen, elaborations of details in avatar profiles, and exhibitions of clothing or animations. Such encounters make it clear that people in virtual worlds are hungry for means of self-expression in order to build social connections with others.

A less explicit although arguably more compelling form of self-expression in *Second Life* is the ability to build objects, from necklaces to cars to castles, and imbue them with a wide range of behaviors. The skill level needed to do so, however, is on par with that needed to build compelling web sites. As such, this form of self-expression is limited to a small proportion of the total virtual world demographic. However, those who can build and script in *Second Life* can express themselves to a far wider audience than those who cannot. For example, attending custom-built virtual dance clubs to meet with other people is a popular activity in *Second Life*. A club's popularity and clientele are largely determined by the personality of the club's creator and how well the club reflects that personality.

Compared to the real world, self-expression in *Second Life* and other virtual worlds is limited; missing are rich sources of information taken for granted in the real world, such as scent, body language, and the telltale signs of daily wear and tear. It's not that these sources of information were forgotten, just that they are difficult to emulate in meaningful ways in the virtual world. For example, virtual wind causes virtual trees to sway, a virtual sun and moon rise and set periodically, and virtual clouds form and disperse in the *Second Life* virtual world, but there is no meaning or cause behind any of these phenomena and their effect on the virtual world is superficial at best. On the other hand, some forms of self-expression that don't normally exist in the real world were added to *Second Life* in order to overcome problems specific to virtual worlds. For example, when using the public chat chennels, the user's avatar mimics a typing motion while the user is composing the message to be sent. This serves to signal to other avatars that the user is actively engaged in the converstation and alleviates the social awkwardness associated with an avatar blankly standing there while the user types a message.

Overall, the demand for richer forms of self-expression in *Second Life* and other virtual worlds is apparent. Data collected from real-world sensor networks can help meet this demand by importing into the virtual world the inherent expressiveness of the real world.

3.3.3 The Vacancy Problem

The **vacancy problem** is the noticeable and profound absence of a person from one world, either real or virtual, while they are participating in the other. Simply put, the vacancy problem arises because people do not currently have the means to be in more than one place (reality) at a time. In the real world, the vacancy problem takes the form of people appearing completely absorbed in their virtual reality, to the exclusion of everything in the real world. In the virtual world, the vacancy problem takes the form of virtual metropolises appearing nearly empty because there are not enough avatars to fill them. In part, this virtual vacancy is due to technical barriers preventing large numbers (hundreds) of people from interacting within the same virtual space. However, the vacancy problem will remain even if processor speeds, network bandwidth, and graphics fidelity increase to overcome these technical difficulties. In a world nearly unconstrained by geography or the laws of physics, the currency of choice is people rather than real estate or possessions.

To give a sense of the scale of the vacancy problem, consider the *Second Life*'s population density. *Second Life* currently has approximately 3300 square kilometers of virtual land and approximately 40,000 concurrent users at any given time [55]. Thus, at any given time the population density of *Second Life* is approximately twelve avatars per square kilometer. In comparison, Manhattan has an area of 87.5 square kilometers, a population of approximately 1.6 million people, and therefore a population density of approximately 18285 people per square kilometer [56].

The vacancy problem is a fundamental characteristic of the current generation of virtual worlds. More closely linking the real world with the virtual world, as the dual reality concept suggests, can mitigate the vacancy problem – just as real cities require special infrastructure to allow for a high population density, so too will virtual cities.

3.3.4 Mapping Between Realities

Aside from the challenges implied by instrumenting the real world with sensor/actuator networks to the extent necessary for realizing a dual reality, there are numerous challenges in designing exactly how the two realities interact and map onto each other. A direct mapping of the real to virtual and virtual to real may not be the most appropriate for all scenarios. For example, the sensor data streams collected from a real person may be better mapped to the virtual land the person's avatar owns rather than to the avatar itself. See Figure 3-5.

One possible mapping strategy is to shape the virtual world according to our subjective

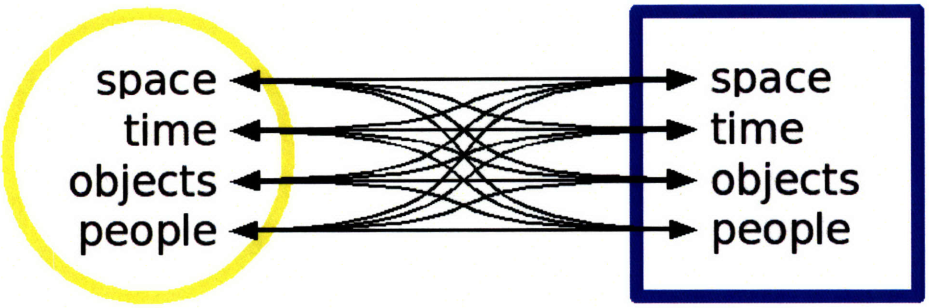

Figure 3-5: Fundamental mappings between the real (left) and the virtual (right).

perceptions of the real world. In essence, the virtual world would be a reflection of the real world distorted to match our mind's eye impressions as discerned by a network of sensors. For example, the buildings on a virtual campus could change in size according to the number of inhabitants and virtual corridors could widen or lengthen according to their actual throughput. Figure 3-6 depicts two classic examples of reality distortion due to perception. The sensory homunculus on the left represents a man whose body parts are sized proportional to the amount of somatosensory cortex devoted to them. On the right, Saul Steinberg's *View of the World from 9th Avenue* on the cover of *The New Yorker* exemplifies how the familiar and immediate occupy more mind share than the unfamiliar and remote [57].

3.3.5 Related Work

Work that couples the real world with virtual worlds falls into several broad categories. There are several efforts to bring a virtual world into the real world by using positioning and proximity systems to cast real people as the actors of an otherwise virtual world, such as Human Pacman [58], Pac Manhattan [59], ARQuake [60], and DynaDOOM [61]. Such work remains almost exclusively within the realm of converting video games into live action games and, aside from location awareness, does not incorporate other sensing modalities. Magerkurth et al. provide a good overview of this genre of pervasive games, as well as other more sensor-rich but physically confined games [62]. In an attempt to make *Second Life* more pervasive in the real world, Comverse has created a limited *Second Life* interface

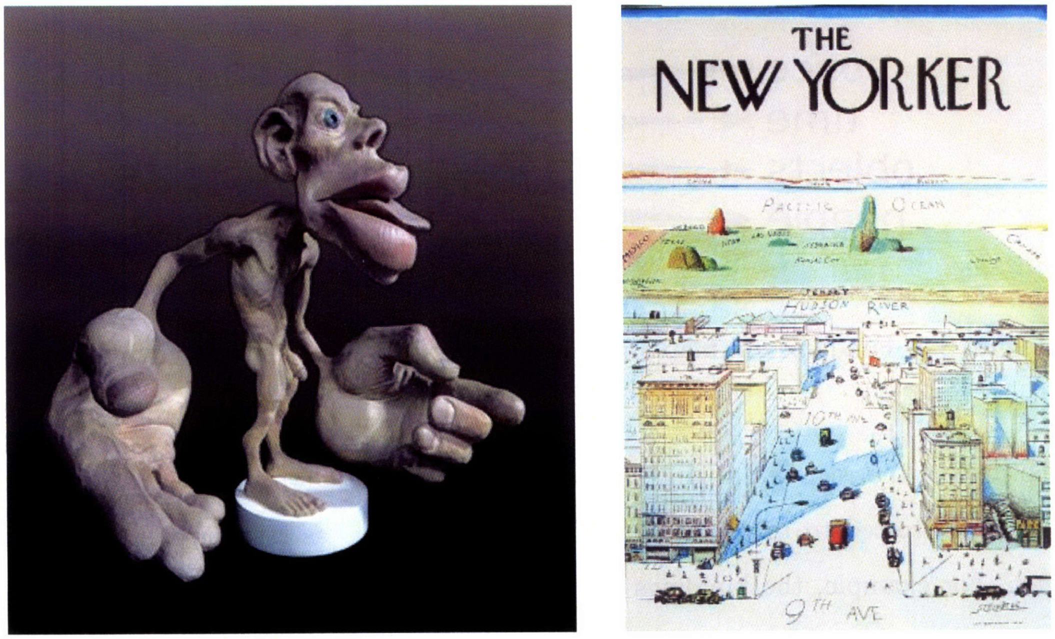

Figure 3-6: Two examples of realities distorted by perception: a sensory homunculus (left) and *View of the World from 9th Avenue* (right).

for cell phones [63]. Virtual worlds are being used to involve citizens in the collaborative planning of real urban areas [64], although this type of system relies more on GIS data than sensor networks embedded in the environment. More advanced and correspondingly more expensive systems are used for military training [65]. Most of the systems mentioned above support only a handful of simultaneous users.

Among efforts to bring the real world into the virtual world, it is standard practice to stream audio and video from live real events, such as conferences and concerts, into *Second Life* spaces built specifically for those events [66]. More ambitious and not as readily supported by existing technologies is the IBM UK Laboratories initiative in which the state of light switches, motorized blinds, the building's electricity meter, and the like in a real lab space are directly reflected and can be controlled in a *Second Life* replication of the lab space [67]. Similar efforts on a smaller scale include a general-purpose control panel that can be manipulated from both the real world and *Second Life* [68], and a homebrewed virtual reality wearable computer made specifically to interface to *Second Life* [69].

The convergence of *Second Life*, or something like it, with popular real-world mapping software to form a "Second Earth" has been broadly predicted [70]. Uses of such a "hyper reality" include analyzing real-world data ("reality mining"), as was done in the Economic Weather Map project [71]. Such ideas have appeared before as interactive art pieces. For example, the Mixed Realities juried art competition organized by Turbulence (a net art commissioning organization [72]) in collaboration with Ars Virtua (a media center and gallery within *Second Life* [73]) recently announced the five winning proposals, all of which mix various aspects of the real and virtual [74]. Sensor network-enabled dual realities may naturally merge with or evolve from the life logging work pioneered by Gordon Bell [75, 76] and popularized by web applications such as MySpace [77], Facebook [78], and Twitter [79].

3.4 An Emerging Medium from Merging Technologies

Many popular theories claim in one form or another that it is the widespread exchange of symbolic information that characterizes us as human [80, 81, 82]. In other words, communications media largely define us, at least on a societal level. Similarly, how we each relate to various media reflects our individuality. Our relation to media, in turn, is informed predominantly by the technologies employed. Within the paradigm introduced in Chapter 1, various technologies have fundamentally altered our capacity to *consume*, *share*, and *create* media. Most notably, television and radio made consumption widespread and the Internet made sharing widespread. In comparison, creation of media is still difficult and limited to a small subset of the population. The primary motivation of this work is to develop, in steps however small, the enabling technology for a medium as easy to create as to share and consume. To this end, sensor networks are employed as a generative tool in the process of transforming our everyday experiences in the real world into content shared and experienced in the virtual world. Just as the data created by a movie camera are shared and consumed in a theater, the data collected from sensor networks will be shared and consumed in virtual worlds.

Although sensor data may be used for creative purposes, it's not at all clear how the transformation from sensor data to self-expression will take place; the ability to generate and record sound says nothing of music. As such, sensor networks provide only the palette from which to work. A similar point was made explicit in the I/O Brush project, where the colors and motions of the physical world are literally captured and used as the palette for a digital brush upon a large touch sensitive display [83]. If data from sensor networks compose the palette, then virtual worlds are the canvas upon which self-expression takes form. The reverse relationship, in which data gathered from virtual worlds forms the palette and actuators spread throughout the real world form the canvas, will also materialize. As with all mass media, dual reality will inevitably be defined both by the average of its entire demographic and the virtuosity of its most gifted artists.

Chapter 4

Design and Implementation

> *"Knowing is not enough, you must apply; willing is not enough, you must do."*
>
> *Bruce Lee*

Most of the work carried out over the course of this thesis qualifies as both infrastructure and prototype. Infrastructure in that gaps in the technology necessary to demonstrate the dual reality concept were filled in, and prototype in that the method of implementation or the implementation itself experiments with some new idea or underlying principle. The resulting whole is a concrete example of dual reality. Magnifying any aspect of the whole reveals the application of various experimental supporting ideas and principles.

This chapter details the what, how, and why of the implementation portion of this thesis. The chapter is split into three roughly non-overlapping sections – one devoted to work implemented primarily in the real world, one devoted to work implemented primarily in the virtual world, and one devoted to work used to connect the two.

4.1 Real World Design and Implementation

Instrumenting the real world with a sensor/actuator network is a necessary step toward enabling a dual reality. Just as microcontrollers can be found in nearly every type of electronic device manufactured today, sensors will soon be embedded in every part of our environment. Many living and working environments could be considered sensor-rich by virtue only of the microphones, accelerometers, and cameras integrated into most personal computers and cell phones today, let alone the various security, safety, and access systems integral to many modern buildings. The problem with using these pre-existing sensor systems lies in their isolation from each other; they were for the most part designed to function as closed systems and retrofitting them with networking capabilities is highly impractical. It is in part for these reasons that a newly designed sensor node was used as the basis for realizing the real-world aspects of dual reality in this work. More importantly, though, was the opportunity afforded by a new design to experiment with a wider variety of sensors, form factors, and usage scenarios. The result is the Plug sensor node mentioned in Section 3.1.

Of course, instrumenting the world with sensors does nothing toward interfacing with those sensors and understanding their output. Representing and manipulating sensor data in a virtual world may offer some remedy to this problem, but an interface situated amidst the sensor network itself may be more appropriate for many usage scenarios, such as debugging or location-aware applications. A handheld sensor network browser called Tricorder was developed for this purpose. The remainder of this section describes in detail both the Plug and Tricorder projects, as well as their several extensions and offshoots.

4.1.1 Plug Sensor Network

Plug Hardware

The Plug[1] platform augments the utility of a standard power strip with sensing, communication, and computational abilities to affect a sensor node for active use in domestic and occupational settings while at the same time forming the backbone of a ubiquitous computing environment [84]. To these ends, each Plug offers four standard US electrical outlets supplying 120VAC at 60Hz. High-turn-ratio transformers sense the current drawn from

---

[1] The Plug and Lug circuit and mechanical designs were primarily carried out by Mark Feldmeier and the radio daughter card was designed by Mat Laibowitz, both of the MIT Media Lab's Responsive Environments Group.

each outlet and triacs allow power to be quickly switched on or off on each outlet. A varistor provides protection against electrical surges. The four current sensors and four switches are monitored and controlled by an Atmel AT91SAM7S64, a peripheral-rich microprocessor based on the 32-bit ARM7 core running at 48MHz with 16KB of SRAM and 64KB of internal flash memory. The same microprocessor controls all other aspects of the Plug as well, including two LEDs, a push button, a small speaker, a piezoelectric cantilever vibration sensor, a microphone, a phototransistor, a 2.4GHz ChipCon CC2500 wireless transceiver, a voltage sensor, a USB 2.0 port, and an extensive expansion port for adding custom hardware to the Plug. The voltage sensor has a dynamic range of ±280V relative to the neutral line and the current sensors have a dynamic range of ±4.1A, but can withstand up to 30A. An analog volume knob ensures that the speaker can be manual deactivated. Figure 4-1 shows a single Plug node. Appendix A details more fully the Plug hardware design. Figure 4-2 shows some of the 35 Plugs built to date.

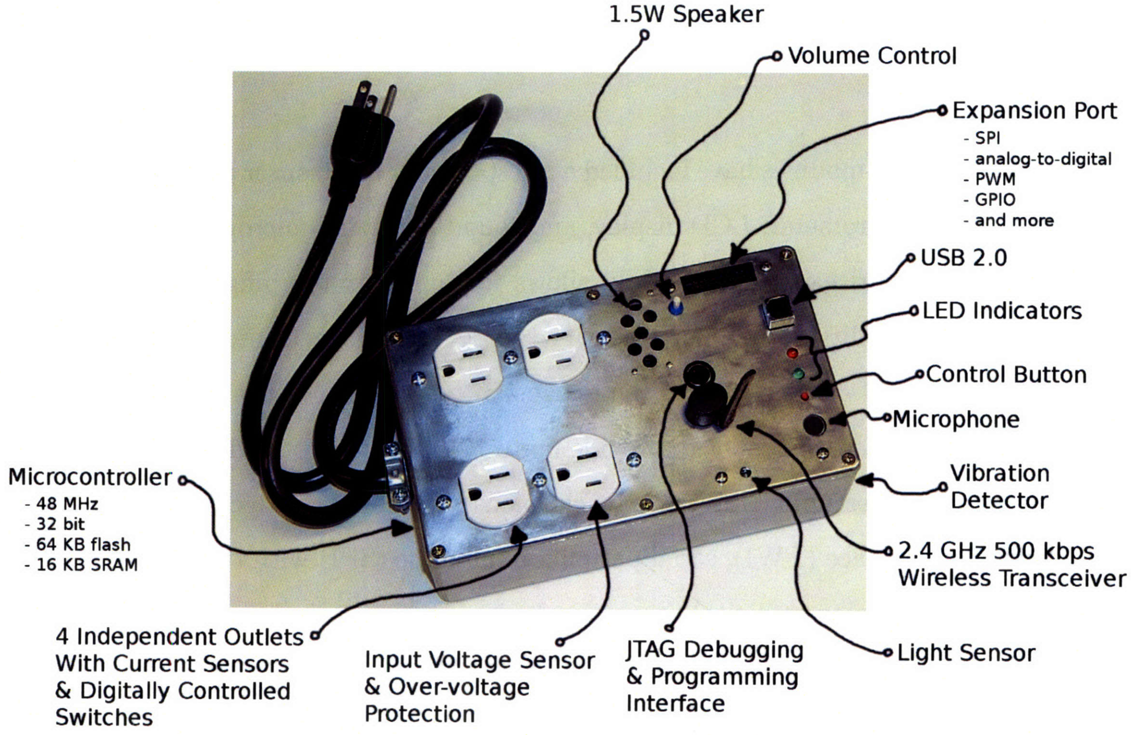

Figure 4-1: A Plug sensor node.

At present, the 20-pin expansion port houses a module, referred to as the data log expansion

Figure 4-2: Plug sensor nodes in preparation for being programmed.

module, with a passive infrared (PIR) sensor for detecting motion, an SD memory card slot for removable data logging, a 4Mbit external flash memory for storing persistent state (e.g., calibration constants and unique identifiers), and a digital 13-bit temperature sensor [85][2]. See Figure 4-3. Unless otherwise noted, all Plugs are assumed to be equipped with the data log expansion module.

Additional expansion modules have included a full spectrum light sensor array, inflatable privacy indicator, accelerometer, LCD display, and sound localizing microphone array, among others. A breadboard expansion module allows for quick prototyping. Another imbues the Plug with the ability to uniquely identify the devices Plugged into it using an RFID reader[3]. See Figure 4-4. The pins of the expansion port can be variously multiplexed by the Plug's internal microcontroller to provide up to 17 general purpose digital input/output (GPIO) pins with interrupts, Universal Synchronous Asynchronous Receive Transmit (USART), two wire interface (TWI), serial peripheral interface (SPI) with two chip selects, one fast interrupt, one analog to digital converter channel, and electrical ground. Of the GPIO pins, three can continuously source up to 16mA and the others can source up to 8mA, so long as the total current sourced is less than 150mA. Of course, an expansion module can also plug in directly to one of the Plug's electrical outlets if it needs more energy.

---

[2] The data log expansion module circuit and mechanical designs and implementations were primarily carried out by Cameron Lewis, also of the MIT Media Lab's Responsive Environments Group.

[3] The RFID reader expansion module was designed and built by Yasuhiro Ono, a researcher in the Responsive Environments Group visiting from Ricoh.

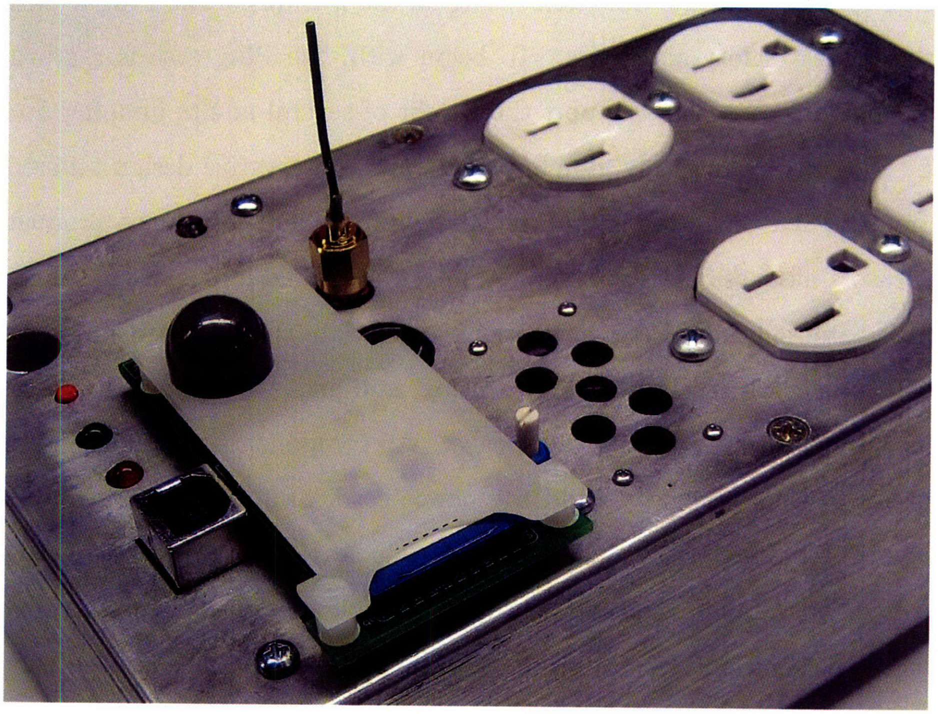

Figure 4-3: A Plug node equipped with the data log expansion module. Visible are the grey dome of the motion sensor on top and the removable SD memory card on the front side.

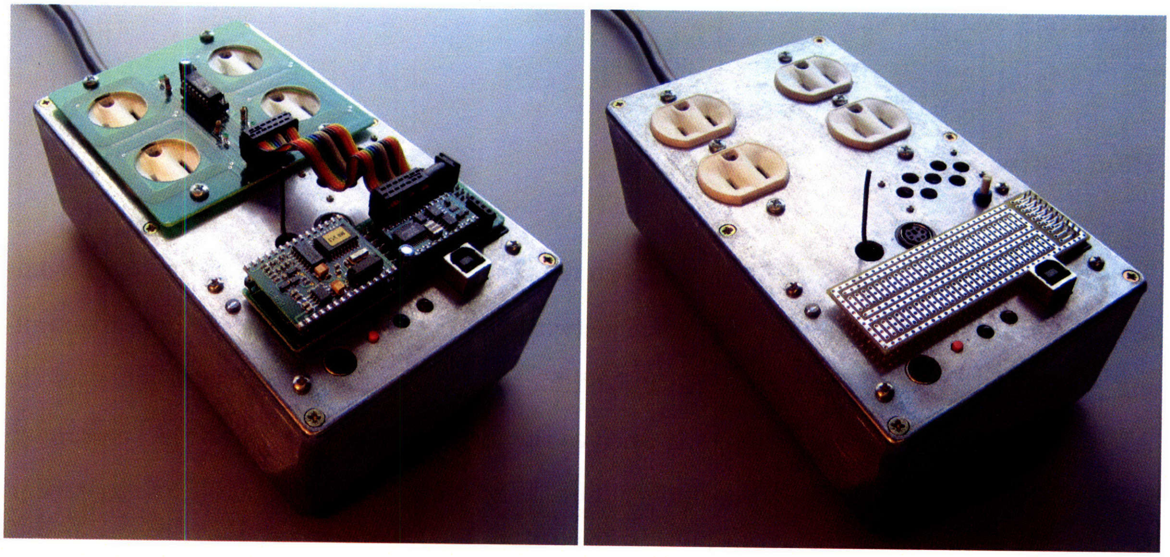

Figure 4-4: A Plug equipped with a RFID reader for each of its four outlets (left) and a breadboard expansion module for prototyping (right).

Significant effort was put into accommodating all the features of the Plug while still maintaining the highest safety standards. To begin with, the Plug case is a sturdy die-cast aluminum box that can easily support the weight of several adults jumping up and down on it. The Plug node's ungainly aluminum case is a conservative design tailored for rapid construction and safe operation within the lab – a proper design for mass production would have the sensors more seamlessly integrated into what would look more like a conventional power strip. Internally, the Plug comprises two separate circuit boards, one that handles all high voltage signals, such as voltage sensing, and another that handles only low voltage digital signals, such as driving the speaker. All components protruding from the case, except for the outlets themselves and the power cord, have connections only to the low voltage board. A transformer on the high voltage board supplies up to 500mA at 3.3V to the low voltage board and expansion port. No high voltage signals are externally accessible through the expansion port or otherwise. The apertures in the case are precision cut by a waterjet cutter to ensure a tight fit around all protruding components. As is standard with conductive enclosures, the entire Plug case is grounded. The total current sourced by a Plug is limited by a slow-blow 8A fuse, which precludes using high current appliances such as heaters. As well as being a safety precaution, the fuse also protects the triac switches from being over driven.

Plug Software

The firmware running on the Plug microcontroller was written from scratch, with some aspects of its design, especially regarding networking, derived from the Pushpin Computing platform [86]. The firmware is optimized for and makes extensive use of many of the AT91SAM7S64's hardware peripherals, such as the direct memory access (DMA) controller, pulse width modulator (PWM) controller, and serial peripheral interface (SPI) controller. In particular, the firmware includes interrupt-driven, double-buffered modules for interfacing to and managing the SPI bus, GPIO pins, radio, USB port, ADC, speaker, LEDs, push button, real-time clock with settable alarms, SD memory card, vibration and motion sensors, random number generator, data flash, and temperature sensor. See Figure 4-5 for an

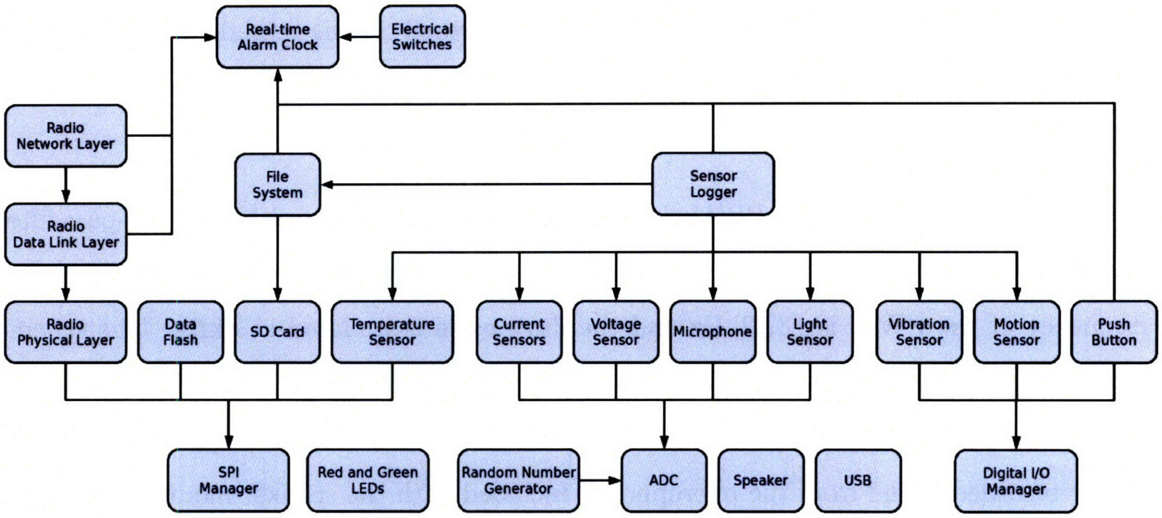

Figure 4-5: A dependency map of the Plug firmware. Arrows from A to B indicate that A depends on B.

overview of the firmware structure. Applications, set at compile time, are pieced together from these modules predominantly through the use of asynchronous callbacks.

Since the Plug platform is essentially free of energy constraints (rate of consumption is limited, but not cumulative consumption), the Plug software modules are designed to be time efficient, not energy efficient. For example, the processor never goes into a sleep mode, although several such modes are available to it. Even so, the Plug only draws approximately 60mA at 3.3V when all subsystems are active.

Each firmware module in Figure 4-5 is described below. More firmware details are given in Appendix B. A side effort to create a high-level interpreted language for sensor networks based on Python is detailed in Appendix C.

Red and Green LEDs: Each LED is independently controlled by a dedicated 16-bit PWM channel updated at 45.7Hz. Patterns can be played by passing to the LED module a function that is called every PWM period. Thus, patterns can be generated algorithmically or read directly from memory. Patterns can also be repeated and interrupted. An interrupt-driven LED manager updates the PWM duty cycle every period until the pattern stops or the LED is turned completely on or off. Included with the LED module are patterns for blinking, pulsing, and fading the LEDs.

Speaker: A dedicated 8-bit PWM channel feeds into a two-pole passive LC low-pass filter tuned to a 20kHz 3dB point to produce the output waveform played by the speaker. The frequency of the PWM is 187.2kHz and the duty cycle is updated at 8kHz. The scheme for playing sounds is nearly identical to that of playing light patterns on the LEDs. The speaker can play single-channel, 8-bit, 8kHz WAV files from the file system module or replay recorded sound from the microphone. Included with the speaker module are sounds for white noise and a 200Hz tone.

USB: The microcontroller acts as a USB 2.0 device operating in bulk transfer mode. Other transfer modes are possible, but not implemented. The microcontroller cannot serve as a USB host. Data transfer to and from the host is interrupt-driven. The data rate is largely dependent on the overhead incurred by the host for making a transfer. For example, in a typical set up, the fastest the host can query the device is once every 2 milliseconds. Thus, due to the high speed of USB 2.0, there is no difference in the time it takes for the host to receive 1 byte or 1000 bytes since both transfers complete within a 2ms window.

Real-time Alarm Clock: A 16-bit hardware timer forms the basis of a 64-bit clock reset to zero upon start up and incremented once every 2.67 microseconds thereafter. The current time can be safely retrieved at any point. An arbitrary number of alarms can be set to trigger arbitrary callback functions. Alarms are guaranteed to trigger only after the specified time and a best effort is made to trigger them as soon after the specified time as possible, depending on interrupt latencies. Alarms can also be unset without causing them to trigger.

Electrical Switches: Each of the four electrical outlets can be independently turned on or off by means of a triac controlled by a digital I/O pin. One of the four switches can be controlled by an interrupt-driven software PWM. The same functionality could be easily extended to the other three outlets, but is not yet implemented. All electrical outlets are on by default after start up or processor reset.

Digital I/O Manager: Any firmware requiring notification of digital input pin change interrupts, such as the push button, receives such notification by first registering with the digital I/O manager module. The digital I/O manager also facilitates the use of pins as outputs.

Vibration Sensor: The vibration sensor interfaces with the microcontroller through a single digital input pin that can either be polled directly or configured to trigger a pin change interrupt. When vibration is sensed, the input pin is driven low, referred to as the active state. The duration of the active state indicates either continuous vibration or the amplitude of a single impulse.

Motion Sensor: The passive infrared (PIR) motion sensor detects change in the heat incident upon a wide-angle plastic lens and can easily detect the motion of a person within a couple meter radius. Like the vibration sensor, the duration of the active state (also low) indicates either continuously changing incident heat or, more probably, the amplitude of the change. The motion sensor can be polled directly or configured to trigger a pin change interrupt.

Push Button: Configurable function callbacks are triggered whenever the push button is depressed or released. Both events are automatically debounced with a 50 millisecond interrupt-driven wait time. Depressing and releasing the button five times within two seconds will cause the Plug to reset.

ADC, Current Sensors, Voltage Sensor, Light Sensor, and Microphone: For simplicity and flexibility, the default behavior of the analog-to-digital converter (ADC) is to continuously sample the light sensor, voltage sensor, microphone, and four current sensors at 8kHz per channel with 8-bit resolution. The ADC interfaces to the direct memory access (DMA) controller to automatically store samples in a memory buffer large enough to hold 64 samples per channel. Once the first buffer is filled, subsequent samples are automatically placed in a second buffer so as not to overwrite previously taken samples while they are being processed by a callback function, after which the roles of the two buffers are interchanged and the process repeated. The default callback function for processing ADC data is part of the sensor logger module. The default ADC configuration can be overridden to have, for example, a single ADC channel sample at as high as 191kHz and with 10-bit resolution, such as might be required to distinguish different kinds of fluorescent light ballasts [87].

Random Number Generator: The random number generator module provides pseudo-random numbers in a variety of formats by using the standard C library random number generator with an initial seed generated bit by bit by XORing the lowest bit of many ADC samples.

SPI Manager: The microcontroller's serial peripheral interface (SPI) speaks to the radio, temperature sensor, SD card, and data flash, but only one at a time. The SPI manager module coordinates the use of the SPI bus among these four peripherals or others, such as an LCD display, depending on the expansion board used. All SPI communication is interrupt-driven and takes advantage of the DMA controller so entire memory buffers can be received or transmitted without processor intervention.

Data Flash: A 4Mbit data flash integrated circuit resides on the data log expansion board. Only rudimentary tests have been implemented since the SD card fulfills much of the same functionality.

Temperature Sensor: The 13-bit temperature sensor on the data log expansion board can report temperatures between -40°C and 150°C with a resolution of 0.03125°C and typical accuracy of ±0.5°C at a sample rate of up to once every second.

SD Card: The data log expansion board holds a 1GB removable secure digital (SD) flash memory card. The SD card module provides interrupt-driven reading and writing of whole and partial 512-byte sectors over the SPI bus.

File System: The entire 1GB SD card is occupied by a customized file system suited to the specific needs of the Plug. Table 4.1 gives an overview of the file system used. In particular, the root of the file system is the first sector (512 bytes) and contains a 16-bit unique identifier (UID) and a table indicating the occupancy status of the dynamic WAV file table located in the second quarter of the SD card's memory. The first quarter of the SD card's memory except the first sector contains the static WAV file table, comprising an arbitrary number of WAV files, each no greater than 4096 sectors in size. The static WAV file table is set at compile time. The dynamic WAV table occupies the second quarter of the SD card's memory, comprising 128 slots each of 4096 sectors. The dynamic WAV file table can be added to or erased during run time. Each WAV file slot, whether in the static or dynamic table, can hold a single 8-bit, 8kHz, single-channel WAV sound file up to approximately 4 minutes and 22 seconds in length. Half of the dynamic WAV file slots are designated as "outgoing" and half as "incoming". Outgoing WAV files are those originating from the microphone or from a neighboring Plug meant to be transferred over the network to an external location for playback. Incoming WAV files are those received over the network meant for local playback on the speaker. The last half of the SD card's memory is devoted to log file entries, each occupying one sector, generated by the sensor logger module.

Radio Physical Layer: The Plug's ChipCon CC2500 2.4GHz low-power radio interfaces to the microcontroller through an SPI bus and two digital inputs. The radio physical layer module uses a carrier sense multiple access (CSMA) scheme with random backoff upon collision detection to transmit packets and won't accept another packet for transmission until the current packet has been successfully transmitted. A successful transmission does not indicate a successful receipt. Received packets are guaranteed to have passed a 16-bit cyclic redundancy check (CRC) and come with a link quality indicator (LQI) and received signal strength indicator (RSSI). All radio physical layer module transactions are interrupt-

SD Card Placement	Name	Description
sector 0	root	16-bit UID and dynamic WAV file table occupancy list
sectors 1 to 495999	static WAV file table	8-bit, 8kHz, mono-channel WAV files up to 4 minutes and 22 seconds long set at compile time
sectors 496000 to 991999	dynamic WAV file table	up to 128 8-bit, 8kHz, mono-channel WAV files up to 4 minutes and 22 seconds long recordable during run time
sectors 992000 to 1983999	log file entries	over 47 days worth of log entries taken once every 4.096 seconds

Table 4.1: The layout of the file system used on the Plug's SD card.

driven. As a debugging and network diagnostic tool, a radio profiler records every state the radio passes through and the result of every transmitted or received packet.

Radio Data Link Layer: Using the radio physical layer as a basis, the radio data link layer module manages single-packet transfers between network-adjacent Plugs. Packets can be broadcast to all neighbors or addressed by means of a one-byte locally unique identifier to a single neighbor, in which case it can optionally be certified so as to require an acknowledgement of receipt. Multiple packets can be queued for transmission at the same time and an arbitrary callback function to be called upon transmission success or failure can be associated with each packet in the queue. The one-byte network identifiers can either be randomly generated and checked against neighboring identifiers for uniqueness or, due to the relatively small number of extant Plugs, derived directly from the Plug's two-byte UID. The radio data link layer maintains a list of up to sixteen neighbors, each entry of which contains the neighbor's UID, network identifier, most recent LQI and RSSI measurements, time of the most recently received packet from this neighbor, and a metric for gauging how many certified packets sent to the neighbor were responded to with an acknowledgement packet.

Radio Network Layer: The radio network layer module builds upon the radio data link layer module to provide a routing scheme for delivering packets across the Plug network from an arbitrary number of source nodes to a handful of sink nodes. The basic scheme involves two phases – one for discovering the route from a source to a sink and another for sending packets along that route. Route discovery involves flooding the network with a query that seeds each node with its minimum hop count from the data sink. In one variation of route discovery, only the identifier for the sink and the number of hops from the sink are recorded by each node. Sending a packet to the sink requires following the path (or paths in some cases) of decreasing hop count until the sink is reached using neighborhood broadcast data link layer packets. In a second variation of route discovery, only the address of the next node in the route and the number of hops from the sink are recorded by each node. Sending a packet to the sink requires tracing back the unique path from the source using certified data link layer packets. In both variations, any set of nodes can be dynamically designated as the sinks, the only limitation being the total number of sinks due to memory constraints. The current implementation has a limit of five sinks.

Sensor Logger: All sensor data from each of the ADC channels (light, sound, microphone, voltage, and current) are processed by the sensor logger module to calculate statistics for each sensing modality and periodically record a log of these statistics and other Plug state to the SD card through the file system module. The sensor logger module processes raw sensor samples in batches of 256 samples per sensor channel by calculating the minimum, maximum, and average for the batch for each channel. The most recent batch of 256 raw samples for each sensor channel is always available through a callback function triggered when the batch is finished being sampled. Similarly, the most recent 256 minima, maxima, and averages for each channel are available as well. The completion of every sixteenth set of minima, maxima, and averages triggers the further aggregation of the previous 32 minima, maxima, and averages into a single minimum, maximum, and average for each channel. Eight such aggregates are recorded in each log entry of the file system on the SD card. That is, each log entry on the SD card contains, for each sensor channel, eight minima, maxima, and averages windowed over 1.024 seconds such that consecutive windows overlap by 0.512 seconds. Thus, a log entry is written to the SD card every 4.096 seconds. For each of the eight 1.024-second windows, the number of times the motion and vibration sensors were active when sampled once every 8 milliseconds is also recorded. In addition to these statistics, each log entry also contains a single temperature sensor reading, the number of times the button has been pressed since the last log entry, the radio profiler log maintained by the radio physical layer module, and the list of neighbors maintained by the radio data link layer module.

4.1.2 Lug Sensor Node

The Lug is sensor node derived from the same microprocessor, radio, and code base as used in the Plug sensor nodes. Its small form factor and fast prototyping amenities place the Lug closer to a more traditional sensor network node and complement the fixed network of Plug nodes. That is, the Plugs act as the always-on backbone network among which Lugs are put to use as wearable or battery powered sensor nodes. See Figure 4-6.

The Lug exposes all of the Atmel AT91SAM7S64's analog and digital pins in a through-hole fashion appropriate for standard headers. The same radio board found on the Plug can be attached either to the back of the Lug or as a protruding extension. The Lug is USB 2.0-ready and, by populating it with a 5VDC to 3.3VDC regulator, can draw power directly from a USB host. Appendix D details more fully the lug hardware.

The microprocessor and all subsystems support low power modes suitable for such battery powered applications, although that has not been the focus of the Lug's use thus far. The Lug's primary use is prototyping on-body or otherwise mobile sensor nodes that interact with the surrounding fixed Plug network. Section 4.1.3 details the most prominent use of a Lug.

4.1.3 Tricorder Handheld Sensor Network Browser

The Tricorder is a location-, orientation-, and network-aware handheld device used to interface in real time to the Plug wireless sensor network [88]. As the name suggests, the

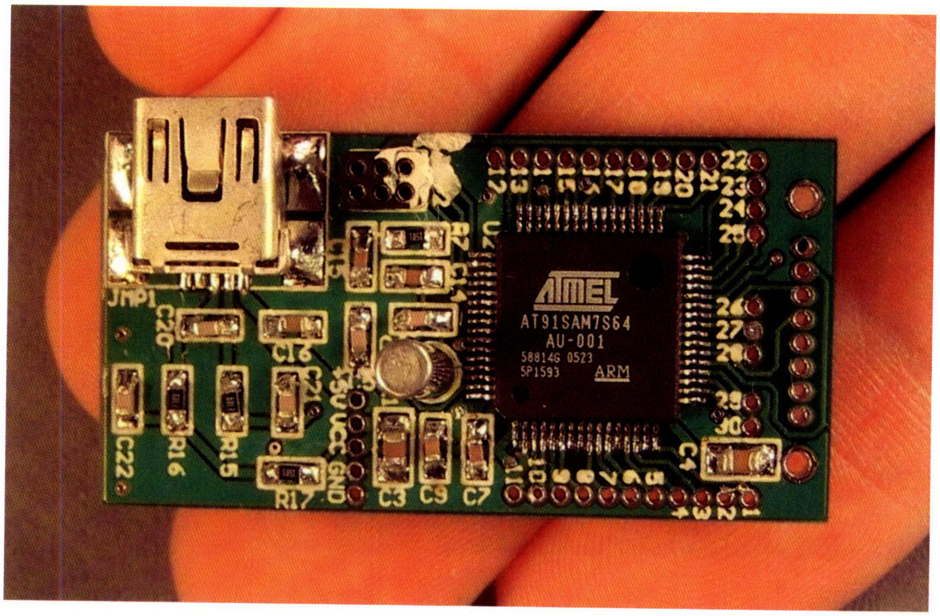

Figure 4-6: A Lug node without radio or battery.

inspiration for the Tricorder comes from the fictional device of the same name from the original Star Trek science fiction television series. The fictional version of the Tricorder was a self-contained device capable of sensing relevant information about whatever it was being pointed at (e.g., life signs 50 meters back, magnetic disturbance above, or plot thickener ahead). The Tricorder device aims to achieve the same goals, but by pulling sensor data off the surrounding wireless sensor network rather than being self-contained. Physically, the Tricorder comprises a Nokia 770 Internet Tablet for display and user input purposes [89], a Lug sensor node for communicating with the surrounding Plug network, a 3-axis compass with electronic gimbaling for ascertaining absolute orientation in three dimensions with up to ±80° tilt [90], a battery pack power supply, and a plastic case to hold it all together. See Figure 4-7. Like the fictional Tricorder, this version knows in which direction it is pointing thanks to the compass. This, combined with coarse localization based on the Lug radio's received signal strength indication (RSSI) from nearby Plugs, allows for real-time point-and-browse functionality while physically roaming within the sensor network itself. The Tricorder's Plug-compatible radio can query nearby Plug sensor nodes directly or distant Plug sensor nodes by means of spreading a multi-hop request through the network and

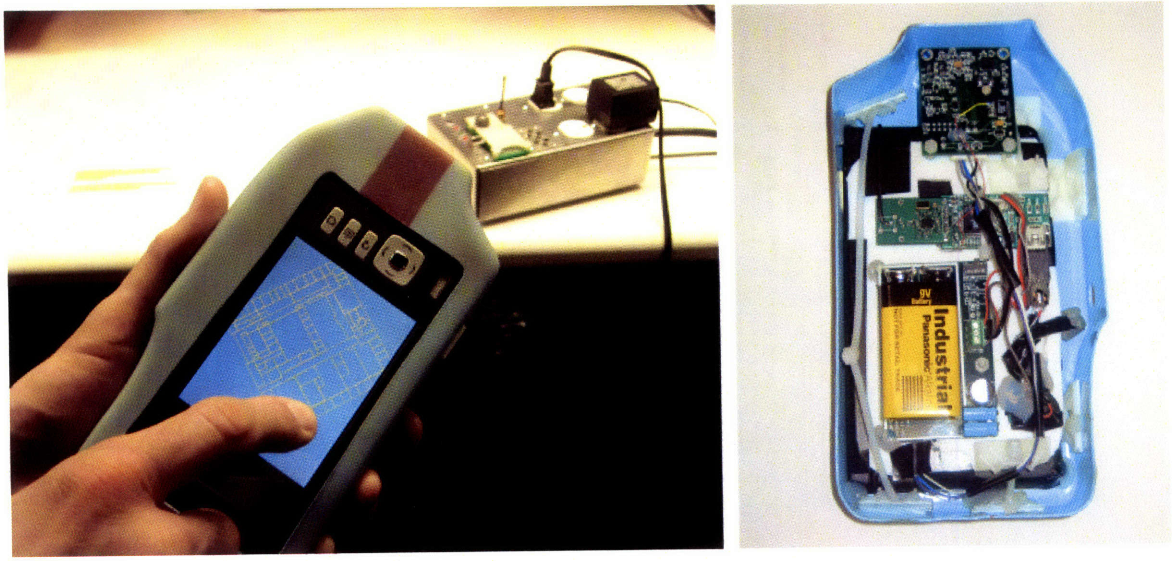

Figure 4-7: The Tricorder, inspired by Star Trek, affords a pointing interface to real-time, situated sensor data provided by the Plug sensor network nodes (left). The opened backside of the Tricorder device shows the battery pack, radio, compass, Nokia 770, and glue logic used to connect the radio to the Nokia via USB and compass to the radio via a USART (right).

waiting for a response.

Figure 4-8 shows a screenshot of the Tricorder's graphical user interface as implemented using the PyGame module [91] of the Python programming language under the Maemo development environment [92]. The interface shows the floorplan of the third floor of the MIT Media Lab overlaid with Plug icons depicting each Plug's most recent measurements of sound (blue concentric circles), light (red radial lines), RSSI (green central circle), current consumption (black central dial), and motion (orange ring around the green central circle). The icons jitter slightly to represent vibration. The map is centered and oriented on the Tricorder's touch screen in real time according to the RSSI location estimate and onboard compass direction reading. The user can pan the map using the touch screen and zoom in or out using hardware buttons. Touching an icon reveals more detailed information about the corresponding Plug sensor node. Specifically, the selected Plug icon is highlighted in a pink square and the minimum, maximum, and average windowed over the previous 4.096 seconds of each of its sensing modalities are displayed in the bar graphs in the side pane.

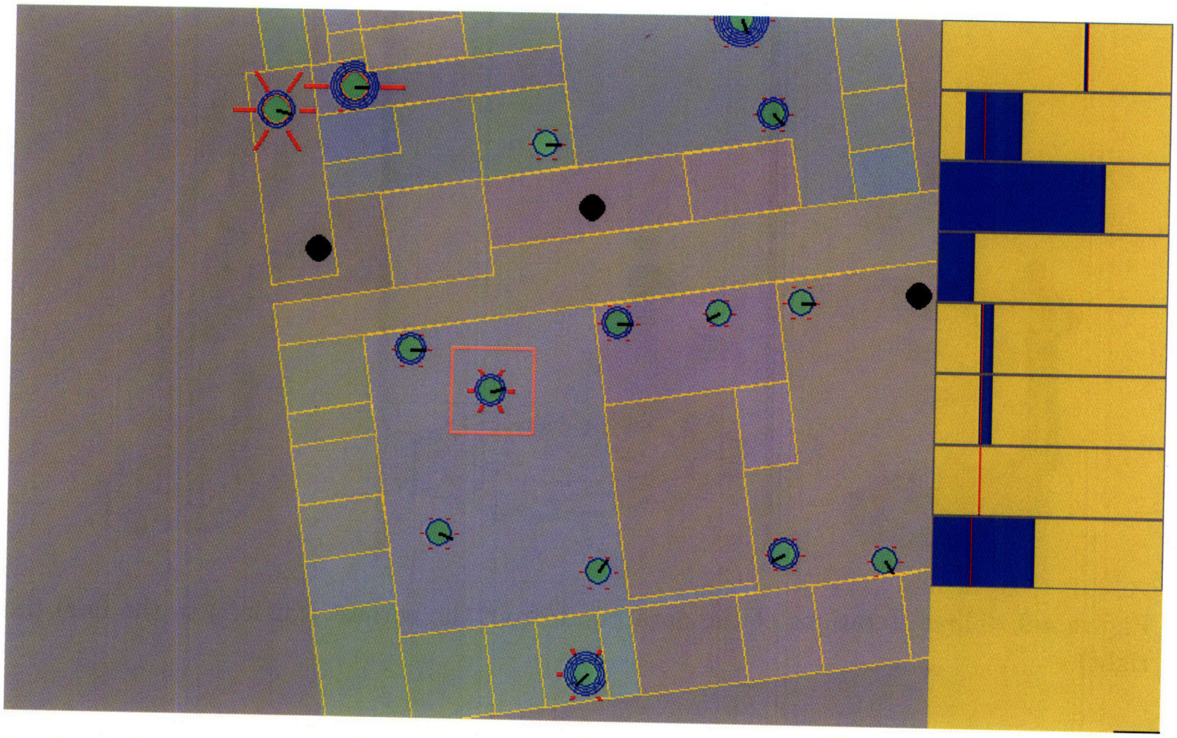

Figure 4-8: A screenshot of the Tricorder device.

The Plug so detailed can be selected automatically according to strongest radio signal or manually via the touch screen.

The Plugs' multi-hop network allows for sensor data to be streamed even from far away nodes. These remote sensor streams, when combined with the orientation measurement from the compass, in a sense grant the Tricorder the power to "see" through whatever walls or other objects obstruct its line of sight. In addition to communicating directly with the Plug sensor network, the Tricorder also has wireless Internet (IEEE 802.11) and Bluetooth capabilities, opening the possibility of accessing more traditional databases, websites, RSS feeds, cell phones, etc.

4.2 Virtual World Design and Implementation

Unless otherwise noted, the virtual world work described in this section exists entirely within the *Second Life* virtual world. Refer back to Section 3.3.1 for a description of the

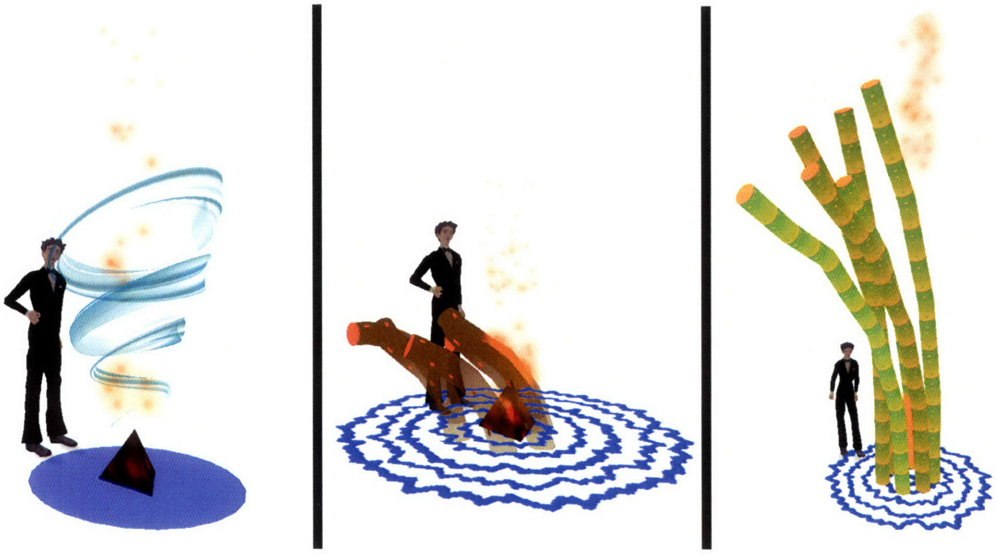

Figure 4-9: Three iterations of the data pond, from the earliest (far left) to the final (far right).

basic design and implementation tools available in *Second Life*. Appendix E contains an example of a Linden Scripting Language (LSL) script taken from Section 4.2.2.

4.2.1 Data Ponds

A single data pond is meant to be an easily distinguishable, locally confined representation of the sensor data from a single Plug node. The final iteration of the data pond design consists of a cluster of waving stalks growing out of a puddle of water and an ethereal foxfire rising from among the stalks, as might be found in a fantastic swamp. The mapping between a Plug's sensor data and its corresponding data pond is, of course, arbitrary. The final design iteration is meant to be as intuitive as possible given the abstract nature of the data pond. In this way, the mapping is easily understood once explained, but still interesting even without the benefit of the explanation. Table 4.2 describes the mapping used in the final design iteration. Figure 4-9 depicts the final data pond design as well as some of the earlier designs.

A real version of the data pond complements the virtual version. The real version follows

Plug Sensor Modality	**Data Pond Attribute**	**Mapping**
light	stalk skin pattern	the number of rings on the stalk is proportional to the maximum light level over the most recent one-second window
temperature	stalk color	the color of the stalks varies linearly from blue to yellow to red within the temperature range 18°C to 29°C
motion	stalk motion	the stalks sway gently when no motion is detected and excitedly when motion is detected over the most recent one-second window
sound	puddle size	the diameter of the water puddle is proportional to the maximum sound level over the most recent one-second window
electrical current	fire intensity	the height and intensity of the fire is proportional to the total average absolute value of the electrical current over the most recent one-second window

Table 4.2: The mapping from a real-world Plug's sensor data to its corresponding virtual data pond.

the virtual's tentacle aesthetic by using a standard desk fan shrouded in a lightweight, polka dotted sheet of plastic. The air flow through the shroud and therefore the height, sound, and other idiosyncrasies of the shroud can be finely controlled by plugging the fan into the outlet of a Plug device and pulse width modulating the supply voltage accordingly. See Figures 4-10 and 4-11.

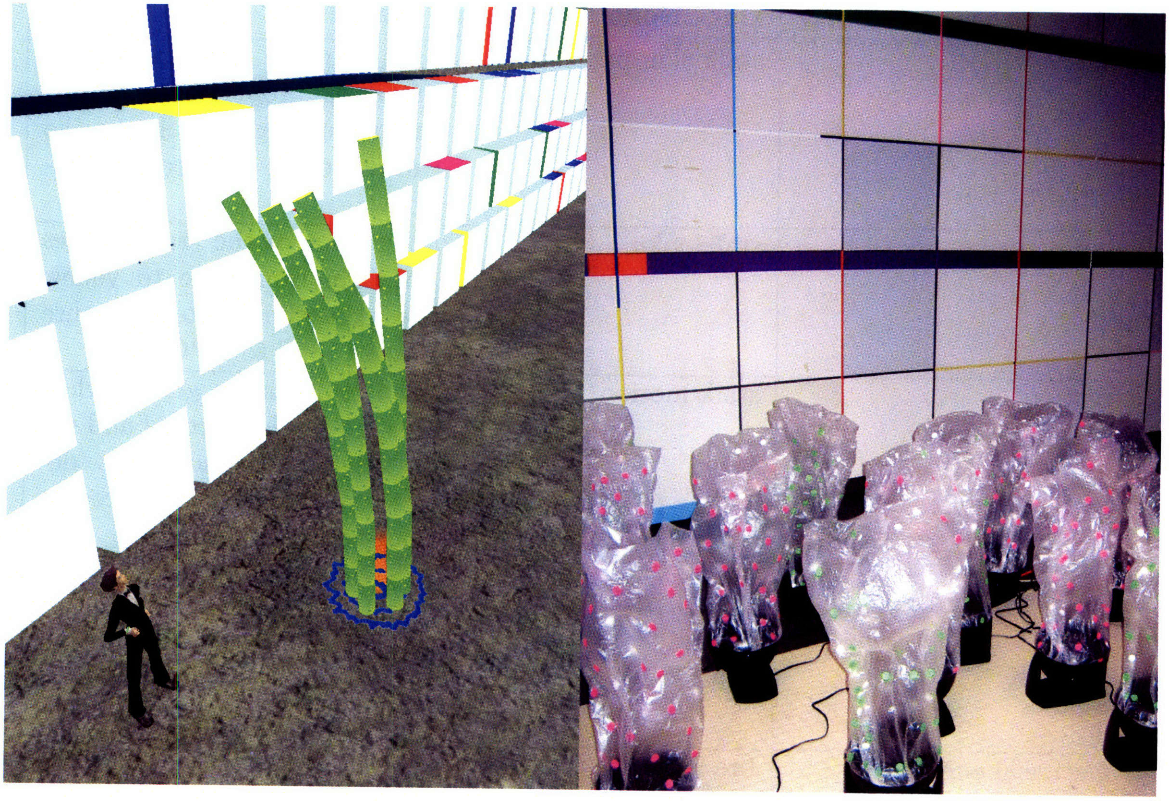

Figure 4-10: A virtual data pond displaying real data in front of a virtual atrium wall (left) and real data ponds displaying virtual data in front of a real atrium wall (right).

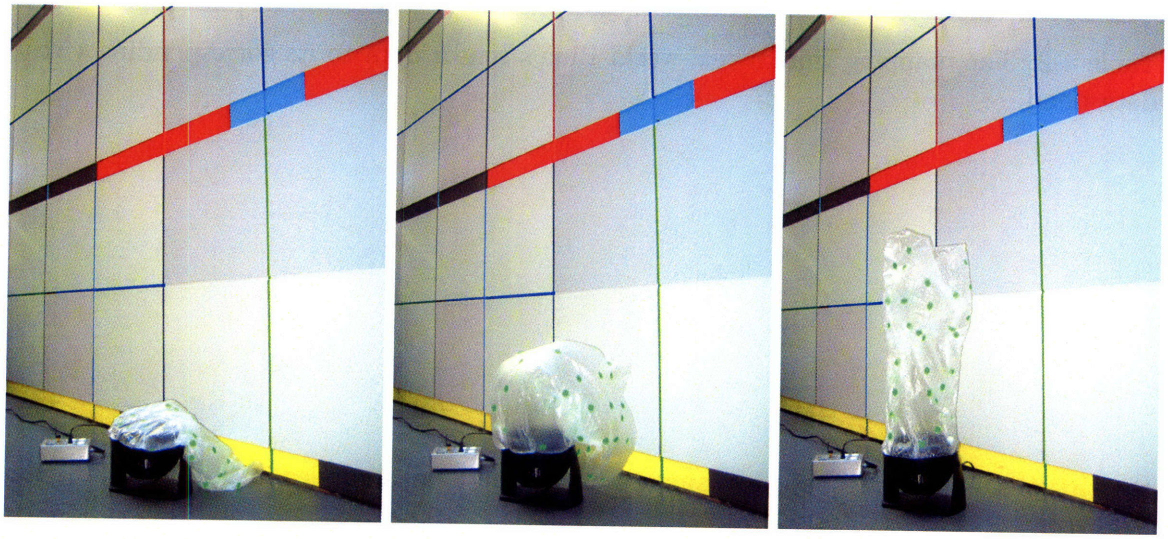

Figure 4-11: The real-world version of the data pond can display information according to how much air flow passes through the plastic shroud, which is controlled at a fine level of detail by the Plug used to power the data pond's fan.

Parameter Sensed	Example
toucher identity	Nix Division (18aebf4a-0461-4308-95fc-a8c9c0e09720)
object location	Ruthenium (43.149315, 34.921227, 23.875257)
object velocity	(0.000000, 0.000000, 0.000000)
object rotation	(0.000000, 0.000000, 0.000000, 1.000000)
Unix timestamp	1181920802.264127

Table 4.3: The parameters sensed by the embedded virtual sensing scheme.

4.2.2 Virtual Sensing

Just as real sensor networks monitor the real world, virtual sensor networks monitor the virtual world. The practical differences between the two are the type and level of detail of the phenomena sensed. For example, in the real world, there are literally an infinite number of ways a person can touch a table, whereas in *Second Life*, there is exactly one. On one hand, this discrepancy makes high-level contextual information in the virtual world readily available – in many cases, it is more useful to know a person is in the room than the motion sensor went off, a difficult task in the real world but trivial in the virtual. On the other hand, the nuance and level of detail afforded by sensing the real world is largely lost when sensing the virtual world. Put another way, the real world overflows with an abundance of information, most of which is not needed for a given sensing application, whereas the virtual world suffers a dearth of information in comparison, much of which is highly valuable for a given sensing application.

For the purposes of this work, two approaches are used to sense the virtual world, corresponding roughly to embedded and wearable sensing schemes. The embedded sensing scheme entails seeding every object of interest in the environment to be sensed with a script that detects when on avatar touches or otherwise interacts with the object and then reports back to a server external to *Second Life* with a full description of the interaction. Table 4.3 details the information gathered by the embedded sensing scheme.

The wearable sensing scheme requires each avatar in the region of interest to wear a sensing bracelet. See Figure 4-12. The sensing bracelet reports back to the same external server every five seconds with a full description of its avatar's location and motion. The bracelet

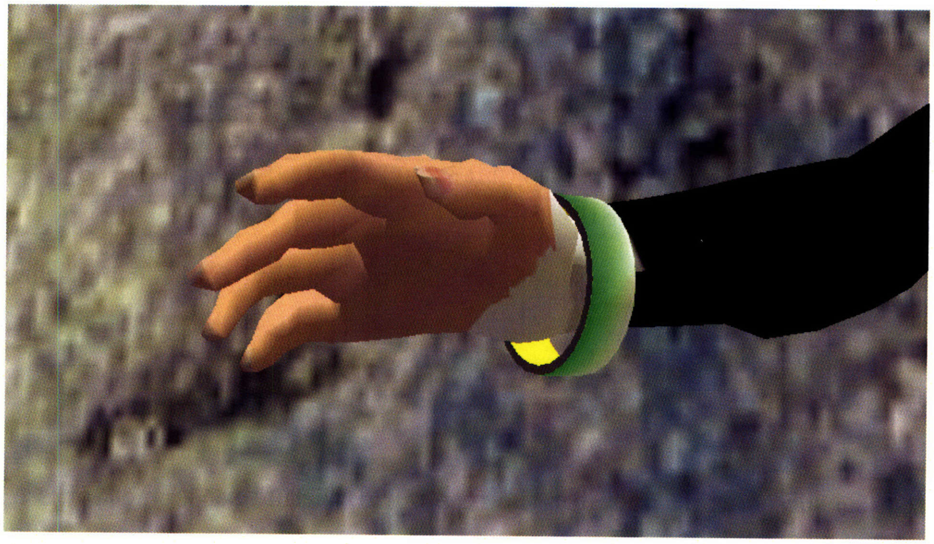

Figure 4-12: A wearable sensing bracelet worn by an avatar.

Parameter Sensed	**Example**
wearer chat messages	Hello world.
wearer identity	Nix Division (18aebf4a-0461-4308-95fc-a8c9c0e09720)
bracelet location	Ruthenium (43.149315, 34.921227, 23.875257)
bracelet velocity	(0.000000, 0.000000, 0.000000)
bracelet rotation	(0.000000, 0.000000, 0.000000, 1.000000)
bracelet identity	Access Bracelet (8e8b5063-aa0a-a302-82f3-27007db9138d)
Unix timestamp	1181920802.264127

Table 4.4: The parameters sensed by the wearable virtual sensing scheme.

also reports everything its avatar says over the public chat channel. Table 4.4 details the information gathered by the wearable sensing scheme. As incentive for avatars to wear the sensing bracelet, the bracelet also serves as an access token without which the avatar will be barred from entering the region being sensed. If the bracelet is removed after the avatar is already in the region, an invisible sentinel object guarding the region will eject the avatar with approximately 30 seconds notice. Appendix E contains one of the LSL scripts used to implement the bracelet's behavior.

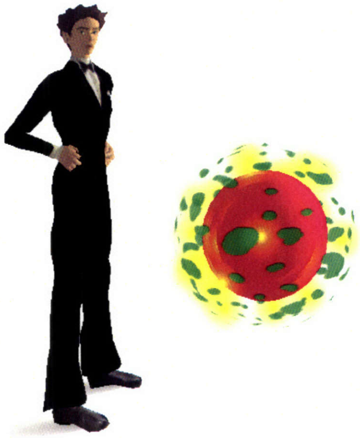

Figure 4-13: A messaging vortex.

4.2.3 Messaging Vortex

The messaging vortex is a portal through which avatars in *Second Life* can send explicit messages into the real world. The vortex can coalesce or dissipate according to any of various stimuli, such as sensor data from the real world, voice commands from avatars, or the approach of an avatar. Once the vortex has formed, an avatar can send a message through it to the real world by touching it and then saying the message over a private chat channel. Once the text of the message is captured, the vortex sends it to an external server where it is converted into a WAV file using the Festival speech synthesis system [93]. The resulting sound file can be played anywhere a WAV file can be played, including over a Plug's speaker. Figure 4-13 depicts a messaging vortex.

4.2.4 Shadow Lab

Shadow Lab is a space in *Second Life* modeled after the third floor of the MIT Media Lab's Wiesner building, where the Plug sensor network is most often deployed. Of the many possible real-virtual mappings depicted in Figure 3-5, Shadow Lab exemplifies one of the most easily understood – the real space to virtual space mapping. The primary feature of Shadow Lab is the to-scale two-dimensional floor plan of the third floor of the building. Only a small portion of the entire space, that corresponding to the Responsive Environments Group lab space, is modeled in three dimensions. In part, this is due to the difficulty and resource drain of modeling everything in three dimensions. However, it is also a design decision reflecting the difficulty in maneuvering an avatar in a to-scale three dimensional space; to-scale spaces invariably feel too confining due to wide camera angles, quick movements, and the coarseness of the avatar movement controls in *Second Life*. Furthermore, the two-dimensional design lends itself more readily to viewing the entire space at once and drawing attention to what few three-dimensional objects inhabit it.

Figure 4-14 shows an early Shadow Lab sketch prototype in which flames represent energy consumption, whirlwinds represent motion, and the raising and lowering of translucent rooms represent aggregate activity in those rooms. Unlike the data ponds, this sketch doesn't evoke the idea of distinct data sources in the real world being represented by distinct instances in the virtual world.

The final implementation of Shadow Lab is a combination of the map of the prototype, approximately 30 data ponds, a message vortex, and a video screen displaying a live video stream, when available, from a next-generation Tricorder device equipped with a camera. The position on the floorplan of each of the data ponds corresponds to its respective Plug's location in the real world.

4.2.5 Metamorphosis

At present, 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 not otherwise moving, such as hands on hips and a bored slouch. All other body language and avatar actions must be intentionally chosen by the user. Clearly, there is room to expand and improve upon this repertoire.

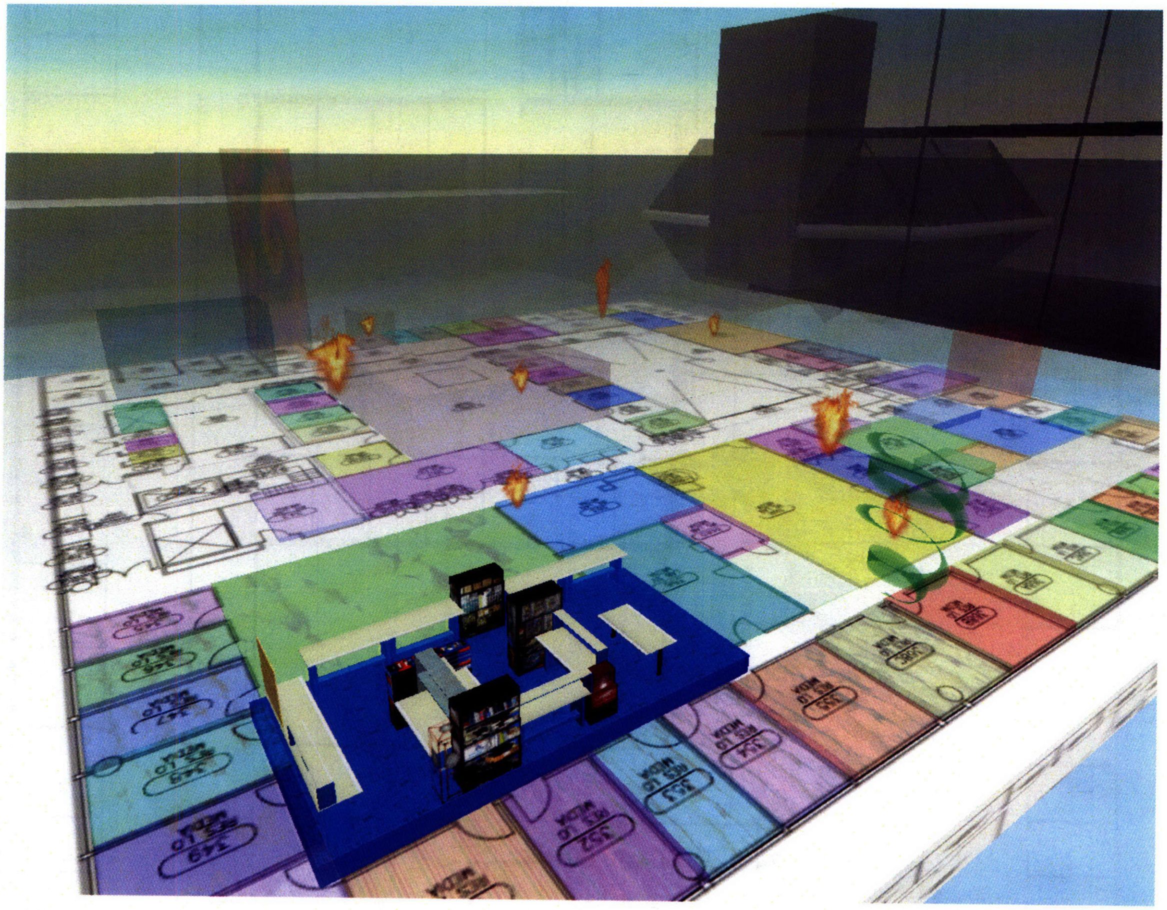

Figure 4-14: An early Shadow Lab prototype.

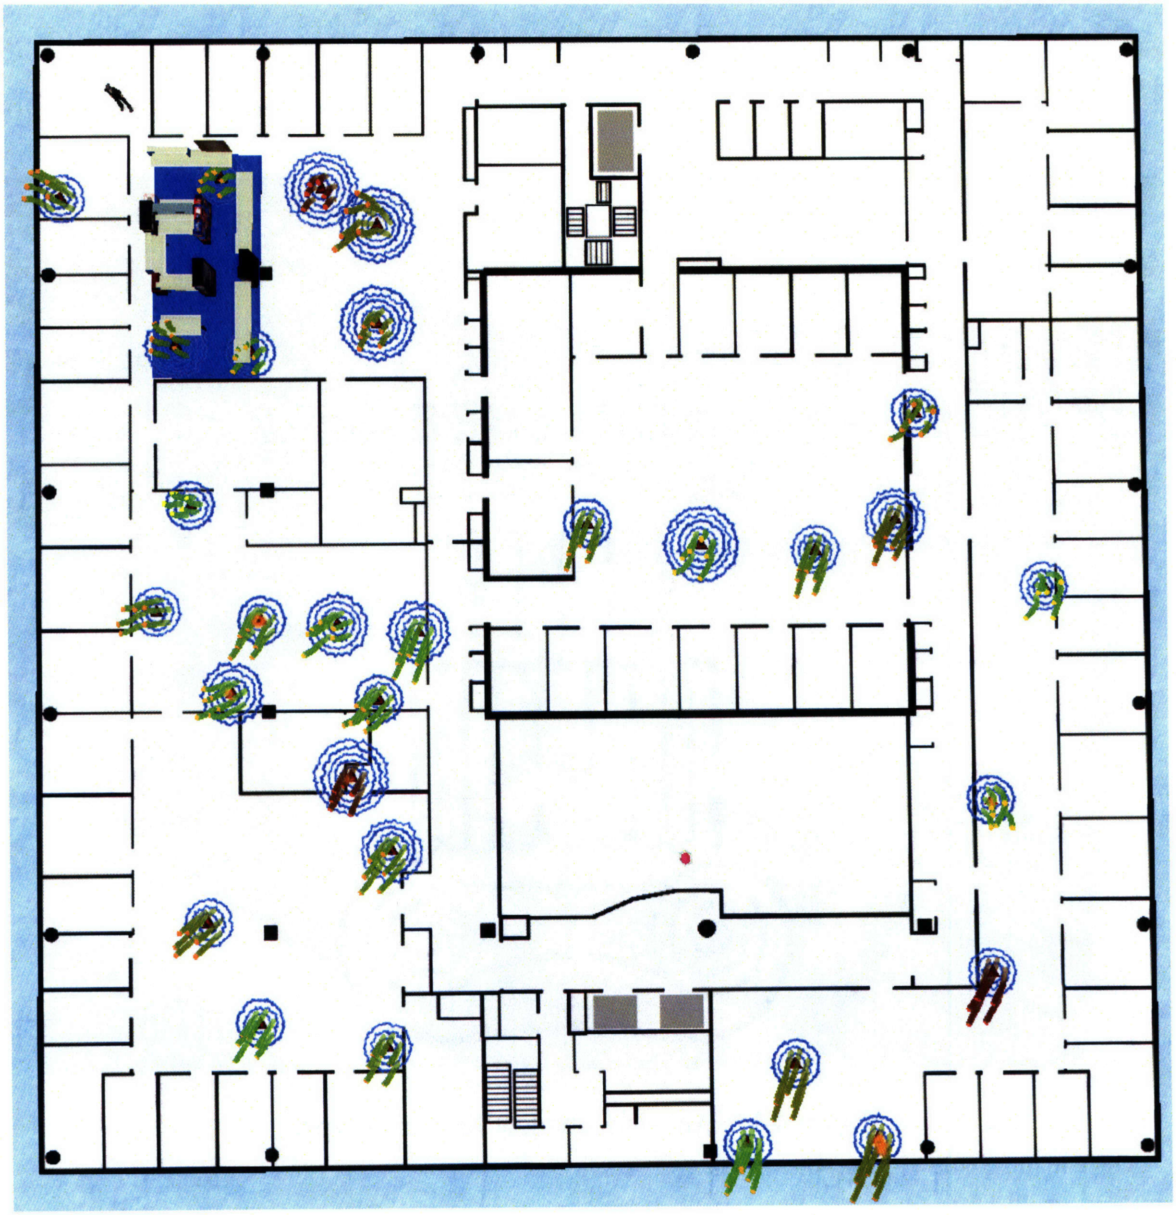

Figure 4-15: Overhead view of final Shadow Lab structure. A human-sized avatar is standing in the upper left corner room.

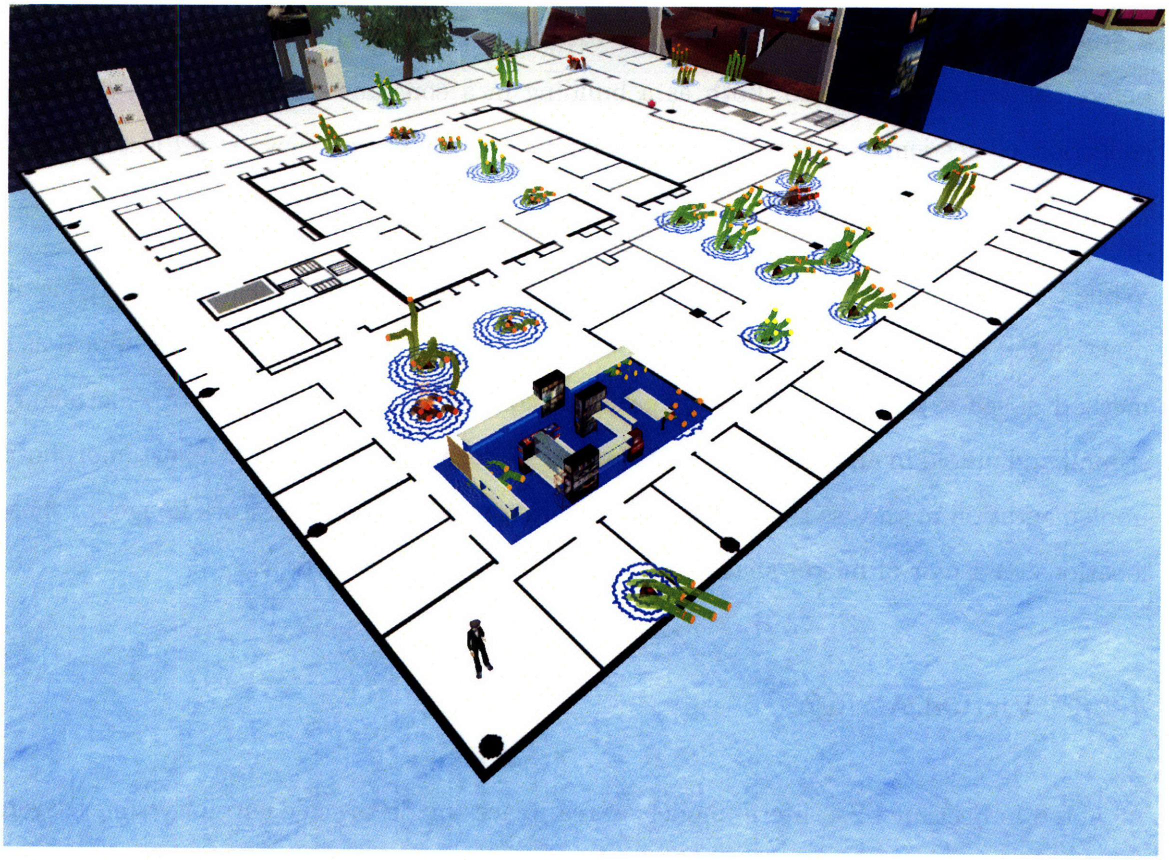

Figure 4-16: Side view of final Shadow Lab structure. A human-sized avatar is standing in the foreground corner room. In the background are buildings belonging to neighboring *Second Life* residents.

Just as Shadow Lab explores mapping real space to virtual space, metamorphosis explores mapping real space to a virtual person. See Figure 4-17. In this prototype, the avatar begins as a typical human and transforms into a Lovecraftian alien according to several parameters drawn from the sensor streams of the Plug sensor network spread throughout the third floor of the Media Lab. As such, the state of a real building is represented as a single avatar in *Second Life*. In some ways, this is the physical analog of synthesizing a group of people's attitudes into a single expert system [94]. Indeed, it may not be long before the entire Media Lab, both as a building in a constant state of activity and as a collection of academics espousing various intellectual points of view, is projected into a virtual world as a single oracle-like avatar.

While the particular metamorphosis depicted in Figure 4-17 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 – the metamorphosis concept can be mapped to other arbitrary stimuli and unfold in an infinite number of ways. For example, a more subtle metamorphosis implementation might use the location of users in a building to reflect stock images of those locations on a pair of mirror shade sunglasses worn by a single avatar.

4.2.6 Virtual Atrium

As noted in Section 4.2.4, literal models of real places made to-scale currently suffer severe usability limitations. Even beyond this practical concern, though, is the more philosophical question of when to impose on the virtual model the limitations of the real space. This is justifiable in many cases, such as when browsing spatially tagged data taken from the real world, or when recreating a real space otherwise inaccessible due to physical hardship, distance, or time. A good example of the latter is the recreation of the Sistine Chapel [95]. However, needlessly carrying the same limitations into the virtual world is somewhat a disservice to both the real space, if any, being modeled and the residents of the virtual space.

With this in mind, the translation of the MIT Media Lab's atrium into *Second Life* attempts

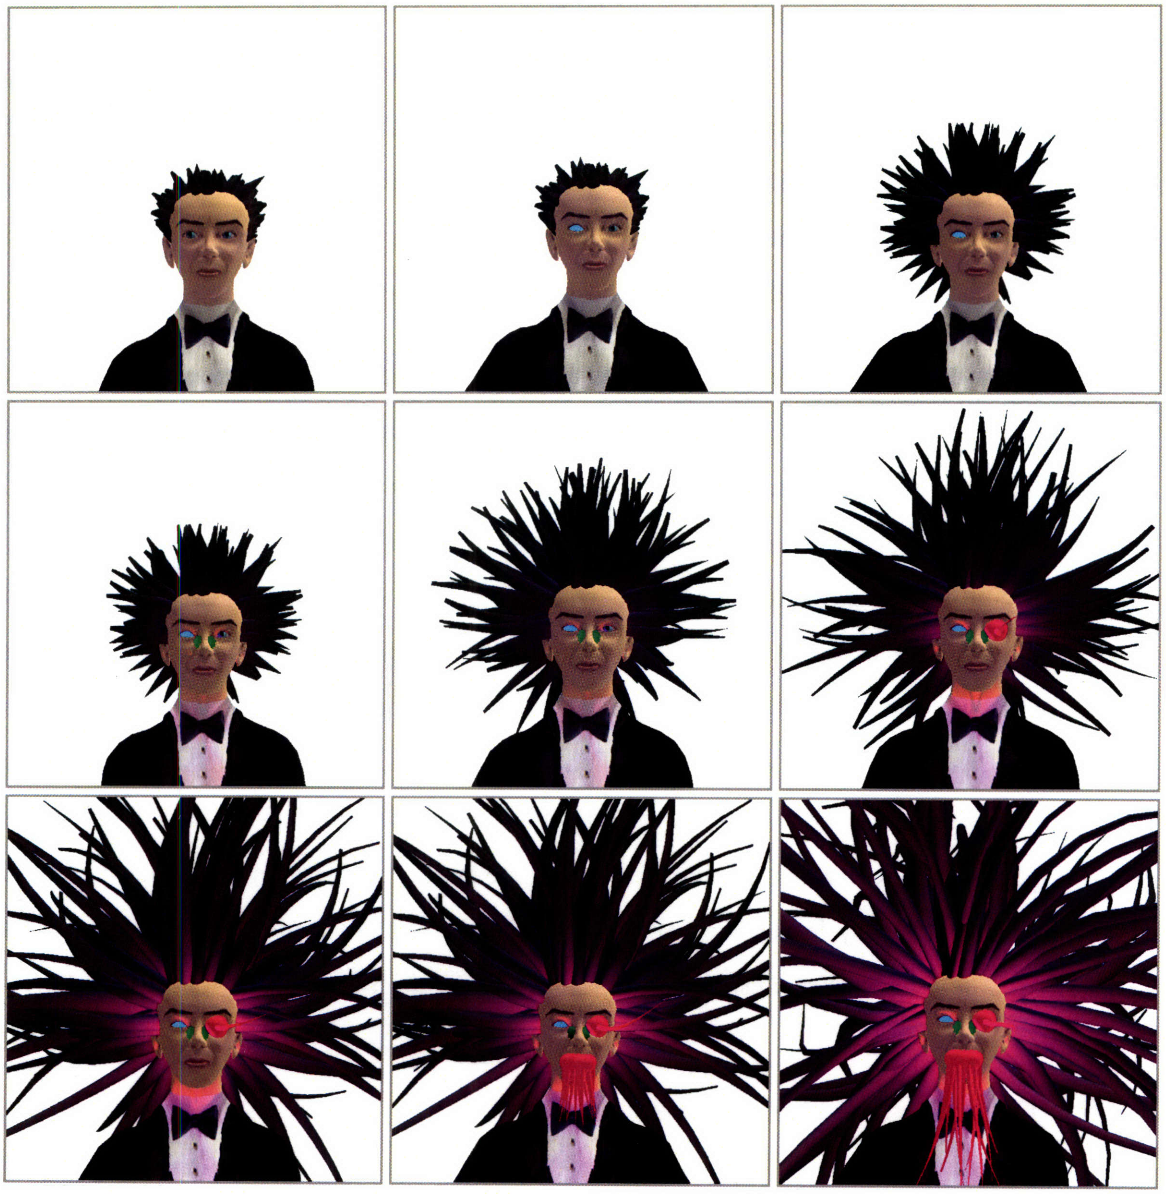

Figure 4-17: A metamorphosis of an avatar (left to right, top to bottom) as activity level in the real world increases.

to retain that which is iconic about the original and at the same time take advantage of the freedom of the virtual world. See Figure 4-18. The virtual atrium is defined by the intersection of two perpendicular walls of tile. A wall assembles itself starting from a single tile that is replicated and positioned as needed according to a set of maximum dimensions. Each tile has a blank white front face, four colored sides, and a black back face. Touching a tile will cause it to flip over, at which point the black back face comes to the front and changes to reveal a hidden movie or image. All tiles in a given wall share the same image or movie when flipped, although the exact image or movie displayed can be dynamically selected from a predetermined inventory chosen by the owner of the wall. See Figure 4-19.

Global commands, such as flipping all tiles at once, can be sent to each wall separately. Physics can be turned on to cause a computationally intensive avalanche of tiles crashing to the ground, only to be re-assembled upon command. The horizontal extent of each wall from the line of intersection is dynamically adjustable, with the trailing tiles at the end popping in and out of existence. The rate at which they do so, as well as the rate at which the four colored sides change their colors is adjustable. Waves of tile flipping or color change can be sent from one end of the wall to the other, or they can simultaneously occur on all tiles. These parameters are easily controlled from external servers, and therefore easily mapped to sensor data from the real world. The motivation for having two walls, aside from helping frame a distinct space, is to have one reflect activity in the real world and the other reflect activity in the virtual world. In so doing the virtual atrium walls can be considered as two axes, each acting as a horizontal bar graph for their respective reality.

4.2.7 Dual Reality Open House

The *Second Life* virtual world is commonly used to host virtual versions of real events, such as academic conferences, movie openings, concerts, and product releases. These events typically rent out a large piece of virtual land on which to erect some information kiosks, an auditorium with a large screen and ample seating, and perhaps some form of entertainment, such as a virtual amusement park ride or virtual game. For live events, the large screen is used to display a slightly delayed audio and video feed of the real event. In the case

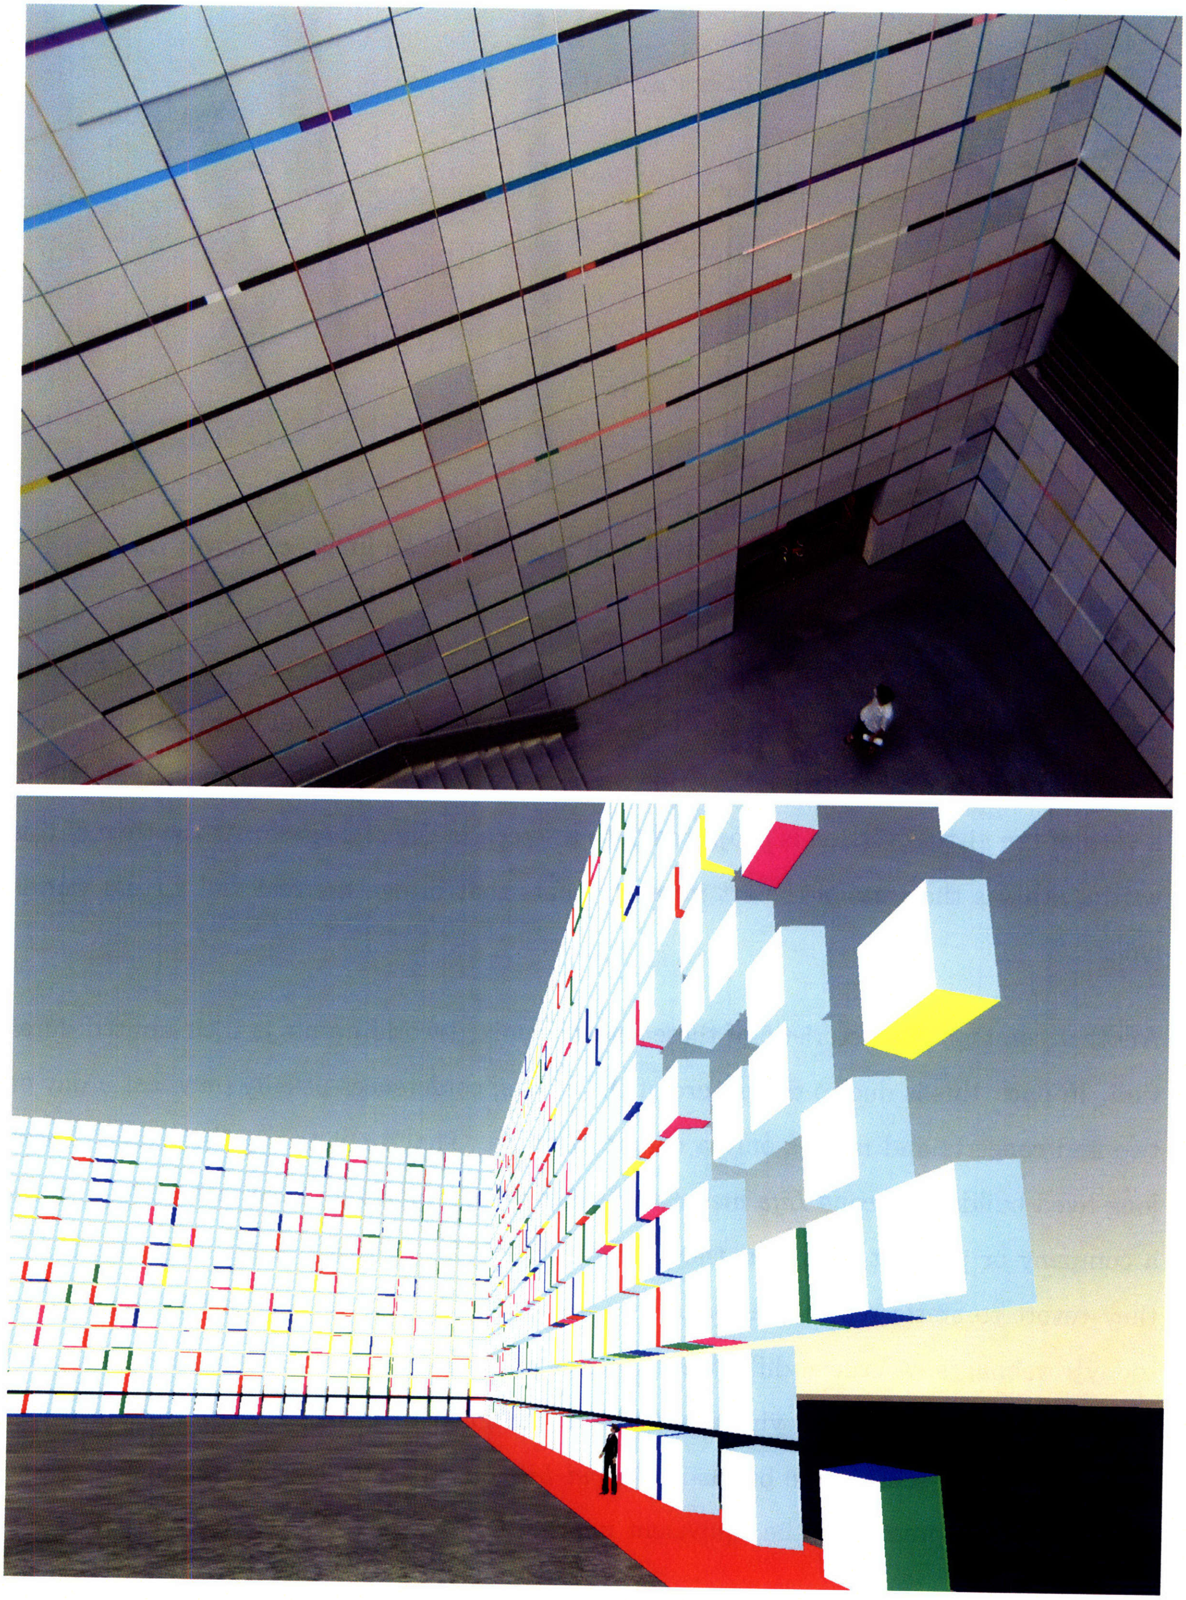

Figure 4-18: The real Media Lab atrium (top) and the virtual version (bottom). A real person and an avatar show their respective scales.

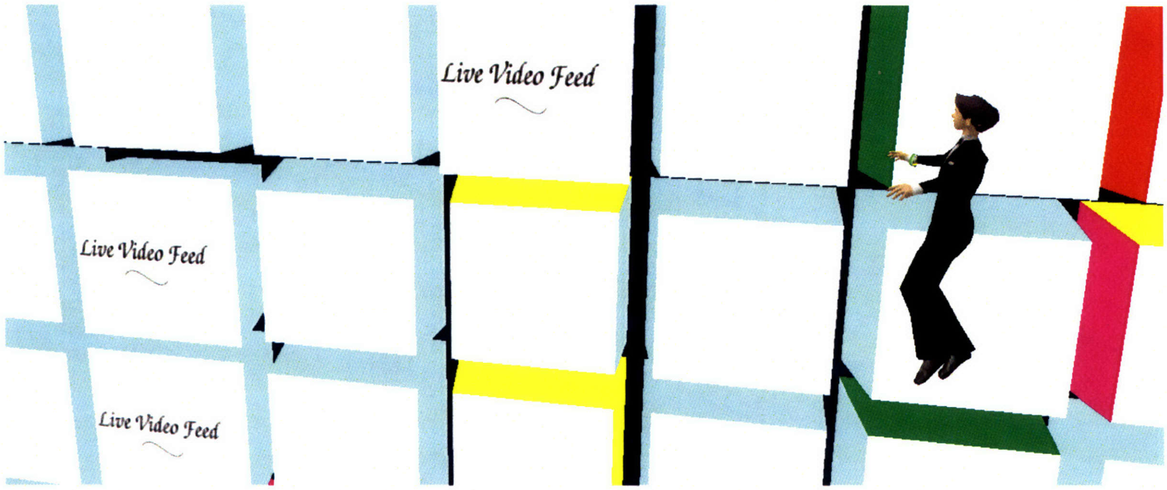

Figure 4-19: Flipped atrium wall tiles display an image or movie.

of interactive events, such as conferences in which the audience can ask questions, there is usually a virtual moderator who acts as a go-between. Real-time audio with no noticeable delay between the real and virtual worlds can assist with audience questions as well. Often, a view of the virtual version is projected onto a screen in the real space. At the time of this writing, this is the state-of-the-art in large events that bridge the real and virtual worlds [96].

Another common characteristic of these so-called mixed-world events is their limited duration. In part, this is due to the nature of the events; conferences usually only last a day or two and concerts and movies only a couple of hours. That's not to say there aren't more long-lived builds in *Second Life* meant to represent some portion of the real world, such as a company or product, just that they don't mix worlds, at least not for very long before they revert to static representations. In a sense, this is very much like the approach of today's webpages. A crucial difference, however, is that the WWW is designed to navigate information, whereas online virtual worlds are designed to navigate virtual physical and social spaces. As such, static builds in *Second Life* quickly lose their allure, much more so than a webpage presenting the equivalent information.

For example, a real-world business might hold a grand opening event for its virtual world branch, which is meant to build brand identity and explore how it might expand into the virtual world market. If they advertise properly and and provide incentives such as give-aways, the opening event will likely be well attended. However, unless they continue to provide incentives for virtual residents to visit, the virtual branch will quickly fall into disuse, in no small part because it is not offering anything customers couldn't otherwise more easily get elsewhere, such as from the business's website. This pattern is well documented and indicates a willingness of real-world businesses to enter virtual worlds for the publicity alone [97]. Of course, formulating viable business plans for real-world businesses venturing into online 3-D virtual worlds is beyond the scope of this dissertation. However, the stagnation of virtual worlds is directly addressable.

Sensor networks offer at least a partial remedy to this problem. Used in the right way, sensor data streams from the real world can enliven the virtual world, and vice-versa. As a prototype demonstration of this, and following the Media Lab's long history of demo culture, a dual reality open house was constructed to introduce residents of *Second Life* to the Media Lab and visitors to the Media Lab to *Second Life*. Specifically, the dual reality open house premiered at a one-day event sponsored by IBM and held in the lower atrium of the Media Lab [98]. The real portion of the event consisted of talks and panel discussions in the building's main auditorium, interspersed with coffee breaks and stand-up meals in the lower atrium among tables manned by Media Lab students demonstrating various Media Lab projects related to virtual worlds.

The virtual portion of the open house is located in the Ruthenium region of *Second Life* [99]. A region is the largest plot of virtual land available in *Second Life*, measuring 256 meters by 256 meters. Although an oversimplification, there is generally a one-to-one mapping between *Second Life* regions and physical servers operated by Linden Lab, the maker of *Second Life*. The capabilities of a region are therefore approximately proportional to the capabilities of the server running the region. The server running the Ruthenium region is limited to 40 simultaneous avatars and 15,000 simultaneous prims. Even the most powerful servers can reliably support only about 100 simultaneous avatars.

In preparation for the open house, Ruthenium was terraformed to resemble a large, four-tiered, terraced amphitheater, as shown in Figure 4-20. Figure 4-21 shows an overhead

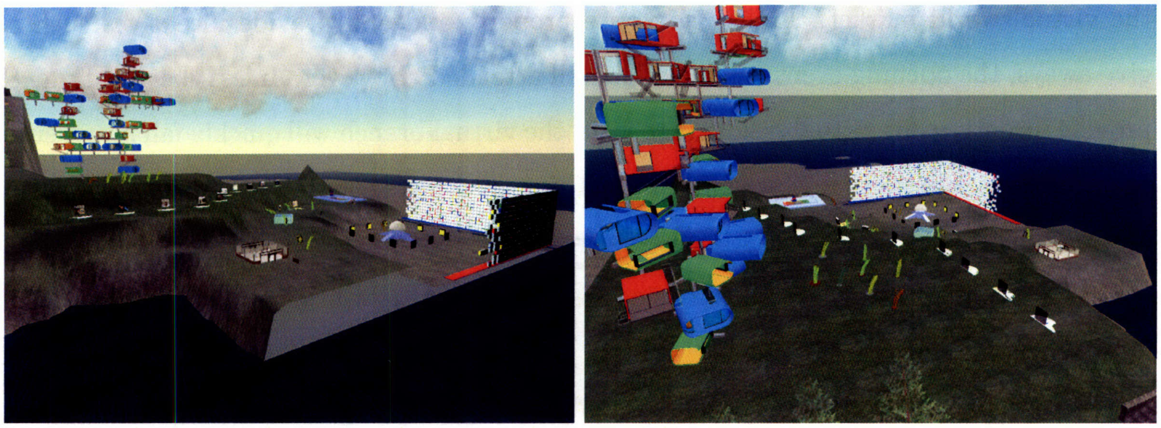

Figure 4-20: Side views of the Ruthenium region.

view of Ruthenium in which the lowest (first) tier is in the upper left-hand corner and the highest (fourth) tier is in the lower right-hand corner. The first tier contains the virtual atrium described in 4.2.6 and a square entry area bounded by three welcome kiosks on each side. See Figure 4-22. The entire Ruthenium region employs the virtual sensing schemes described in Section 4.2.2. The entry area is the only place in the region avatars can enter without a sensing bracelet. All attempts to teleport into the region are redirected to the entry area, where avatars can receive a free sensing bracelet by touching any of the twelve welcome kiosks. The sensing bracelet grants its wearer access to the remainder of the region. See Figure 4-23.

The gazebo in the center of the entry area is the virtual half of "Sousreality", a visual portal between the real and virtual worlds [100][4]. Avatars looking up into the dome roof of the gazebo from below see live video taken from the top of a miniature real-world dome looking up through a fish-eye lens. Similarly, people in the real world looking down from above the miniature dome see a live video projection from a virtual camera at the top of the virtual dome looking down. The effect is one of treating the virtual world as a miniature world that can be peered into as a child would peer into a dollhouse; in this case the dolls can peer back out into the real world. Figure 4-24 shows the real half of the Sousreality prototype, consisting of a one-meter diameter frosted acrylic hemisphere with a fisheye camera at its

---

[4] Sousreality was designed and built by members of the MIT Media Lab's Sociable Media Group.

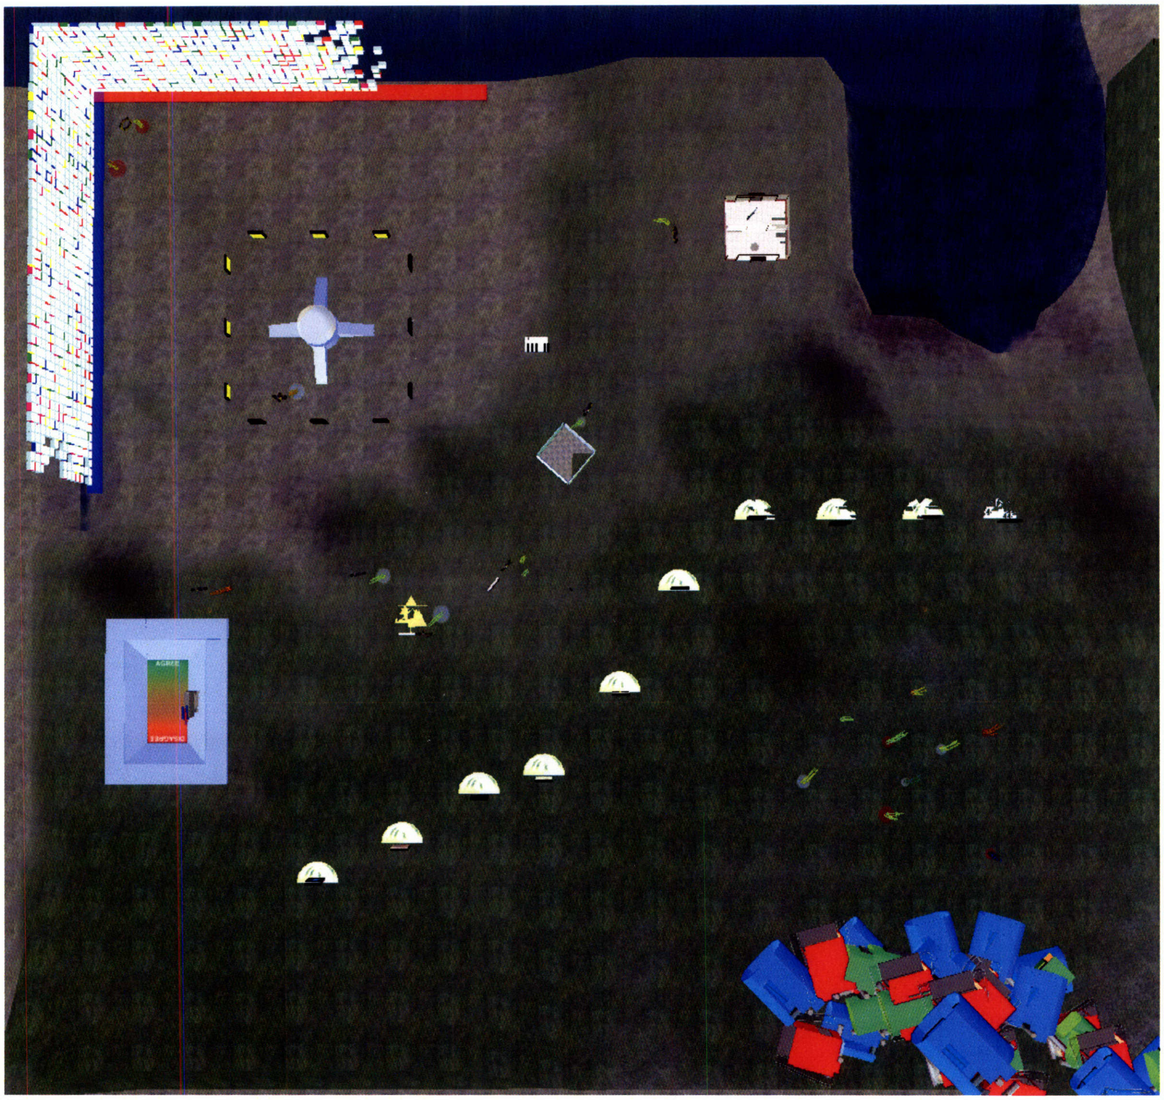

Figure 4-21: An overhead view of Ruthenium, a 256-meter by 256-meter region in *Second Life* where the virtual portion of the dual reality open house took place.

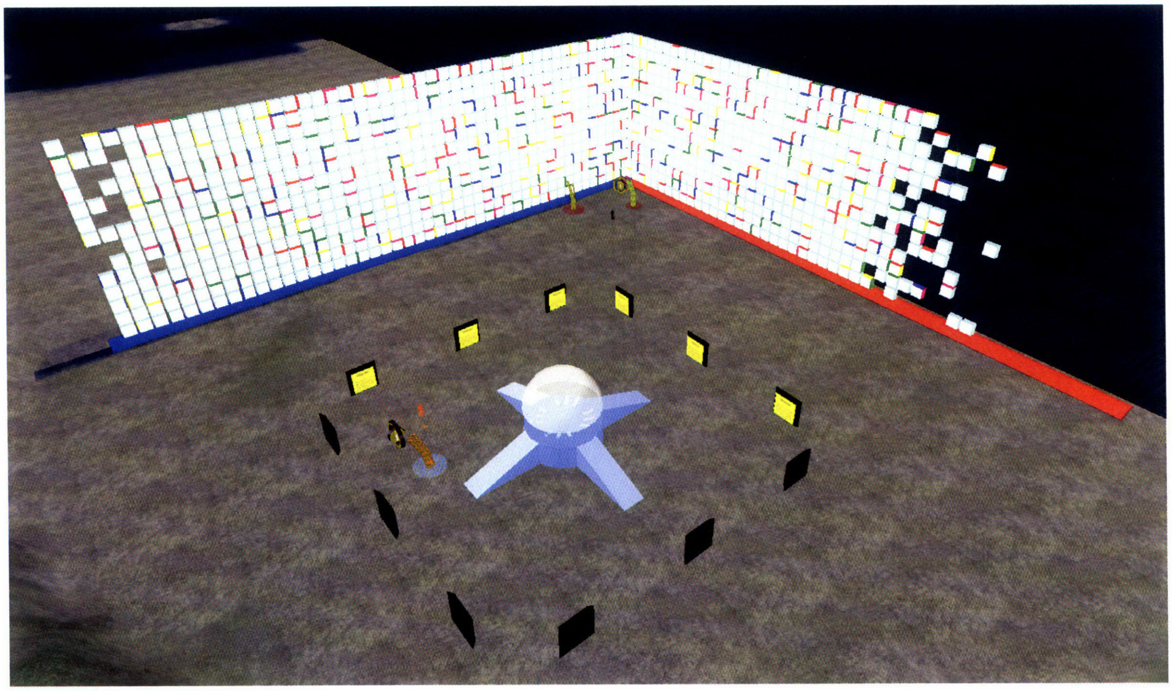

Figure 4-22: The first and lowest terrace of the Ruthenium region contains the entry area and the virtual atrium walls.

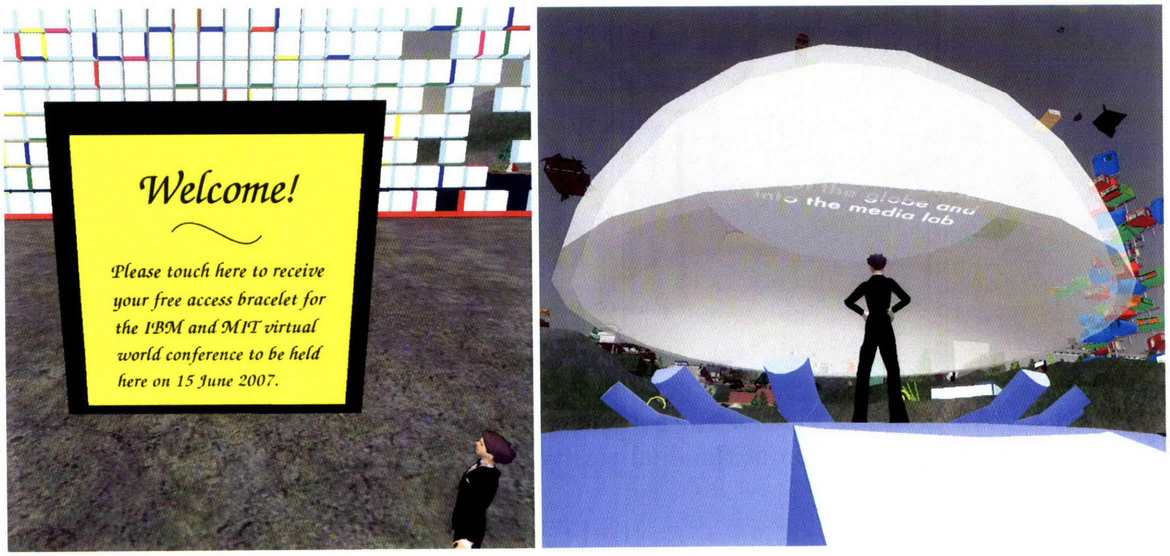

Figure 4-23: The entry area to the Ruthenium region is ringed with welcome kiosks (left) and contains the virtual half of the Sousreality video dome (right).

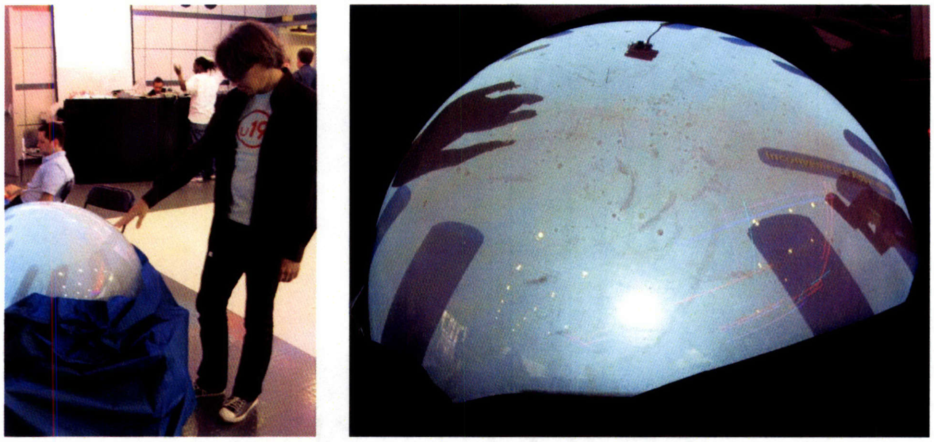

Figure 4-24: The first prototype of the real portion of the Sousreality domed visual portal.

peak and video projector below. Variations of this theme might include reversing which world is considered miniature or adding recursion by providing both worlds with a dome to look down upon and a dome to be surrounded by such that each world gets both a miniature and enlarged view of the other.

The second tier of the Ruthenium region showcases a variety of virtual world-related projects carried out by other members of the Media Lab and demonstrated live and in-person in the real Media Lab atrium during the conference event. In addition to the virtual version of the exhibit itself, each project showcased in Ruthenium has a standardized information icon that dispenses explanatory note cards to avatars upon being touched. See Figure 4-25. All projects with a *Second Life* component were open for virtual visitors to try. Those without instead showed a video within *Second Life* describing the project. The projects exhibited were:

- **Story Cafe:** An interactive narrative generation, storage, and retrieval tool with both *Second Life* and web interfaces [101].
- **The Restaurant Game:** A stand-alone 3-D virtual environment used to record the interactions between thousands of human players and software agents in order to algorithmically combine gameplay experiences to create a new game and better software agents [102].

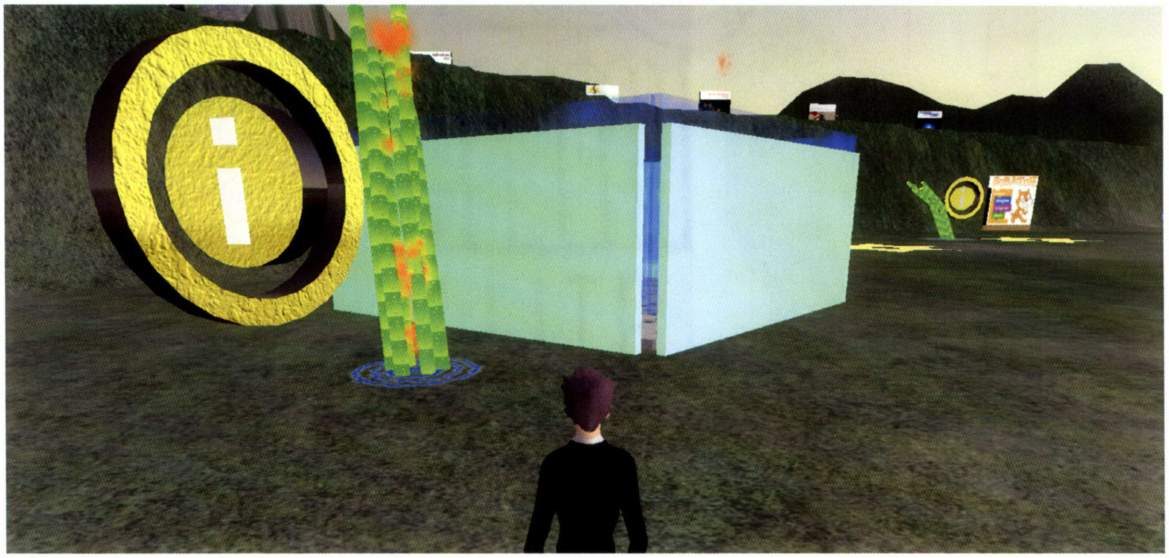

Figure 4-25: A view of part of the second terrace of the Ruthenium region. Each circular yellow information icon identifies a Media Lab project related to virtual worlds and dispenses explanatory notes to avatars. In the foreground is the Autonomous Avatar project and in the background is the Scratch project.

- **Information Spaces:** A *Second Life* meeting space that gives participants a spatial vocabulary for responding to the meeting and visualizes the history of the meeting itself [103].
- **Scratch:** A new programming language designed to be easy to use by young people for creating and sharing on the web interactive stories, animations, games, music, and art [104].
- **Autonomous Avatar:** A primitive autonomous *Second Life* avatar created from a customized version of the *Second Life* client software[5].
- **The Projects:** An extensible cellular building structure in *Second Life* formed by $10m \times 10m \times 8m$ customizable pods that provide various social utilities [105]. (Due to the size of The Projects, it is exhibited on the fourth tier. See below.)

The third tier of Ruthenium is occupied by a series of ten standardized video kiosks, each

---

[5] The Autonomous Avatar was created by Piotr Fidkowski and Andrea Lockerd Thomaz.

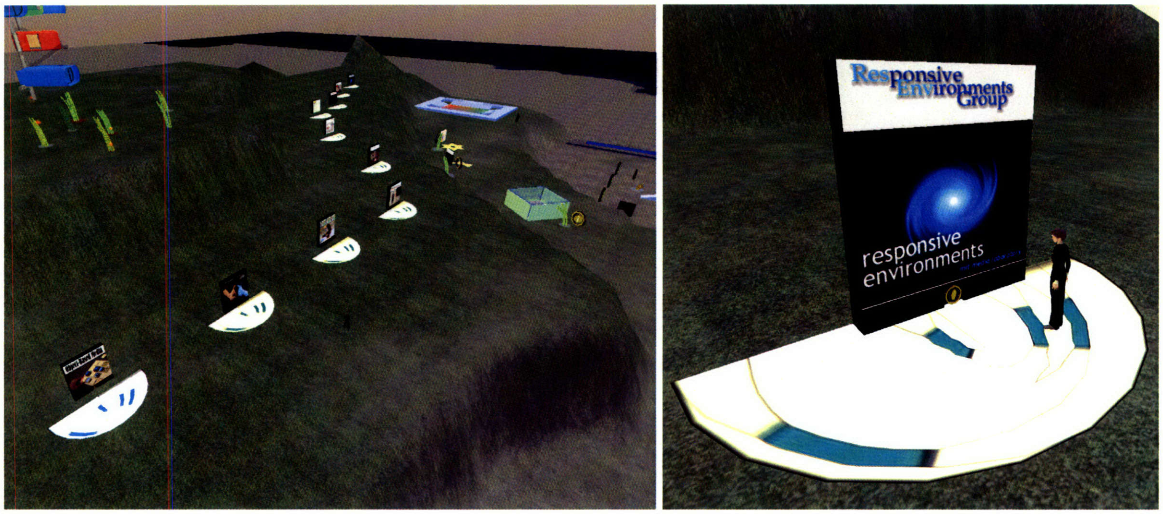

Figure 4-26: The third terrace of the Ruthenium region (left) and one of the video kiosks contained therein (right).

describing one of the MIT Media Lab's research groups. The content at each kiosk[6] is taken directly from that research group's website and formatted for presentation within *Second Life*. See Figure 4-26.

The fourth and smallest tier of the Ruthenium region is devoted to The Projects installation of extensible social architecture and to a field of data ponds. The fourth tier can be clearly seen Figure 4-20.

4.3 Middleware Design and Implementation

The infrastructure connecting the Plug sensor network to the *Second Life* virtual world is necessarily specific to both to some extent, but could be easily modified or extended to accommodate other virtual worlds and sensor network platforms. Figure 4-27 gives a schematic overview of the middleware that allows information exchange between the Plug sensor network and *Second Life*.

---

[6] The video kiosks and their content were designed and assembled by Ally Lee, Daniel Jang, and Jason Uh of the Story Cafe project and Becky Bermont, the Media Lab's Director of Sponsor Management.

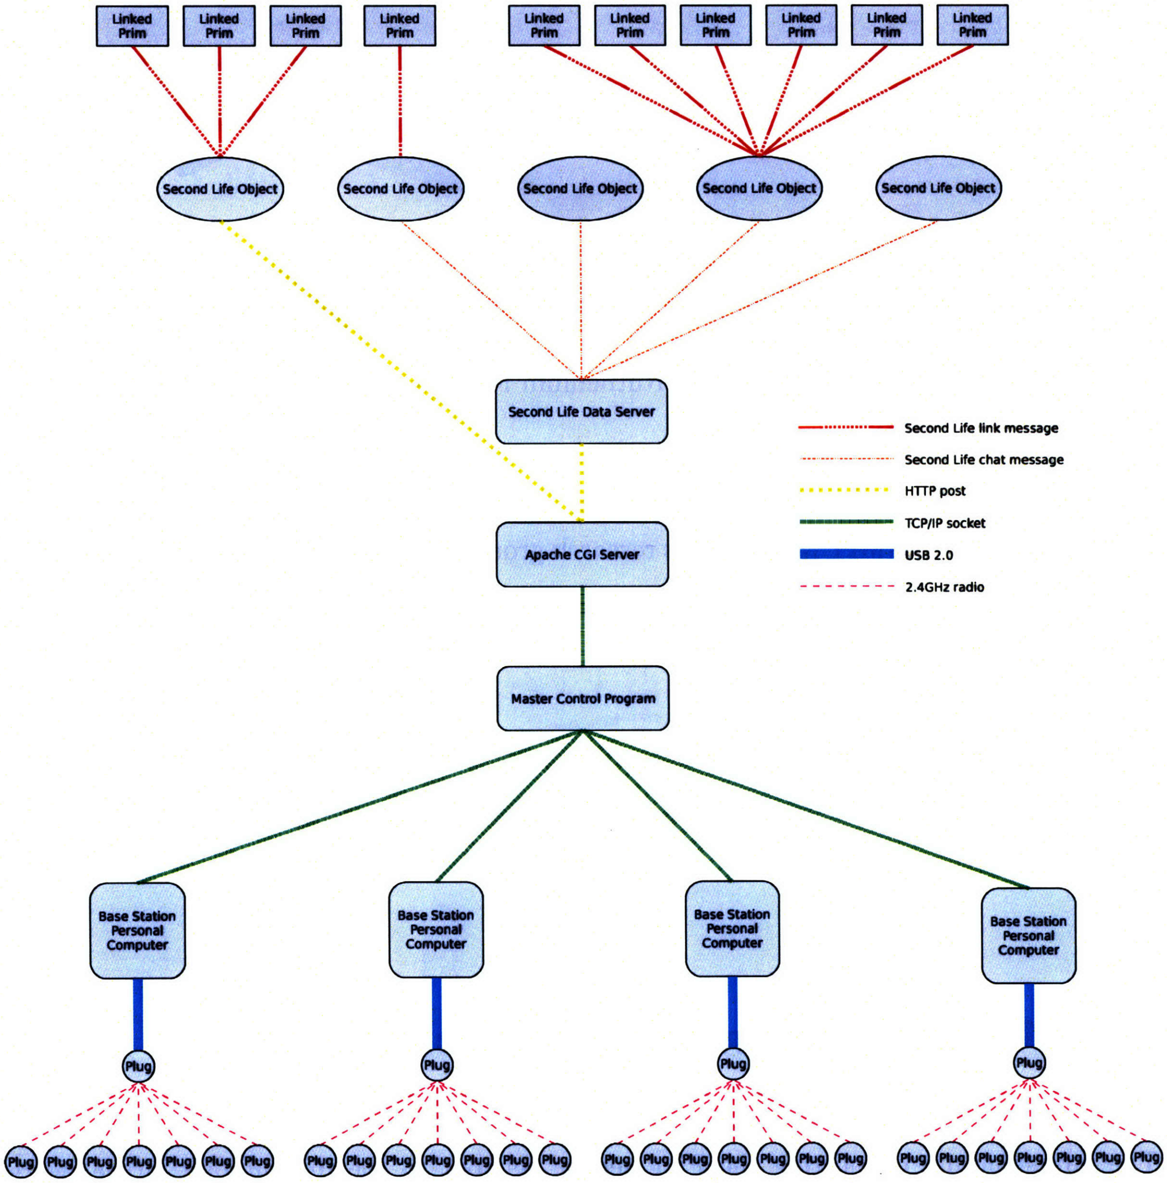

Figure 4-27: The general structure of data flow for communication between *Second Life* and the Plug sensor network.

The middleware revolves around a single master control program (MCP) that makes all decisions about what data to send where. The MCP responds to user input from the command line, sensor data payloads from the Plug sensor network, sensor data payloads from the *Second Life* virtual sensing systems, and requests for data from *Second Life* objects. The user can query the state of the MCP and start and stop it from the command line. Upon receipt at the MCP, all sensor payloads are timestamped, used to update a model of the current state of the source of the payload, and logged to a flat file. The only way for the MCP to transfer data to an object in *Second Life* is to respond to a request for data from that object. In essence, *Second Life* must poll the MCP for updates; the MCP cannot initiate a data transfer to a *Second Life* object. This limitation is due to the modes of communication available for interfacing to *Second Life* from the outside world. An XML-RPC interface is available, but notoriously unreliable and often slow, on the order of several seconds per transfer. E-mail has more reliable delivery, but unreliable speed. The only method with both reliable delivery and speed is HTTP requests originating from within *Second Life* directed to servers outside of *Second Life*.

The HTTP requests made from *Second Life* are directed toward an Apache web server and initiate a CGI script that opens a TCP/IP socket connection to the MCP to relay the request. The MCP immediately responds over the same socket, the socket is closed, and the response is returned by the web server as the result of the original HTTP request. In order to limit server load caused by HTTP requests and to curb the use of HTTP request for malicious purposes (e.g. denial of service attacks), *Second Life* limits the rate of HTTP requests made on behalf of each user. In practice, this means requests for data from the outside world must be aggregated by a central *Second Life* data server object and then dispatched as a single HTTP request. The response from the outside server is then parsed and distributed to the intended recipient objects using *Second Life*'s chat system, which can then further disseminate the data to their constituent prims.

Data transfer between the MCP and nodes in the Plug sensor network is mediated by a group of base station personal computers, each connected via USB to a single Plug device. The Plug sensor network is partitioned into non-overlapping sets, each reporting to a single base station. As discussed in Section 4.1, the Plug nodes communicate among themselves using a low-power 2.4GHz radio. For reasons of network reliability and speed, a base station only requests data from Plugs within a single radio hop of the base station's USB-tethered Plug. Other variations of the middleware were implemented that omitted the base station personal computers in favor of the MCP communicating directly with a single Plug device that acted as the only data sink in a multi-hop mesh network of Plugs, but communication was not as reliable and much slower due to the bottleneck of a single network sink. Without supervision of the MCP, base stations periodically query their Plug nodes in round-robin fashion and report the results to the MCP via a persistent TCP/IP socket.

As a rudimentary example, consider a situation in which the behavior of a virtual flower in *Second Life* is determined by the average light reading from all Plug nodes, and the sound played over a particular Plug's speaker is determined by the number of avatars near the virtual flower. All Plug nodes continuously gather sensor data and the base stations periodically query each node approximately once a second. For each query of a single Plug, the response is radioed by the queried Plug to the Plug tethered to the base station, from there by USB to the base station, and then by TCP/IP socket to the MCP. The MCP stores the response both in a log file and as the most recent update from the queried Plug. Meanwhile, the *Second Life* data server object makes requests from within *Second Life* to the Apache web server approximately once per second. Upon receiving a request the web server forwards it by TCP/IP socket to the MCP, which responds by sending the web server the most recent updates for all Plug devices. The web server returns this payload directly to the *Second Life* data server as the response to the data server's original HTTP request. The *Second Life* data server object passes the response to the virtual flower, and any other interested virtual objects, via *Second Life*'s chat system. The flower object then parses the response and sends *Second Life* link messages instructing each of its petals to open a certain amount, its stem to grow a certain amount, and its leaves to turn a certain shade of green, all according to the Plugs' average light reading as indicated by the parsed response.

In the other direction, the sensing bracelet of each avatar in the region periodically reports directly to the Apache web server with the avatar's location and other parameters. The web server relays these reports to the MCP, which logs them to a file and calculates how many avatars are within five meters of the location of the flower and how close they are. Given n avatars within five meters, each at a distance d from the flower, the MCP sends to the proper base station an instruction for a particular Plug to play a sound of footsteps at a volume proportional to $\sum_{k=1}^{n}(5-d)$, which is always greater than zero since d is assumed to be less than five meters. This instruction is in turn relayed to the base station's Plug and finally to the destination Plug. Alternatively, the virtual flower itself could have sensed nearby avatars and communicated directly to the web server without having to go through the *Second Life* data server object.

Chapter 5

Evaluation and Discussion

"Pick battles big enough to matter, small enough to win. Lucky Numbers 7, 14, 25, 29, 36, 44"

Golden Bowl Fortune Cookies

The vision presented in Chapter 3 of sensor networks being the palette and online virtual worlds being the canvas by which mass consumption and mass communication of media are finally complemented by mass creation of media is very far from being realized. The technical infrastructure alone presents a daunting challenge – sensor networks as pervasive as today's electrical and lighting systems and online virtual worlds as populated and useful as today's WWW are at least years away, if not decades. Even if the requisite technologies come to pass, their merger and adoption for purposes of media creation are not certain. Technologies promising to change the world have accumulated in the dust bin as human nature and chance dictate otherwise; witness the Segway Human Transporter and the International Space Station.

This dissertation attempts to take advantage of the nascent technologies available today to explore the landscape of possibilities and push, however modestly, toward the vision of mass media creation. The gulf between vision and reality poses certain challenges to evaluation.

On one hand, the area of inquiry is new enough to have very little to compare against. On the other hand, what is accomplished here is by no means a mature body of work and will likely differ considerably from the final version, if such a thing comes about at all. In this light, the remainder of this chapter evaluates on several levels the work described in Chapter 4. The evaluation ranges from the analytical to the anecdotal and tries to gauge the overall success of the tools and examples presented here.

5.1 Plug Sensor Network Evaluation

5.1.1 First Impressions

In the fall semester of 2005 at the MIT Media Lab, a graduate-level introductory course about sensor networks was offered based around the Plug [106]. The 12 students in the course built Plugs and Lugs, and then prototyped various applications with them. The class participants were a largely graphic, industrial, and interaction design group of people with minimal sensor or electronics experience (which guaranteed the time debugging their Plugs far outweighed the time they spent building them, not that this was ever meant as a means to save labor or time). In introducing sensor networks in general and Plugs in particular to a broad group of designers, it became clear that their primary interests were developing interfaces for browsing the data collected from the Plugs and using the Plugs as actuators. A concrete result of the class was the incorporation of the Plug codebase and aspects of the Lug hardware design into wearable badge devices that measure human behavior according to various sociometrics [107].

5.1.2 Single Plug Evaluation

Figure 5-1 shows some of the data taken from a single Plug during a rudimentary scenario – a desk lamp is plugged into a Plug node's electrical outlet and turned on. Even scenarios as simple as this make clear the value of the Plug's multi-modal sensing abilities for disambiguating the context of the event. For example, the current sensor data indicate precisely

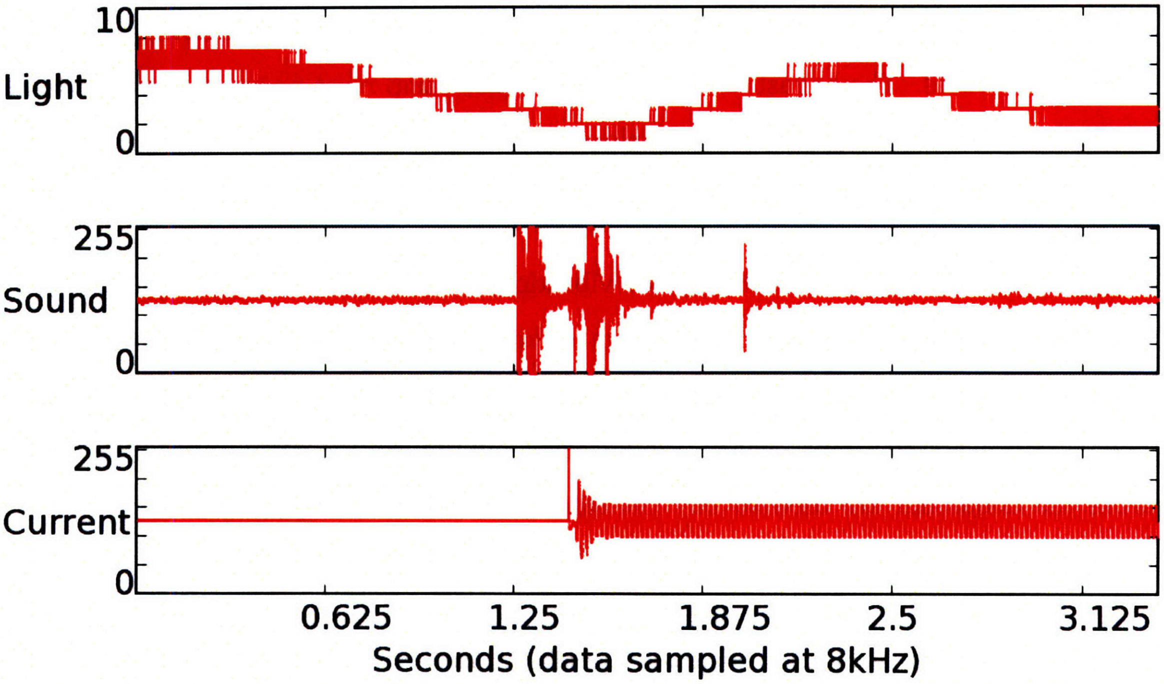

Figure 5-1: Data taken from a single Plug node as a desk lamp is being plugged into one of its electrical outlets. The vertical axes are in arbitrary units. The vertical axis of the "Light" plot has been scaled to show greater detail. All data were sampled at 8-bit resolution.

when the lamp was turned on, but cannot be used to discern whether or not the lamp was already plugged in but turned off. The microphone data, on the other hand strongly indicate the event included the lamp being plugged into the outlet. This is further corroborated by data from the light sensor, which show the shadow of the hand passing over the Plug as the lamp is being plugged in. These hypotheses are strengthened by looking at the binary motion and vibration indicators. Finally, given that a lamp was plugged in, which might be inferred directly by a more detailed analysis of the current signature, it is a safe guess from the relatively constant light reading that the lamp is not shining on the Plug directly and may be far removed due to an extension cord, as was the actual case. Automated methods for such qualitative single event accessment from many sensor channels are discussed in [108].

Figure 5-2 clearly shows how individual and classes of electronic devices can be identified and classified by their current signature alone, as mentioned in Section 3.1.1. For example, the digital oscilloscope, LCD monitor, and desktop computer all have distinctive, predictable startup sequences. The ballast of the fluorescent desk lamp must power up before the gas in the bulb will fluoresce, whereas the halogen desk lamp is an almost purely resistive load and therefore its current draw is proportional the voltage applied.

Figure 5-3 shows the electrical current, vibration, and sound levels of a typical snacks vending machine over the course of several minutes. Three events take place over this period. First, someone walks by and hits the machine (see vibration graph). Second, two people in conversation buy a snack using a dollar bill. Third, the coins returned as change from the dollar bill are used to buy the same snack. For each purchase, there are two current spikes, one for the exchange of money and one for the actuation needed to dispense the snack. Conversation ensues throughout the second two events. The vending machine scenario highlights the Plug's potential to easily retrofit appliances otherwise difficult to instrument.

Naturally, not all inferences can be made at the node level; some inferences are either too computationally demanding or in need of extra information. In such cases, a likely solution is to reduce the raw data to features at the node level and then communicate these features elsewhere for further processing. As proofs of concept, simple versions of such algorithms were developed for the Plug by other Media Lab researchers to classify types of light (e.g. fluorescent, incandescent, halogen, and natural) and types of electrical devices (e.g. resistive, switching, and inductive) [109, 110].

5.1.3 Plug Network Evaluation

Looking at the network as a whole, general trends of activity across the building are easily discernable. Figure 5-4 shows a map of the third floor of the Media Lab and the locations of each Plug during a data collection run lasting about 20 hours starting very early Monday morning and going until late Monday night. For this run, 31 of the 35 Plugs were deployed. The data collected from each Plug include five-second windowed and rectified minima, maxima, and averages of light, sound, voltage, and current, as well as cumulative motion

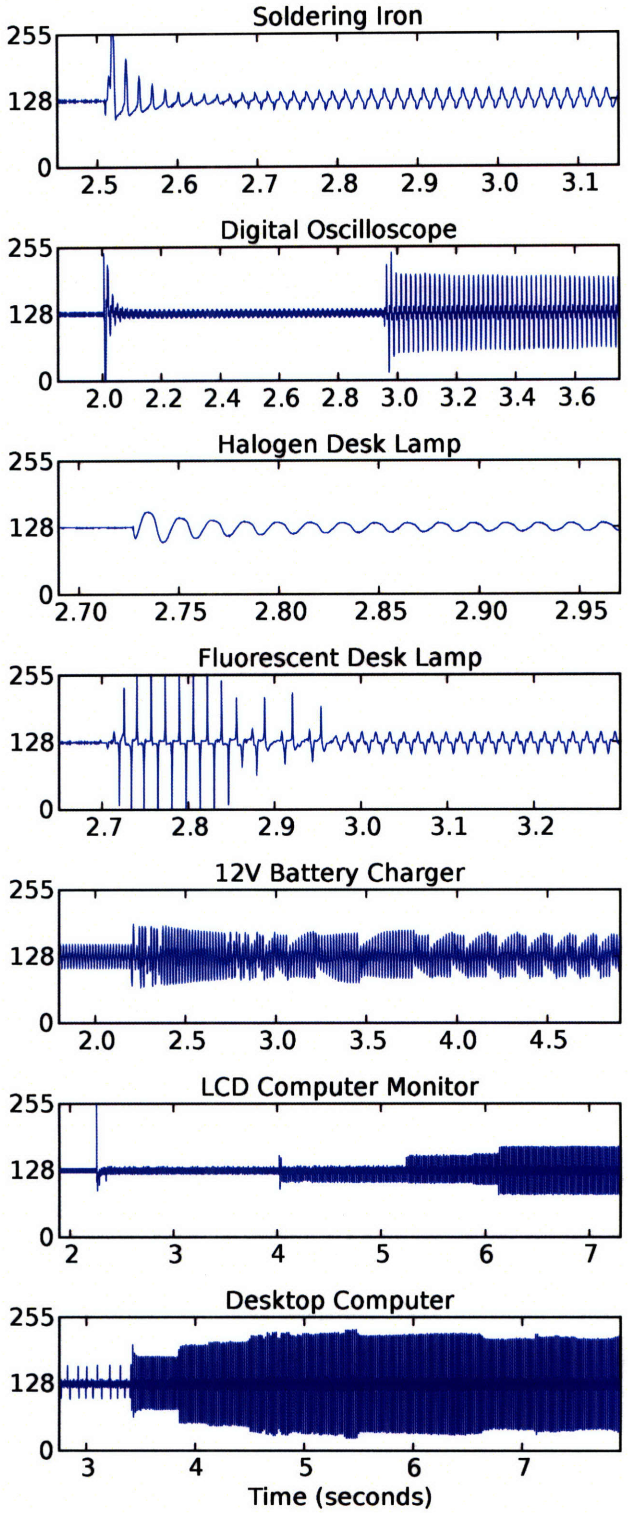

Figure 5-2: Plots of current versus time taken from a Plug sensor node for a variety of common electronic devices. Each plot shows current data from a several second window encompassing the time at which the device was plugged into the Plug node's electrical outlet. All data were sampled at 8kHz. All data are in arbitrary units.

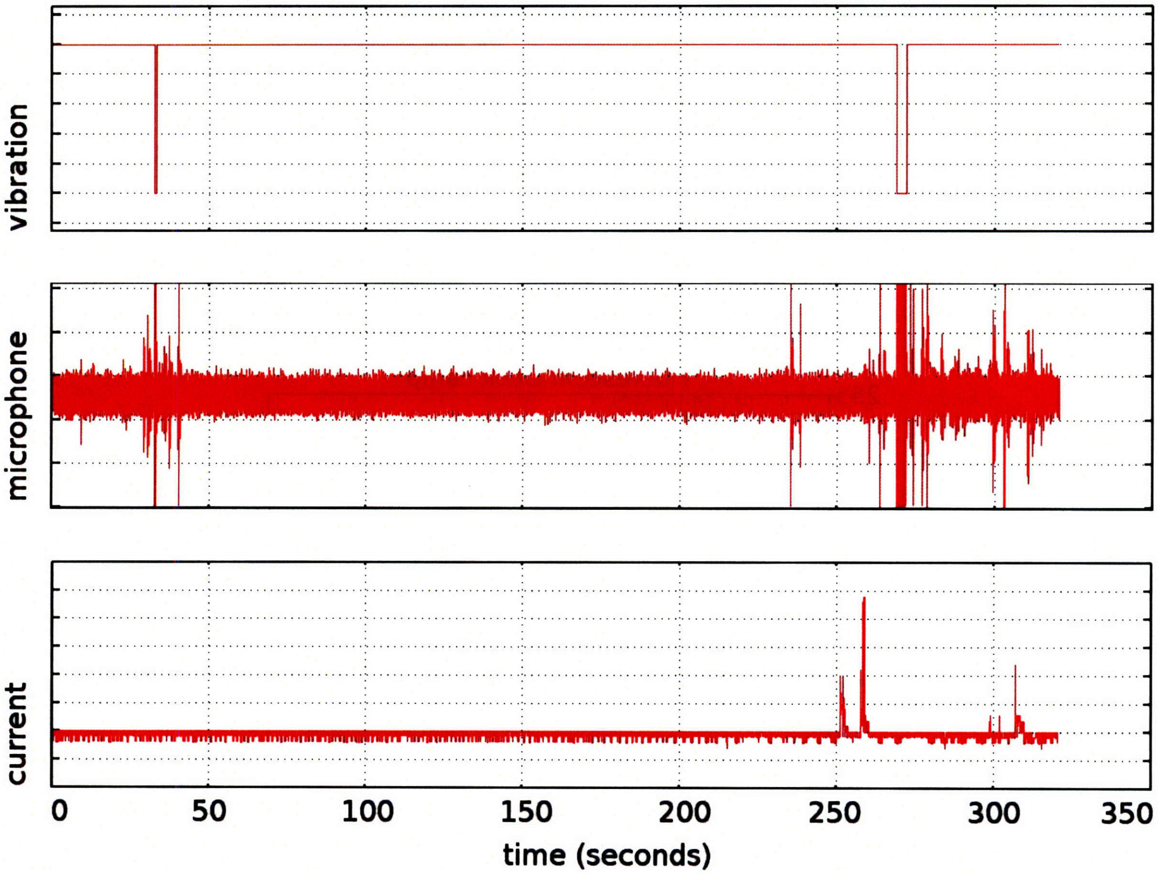

Figure 5-3: Plots of vibration (top), sound (middle), and electrical current (bottom) versus time in a typical snacks vending machine. The vibration sensor takes on an active-low binary value and the microphone and current sensor were sampled at 500-Hz with 8-bit resolution. The electrical current plot is a low-passed, peaked version of the original AC signal. All data are in arbitrary units and scaled for ease of viewing.

and vibration activations (the motion and vibration sensors being binary) and the route by which the network packet arrived to its destination. The samples from which the extrema and averages were calculated were taken at 8kHz. The shade of each circle represents the total "activity" seen over the entire course of the data collection, where activity is defined as an equally weighted sum of total motion and total vibration activations. All data were routed to a single Plug over the radio network (the dark circle in the upper right corner, number 18) and siphoned off to a personal computer via a single USB connection. This graph of activity level corresponds well with intuitive impressions of the activity level of different parts of the building. For example, Plug 03 was placed next to a heavy door leading to the main kitchen and cafe area, making its high activity level unsurprising.

Figure 5-5 shows light (maximum over five-second window), sound (maximum over five-second window), and current (rectified average over five-second window, averaged across all four outlets) data over time from three specific Plugs during the same data collection period mentioned above. In this case, trends taking place over an entire day become apparent. For example, the light readings from Plugs 22 and 30 clearly show the sun rising and setting. (Although the map indicates Plug 22 is located near the middle of the building, it was placed next to a window that overlooks a large atrium with sky lights). Plugs 23 and 30 show office lights being turned on and off, albeit at different times of day. The current draw from Plug 30 shows a desktop computer being started, shutdown, and rebooted at various points in the evening, whereas Plug 23 only had small DC converters plugged into it. The sound level graphs indicate discrete events (spikes in the graph) and general activity level. The several visible gaps in Plug 23's data sets are cases of packets being dropped in the network. In general, packet loss increased with increased network distance from the data collection node, in this case most likely due to node-to-node packet transfer using unacknowledged packets.

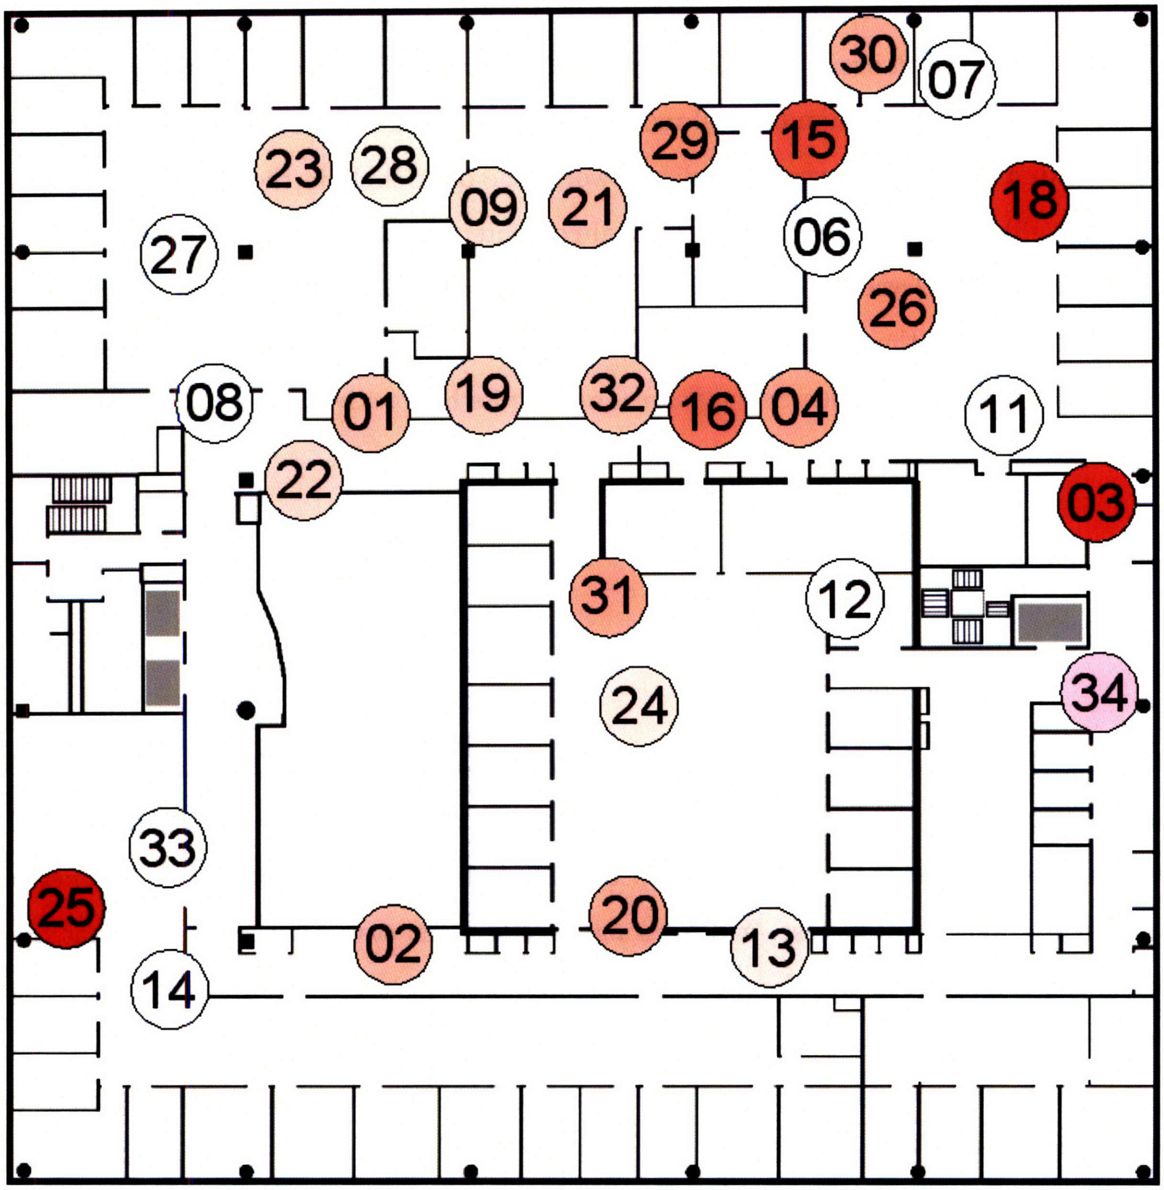

Figure 5-4: A map of the third floor of the MIT Media Lab. The 31 large circles indicate the locations of Plug sensor nodes. The number within each circle is the ID of the Plug at that location. The darker the circle, the more activity occurred at that node over the span of a 20-hour data collection period. Here, "activity" is defined as the sum of the number of motion sensor and vibration sensor activations.

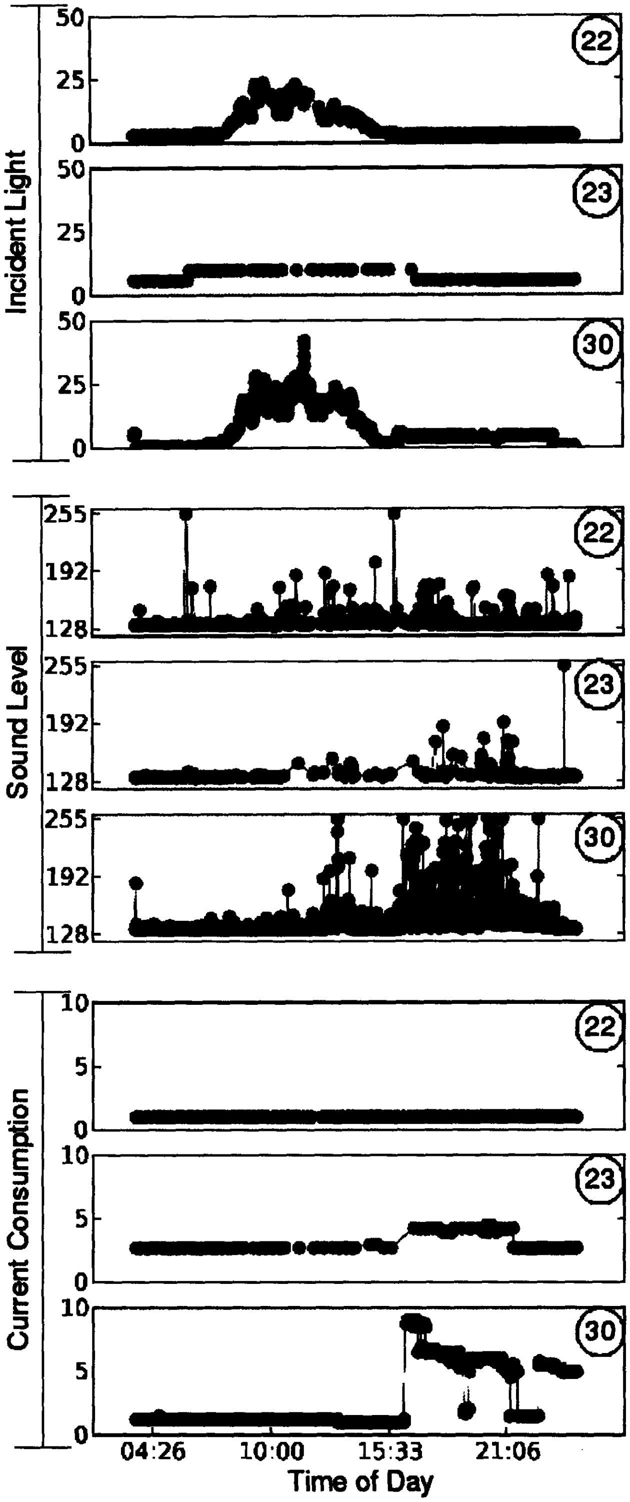

Figure 5-5: Light (top three graphs), sound (middle three graphs), and current (bottom three graphs) data versus time of day for three of the Plug sensor nodes shown in Figure 5-4. The circled number in the upper right-hand corner of each graph is the ID of the Plug and corresponds to the location shown in Figure 5-4. All data are in arbitrary units.

5.2 Dual Reality Evaluation

5.2.1 First Impressions

Over the course of the construction and testing of the various virtual world concepts described in Section 4.2, on the order of 100 passersby within *Second Life* stopped to interact with the space and Nix Division, the author's avatar managing the space. Some of the passersby own neighboring virtual land, but most happened across the construction site by chance. Visitors would typically explore the area on their own before asking what it was all about. Their responses to the explanation of dual reality and the various concepts under construction were invariably positive and often accompanied by ideas of their own. In general, visitors were eager to interact with the various installations, even while under construction; Figure 5-6 depicts examples of the frequent "graffiti" found on the virtual atrium walls described in Section 4.2.6.

Several of the visitors own clubs in *Second Life* and were interested in the possibilities of bringing dual reality ideas to their club in order to attract and maintain a customer base, as discussed in Section 3.3.2. With the herd mentality being particularly prevalent in *Second Life*, differentiation is an important asset among clubs trying to capitalize on the virtual social scene. Typically, a club's owner is a very visible presence in the club's day-to-day operation and the owner's personality and expression thereof within *Second Life* is an important factor in attracting customers. As such, even though most virtual club owners don't own real clubs, they could still benefit from enlivening their virtual clubs with sensor data as a means to further build and express their virtual personas.

Other visitors had just started using *Second Life* and still hadn't acclimated to the environment or come to understand what was and wasn't possible in *Second Life*. With very little deviation, conversations with these newcomers revolved around four topics, usually in quick succession: the origin of avatar names, the meaning of the immediate virtual environment, avatar appearances, and real-world locations of avatar owners. These are essentially the only four topics a newcomer to *Second Life* could possibly discuss in an attempt to develop

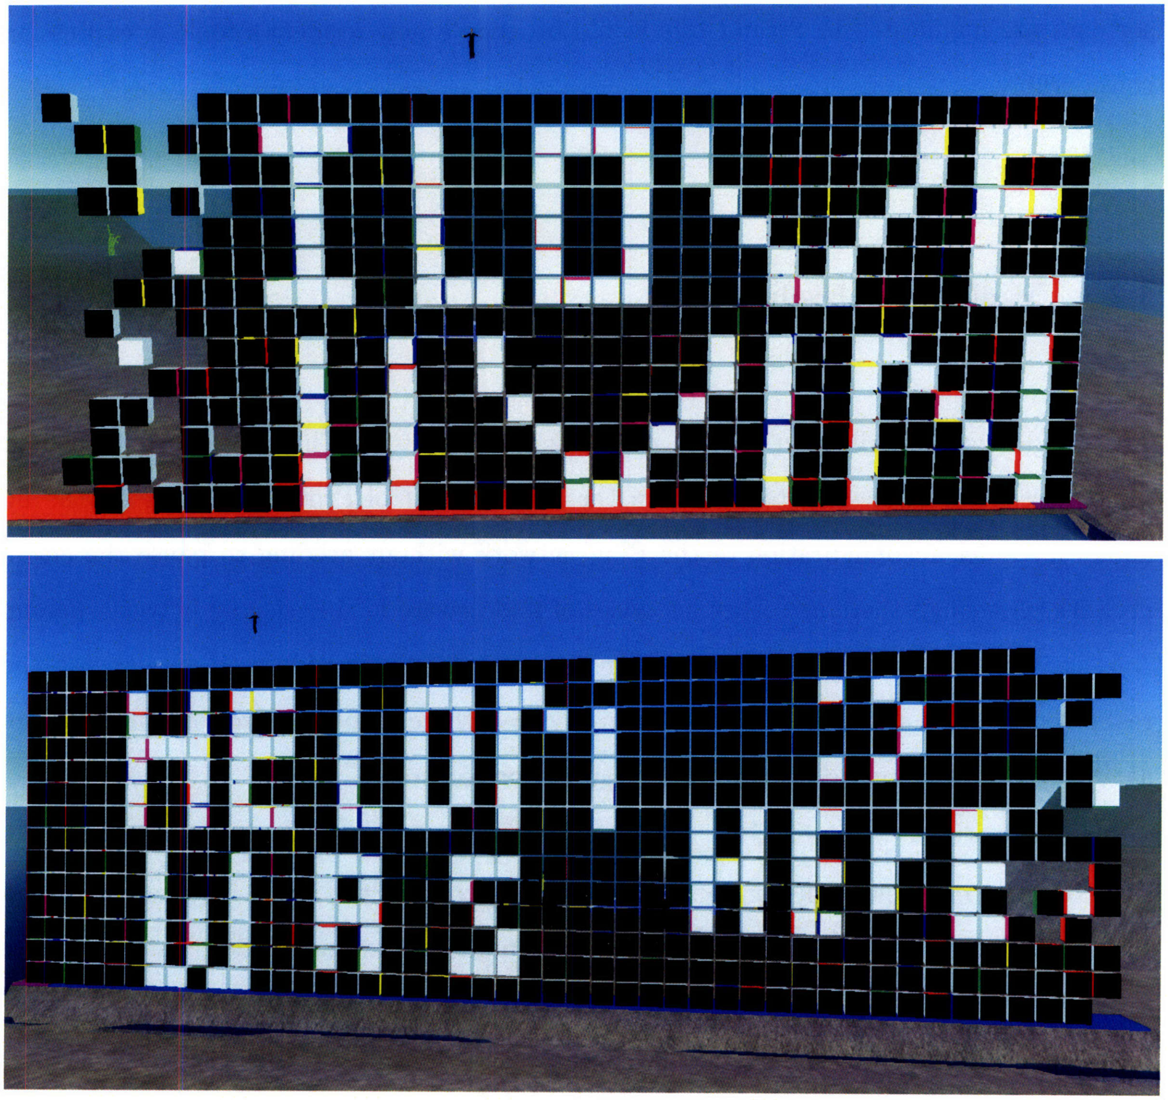

Figure 5-6: The reverse sides of the two virtual atrium walls after a visitor, presumably named Helori, strategically flipped tiles to spell out a large message, presumably for someone named Vin. Allowing anyone to flip tiles was originally intended to be a mechanism for hiding or showing the current video stream or image. An avatar floating above each wall gives a sense of scale.

a rapport with another *Second Life* resident. An avatar's name and basic appearance are the only characteristics a new *Second Life* resident can use to express themselves and shape their virtual persona. Allowing new residents to dictate how even basic real-world sensor data are manifested in *Second Life* would broaden a new resident's tools of expression considerably.

5.2.2 Middleware Evaluation

The middleware described in Section 4.3 works as expected. The ability toconnect and disconnect basestation PCs to and from the Master Control Program (MCP) on the fly allows for quick reconfiguration of the Plug network without needing to pause the entire network. Although there are other instances of middleware for interfacing between *Second Life* and external servers, very few are concerned with real-time interactions between the real physical world and the virtual world. In any case, there is no standard set of middleware against which to compare performance or even architecture.

The complete middleware system, including the basestation PCs, MCP, Apache server, and *Second Life* data server, can be set up relatively quickly on standard hardware using free and readily available software components. Actual instances of middleware deployments have remained stable for as long as they were needed, which is on the order of weeks at a time at the longest. There has not yet been an instance of the middleware unexpectedly failing. At present, the primary performance limitation is the restriction of updating *Second Life* from the MCP at a maximum rate of once per second. This limitation could easily be overcome at the expense of a more complex framework within *Second Life* involving multiple *Second Life* data servers belonging to distinct avatars. In principle, this limitation could be overcome to the extent that the next bottleneck would most likely be the data rate of the Plug wireless network, an improvement on which would require considerable work modifying the underlying wireless protocol or radio hardware.

5.2.3 Dual Reality Open House Evaluation

The dual reality open house described in Section 4.2.7 incorporates many of the other projects described in Section 4.2, and serves as a good test of the real data and virtual data collection systems. Sensor data from both the real world and virtual world were collected during the day-long event. The real-world data originated from the Plug sensor nodes used throughout the Media Lab atrium at the various demo stations described in Section 4.2.7. Table 5.1 identifies the particular Plug used for each of the demos. These identifiers are used throughout Figures 5-7 – 5-10. All data shown in these figures were retrieved from the Plugs' SD cards, as described in the portion of Section 4.1.1 detailing the sensor logger software module. The virtual-world data originated from the virtual sensing system detailed in Section 4.2.2 as deployed throughout the virtual portion of the dual reality open house described in Section 4.2.7.

Such an extensive data set from a single event spread across both the real world and the virtual world had not previously been collected. By the nature of the event and its presentation in each world, very little correlation between the real and virtual data was expected. However, each data set does speak to how people interact within each world separately and what the possibilities are for using data from one world in the other.

The real-world sound data shown in Figure 5-7 clearly follows the structure of the event as attendees alternate between the atrium during break times and the auditorium during the conference talks. Similar structure, to varying degrees, is seen in the motion, vibration, and even electrical current data (Figures 5-10, 5-11, and 5-9, respectively). On the other hand, the light data (Figure 5-8) indicate physical location more than attendee activity – the Scratch, Story Cafe, and IBM Demo tables (Plugs 0x001, 0x0008, 0x0010) were all located in nearly direct sunlight modulated by cloud cover, whereas the Virtual Atrium Wall table (Plug 0x000C) was located in the path of an LCD projector.

Of the various data collected from the virtual world during the day-long event, Figure 5-12 shows the distribution over time of touch events (avatars touching a virtual object equipped with the virtual embedded sensing system) and avatar movement events (the

Plug UID	**Demo Station**	**Demo Description**
0x0018	Promise Server	an online tool for managing informal contractual relationships
0x0008	Story Cafe	a collaborative narrative generation tool in *Second Life* and on the web
0x0010	IBM Demo Table	a venue for the organizers of the conference to show their work in virtual worlds
0x0013	The Restaurant Game	a downloadable game designed to teach autonomous software agents based on the actions of real players
0x0001	Scratch	a programming language accessible to children
0x000F	Sousreality	a two-way portal between *Second Life* and the real world
0x000E	Information Spaces	a space in *Second Life* designed to facilitate business meetings
0x000B	Real Atrium Wall	a wall of tiles in *Second Life* that reflects the aggregate state of data collected from the real world
0x000C	Virtual Atrium Wall	a wall of tiles in *Second Life* that reflects the aggregate state of the data collected from the virtual world

Table 5.1: The data in Figures 5-7 – 5-10 were collected from the above Plugs and their associated demo tables during the IBM Virtual World Conference held in the atrium of the MIT Media Lab. Those demos with *Second Life* components were described in Section 4.2.7.

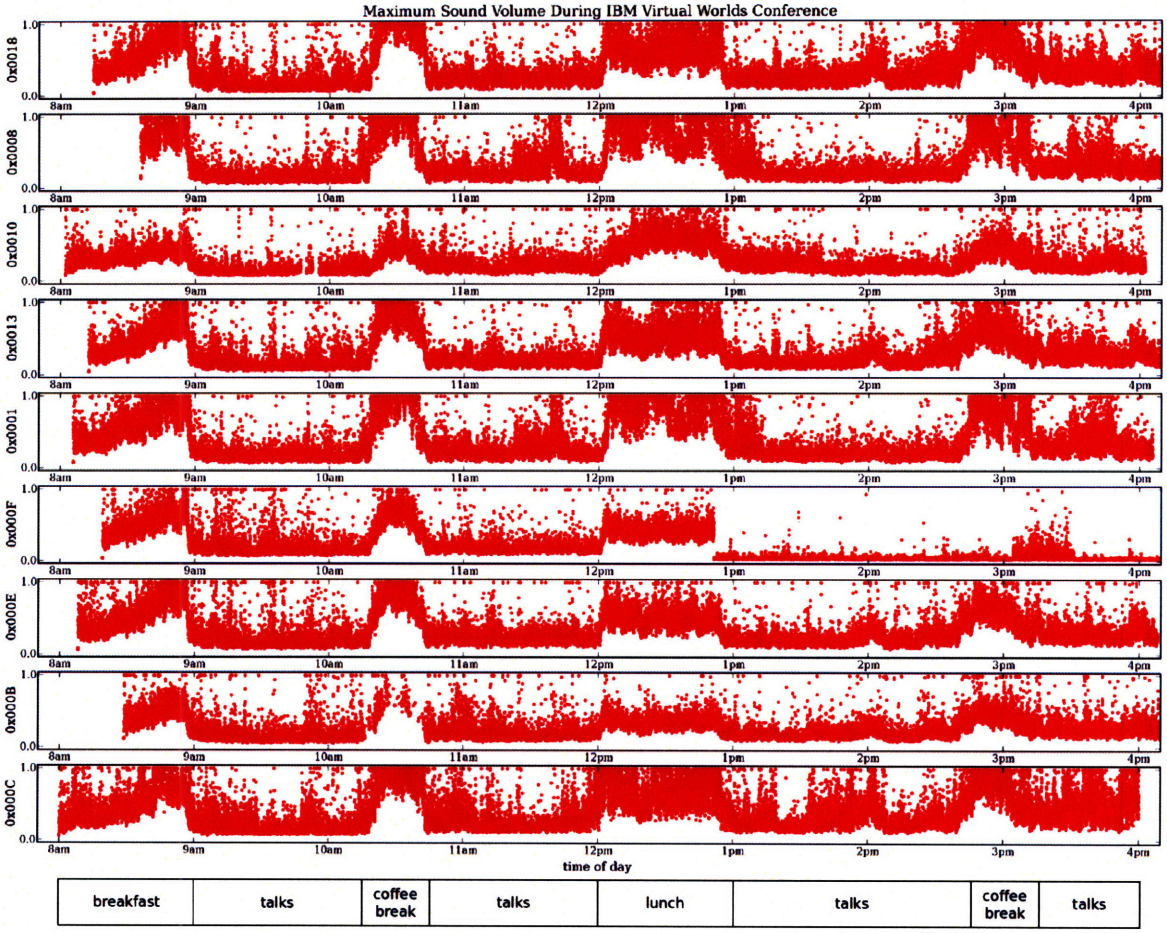

Figure 5-7: Sound data from each of the Plug nodes used at an atrium demo station at the real portion of the IBM Virtual Worlds Conference event. The y-axis of each graph is labeled with the unique identifier for the Plug node in question.

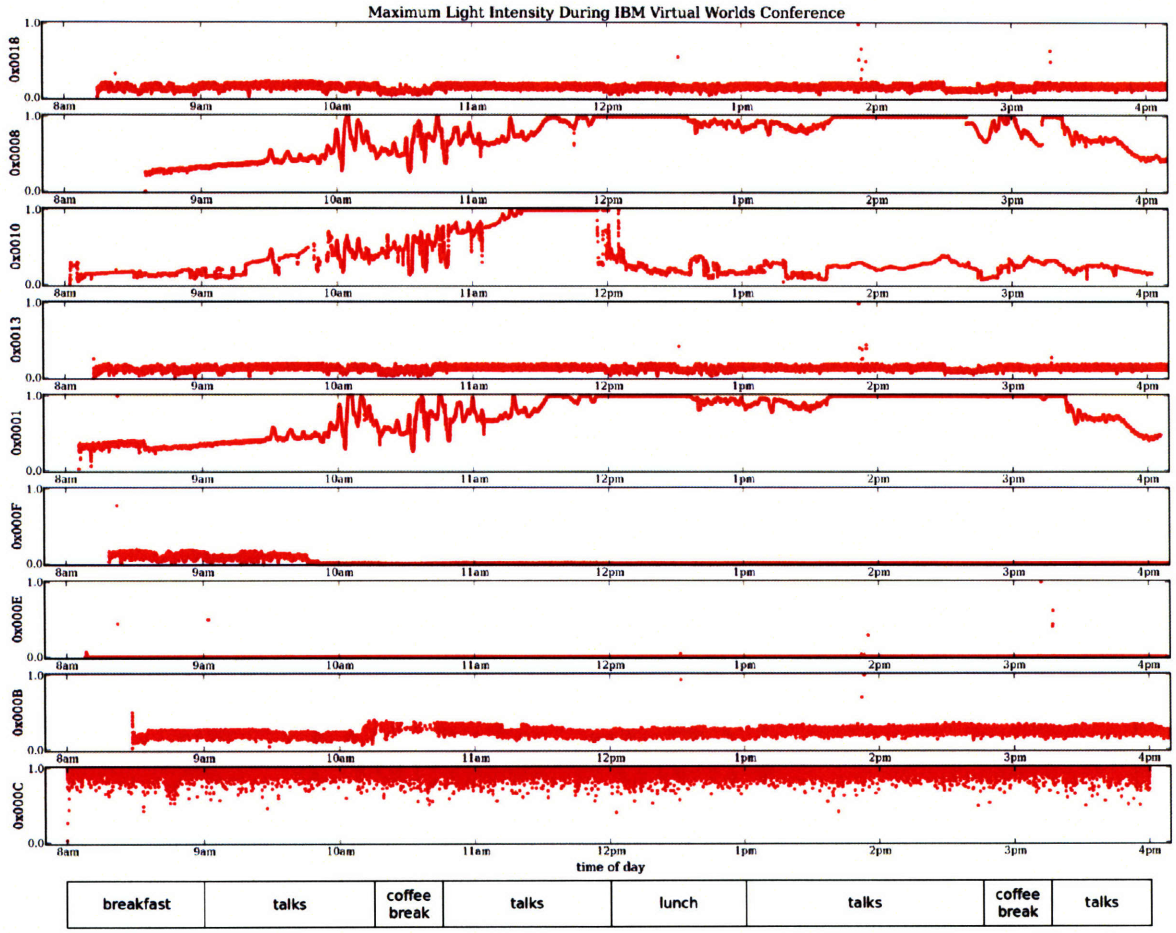

Figure 5-8: Light data from each of the Plug nodes used at an atrium demo station at the real portion of the IBM Virtual Worlds Conference event. The y-axis of each graph is labeled with the unique identifier for the Plug node in question.

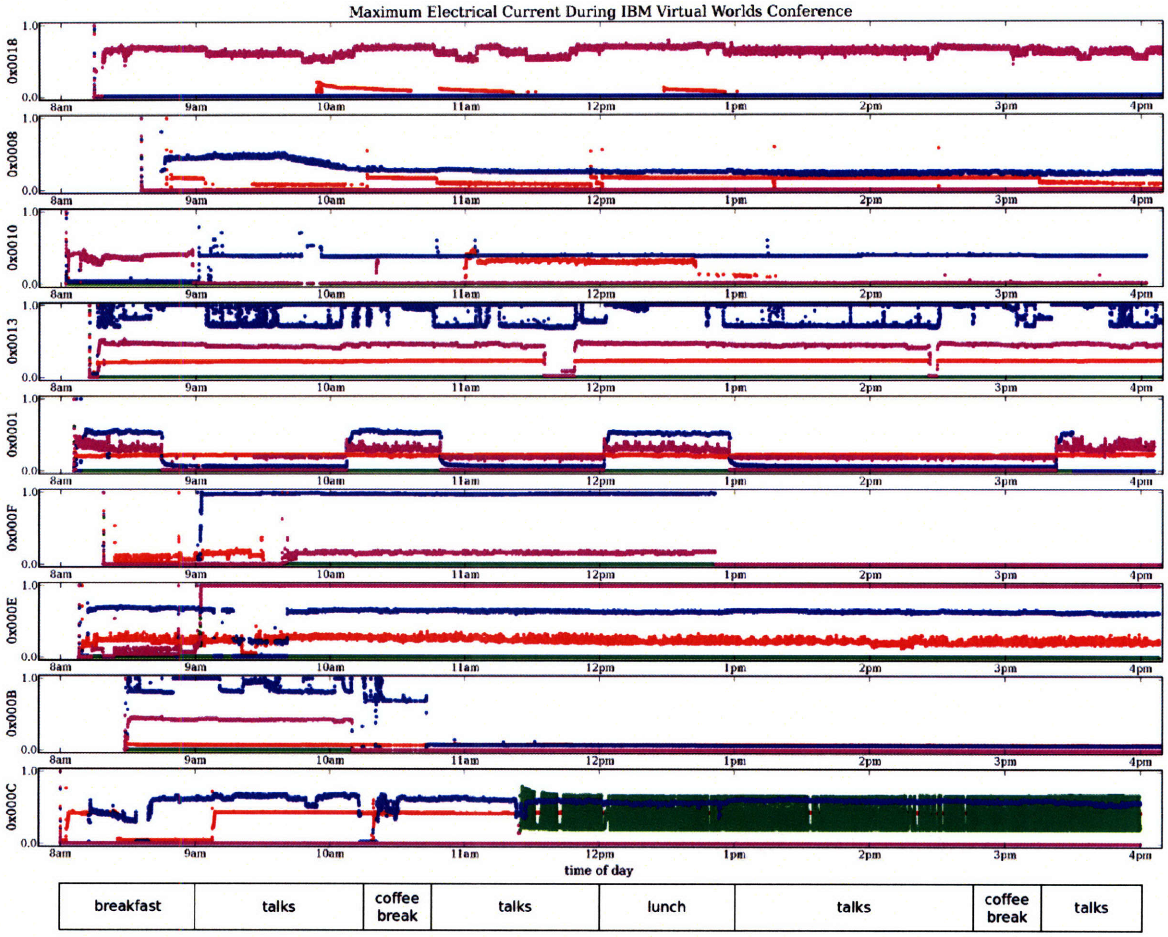

Figure 5-9: Electrical current data from each of the Plug nodes used at an atrium demo station at the real portion of the IBM Virtual Worlds Conference event. The y-axis of each graph is labeled with the unique identifier for the Plug node in question. Each color represents one of the four electrical sockets on a Plug.

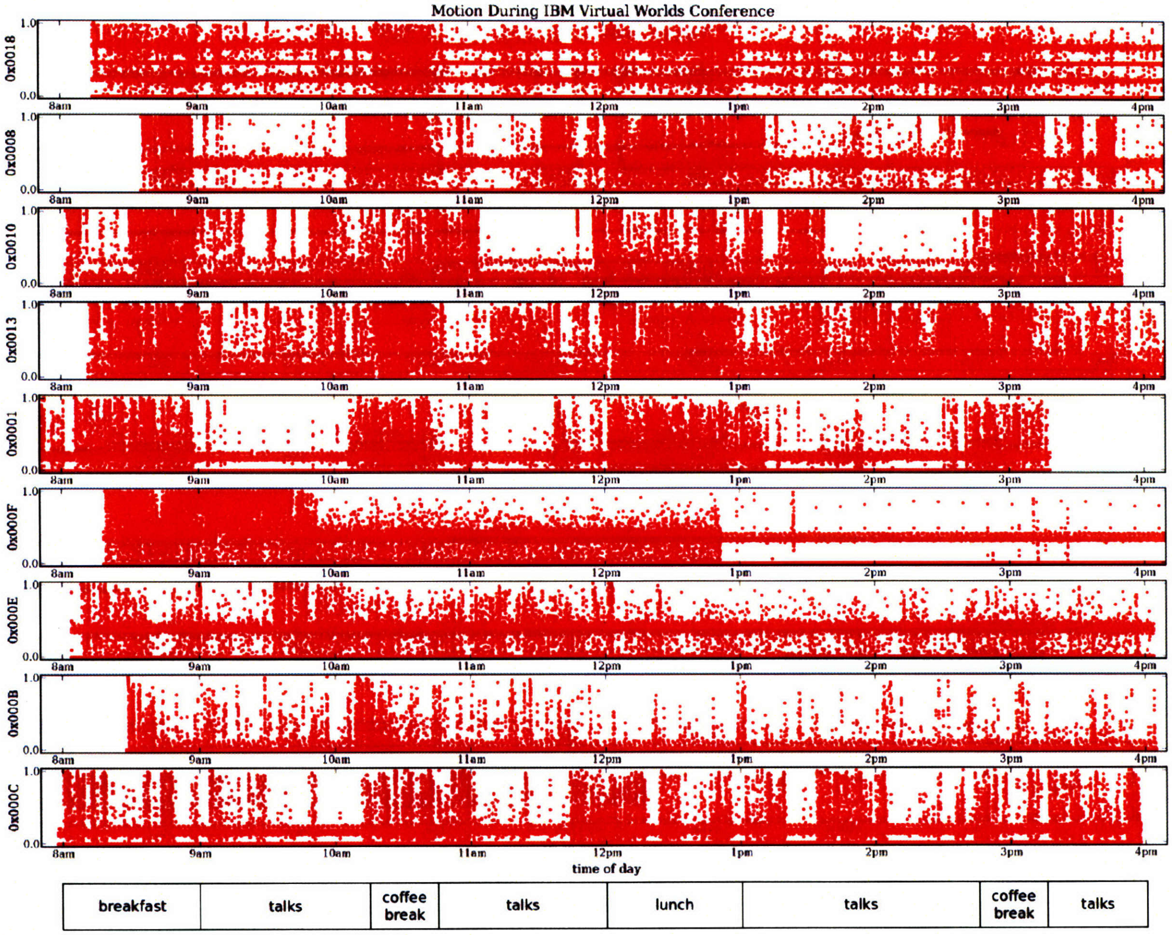

Figure 5-10: Motion data from each of the Plug nodes used at an atrium demo station at the real portion of the IBM Virtual Worlds Conference event. The y-axis of each graph is labeled with the unique identifier for the Plug node in question.

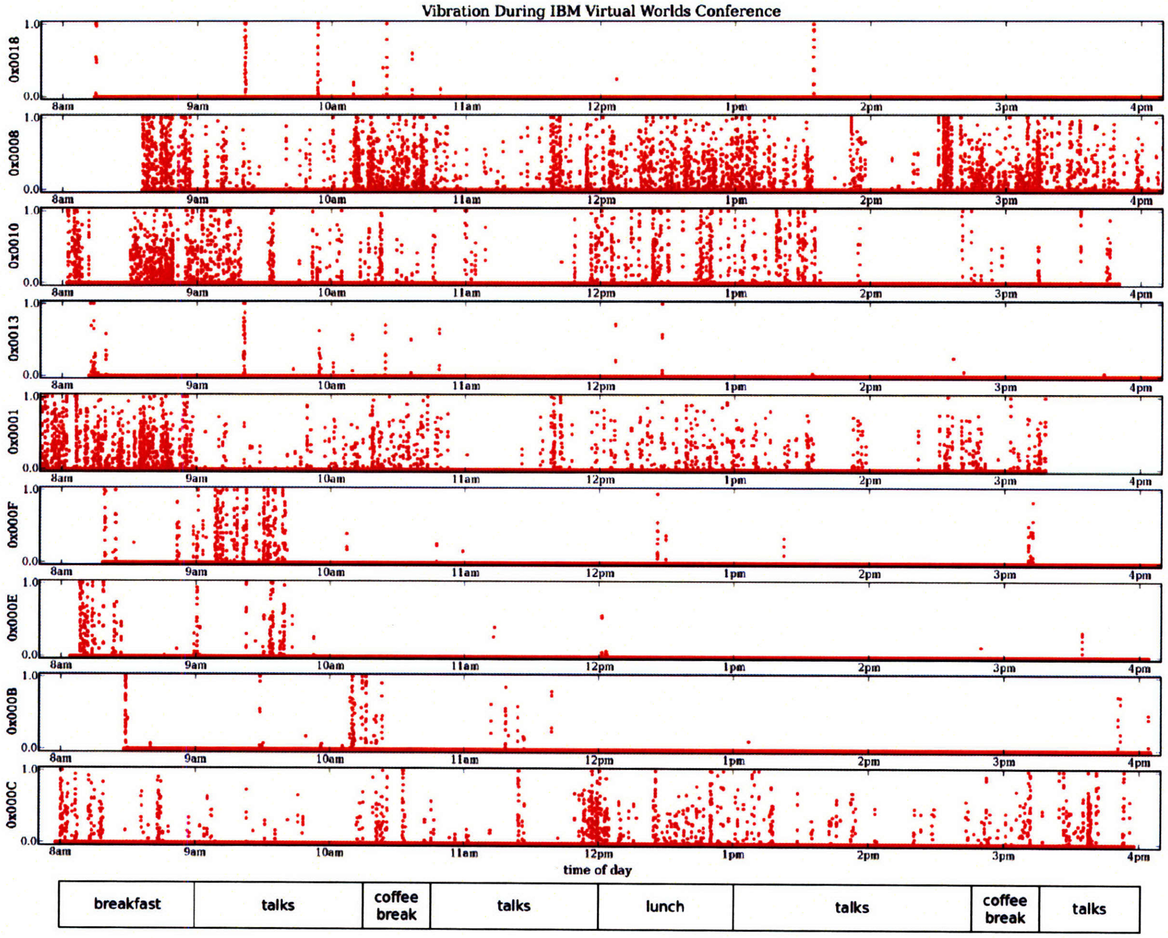

Figure 5-11: Vibration data from each of the Plug nodes used at an atrium demo station at the real portion of the IBM Virtual Worlds Conference event. The y-axis of each graph is labeled with the unique identifier for the Plug node in question.

Type of Object	**Number of Avatar Interactions**
Atrium Tiles	240
Information Icons	27
Data Ponds	12
Welcome Kiosks	11

Table 5.2: The types of virtual objects instrumented with a virtual sensing system and the number of times they were touched by avatars during the dual reality open house.

virtual wearable sensing system checks if its avatar is moving approximately once per second) collected from 22 avatars, of which 16 chose to wear the access bracelet virtual sensing system. Due to the server hosting the Master Control Program described in Section 4.3 accidentally being assigned the wrong IP address, data collected from the virtual sensing system started being logged at approximately 11 AM rather than at 8 AM, when the event actually started. The spike of avatar movement at around noon is likely due to the pause in the live video stream from the auditorium when the talks broke for lunch, thus giving avatars watching the video stream incentive to move to another location to interact with other aspects of the virtual space. The relatively constant motion thereafter might indicate the exploratory nature of the participants and/or the space. Table 5.2 lists the number of times the various types of virtual objects had interactions with avatars. The number of interactions indicates interest level to some extent, but also reflects the number of objects (e.g. there are thousands of atrium tiles, but only a dozen or so information icons). Not surprisingly, some combination of the streaming video and easily understood interaction mechanism (touch to flip) made the atrium the biggest draw, followed by the notecards provided by the information icons of the various virtual demo stations.

Although a valuable source of data from the real world and virtual world, the dual reality open house could have been improved in several respects. For example, the number of participating avatars could have been increased with better advertising; this is the type of event that would typically be popular in *Second Life* if only people had known about it. Also, a stronger connection between the real and virtual premises could have been made and "connectedness" metrics formulated and tested, even if this type of event isn't expected to engender much of a connection between the real and virtual worlds. In a sense, such a

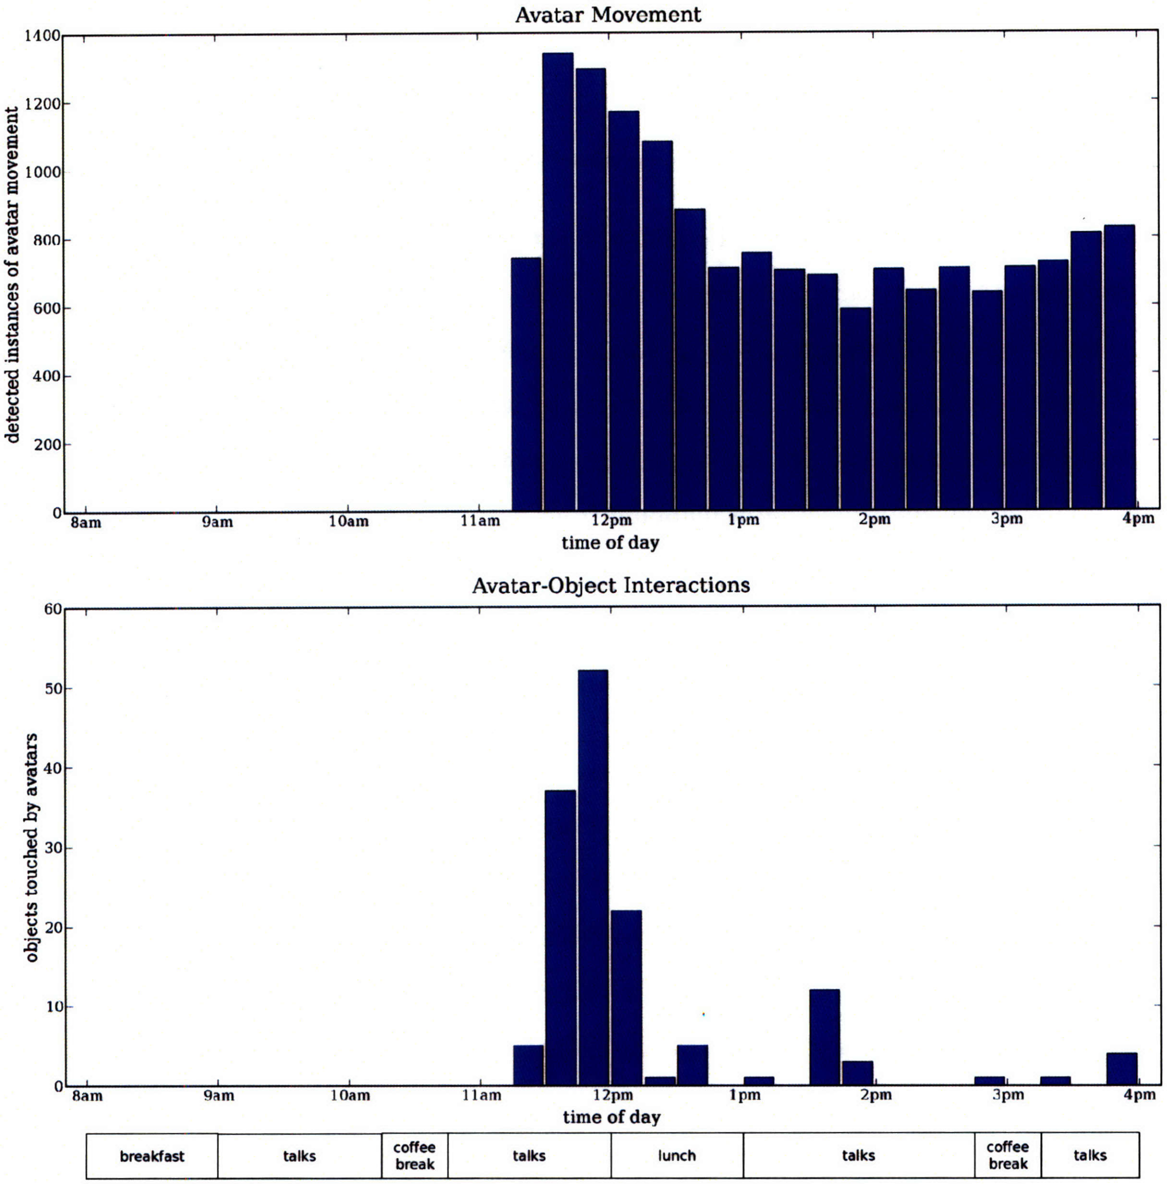

Figure 5-12: Some of the data collected from the *Second Life* portion of the IBM Virtual Worlds Conference event.

weak connection strengthens the argument that the primary use of sensor data from either world will not be for real-time, two-way, direct interaction, but rather as a more subtle means of communication between worlds.

5.3 Discussion

In a completely fabricated virtual world, the entropy of a real-world data stream can dramatically alter the virtual ambiance. Certainly, a cleverly utilized pseudo-random number generator could do the same, but meaning derives more from perception than from the underlying mechanism and it is much easier to weave a story from real data than from pseudo-random numbers. The act of weaving a story from sensor data is essentially the act of designing and implementing a mapping from data to a real or virtual manifestation of the data. A successful story must be meaningful to tell as well as to hear, and using sensor data grounded in either the real or virtual world helps achieve this. In essence, the act of creation must be as gratifying as the act of consumption.

The creative aspects of dual reality, the mapping of real or virtual sensor data to some manifestation, will likely follow the trend of another recent medium – blogs. While blogs have allowed some creative geniuses an outlet and given them a wide, appreciative, and well-deserved audience, the quality of most blogs, at least as a consumptive medium, is far below previous mass media standards. Of course, their quality as a creative medium and the value they bring to their creators in that regard far exceed previous standards by virtue of their relatively low barrier to entry alone. These trends will be exagerated in the context of dual reality for two reasons. First, the medium is much richer, involving virtual 3D worlds and complex social interactions and is therefore accessible to a wider audience. Second, once the mapping of data to manifestation is set, the act of creation is nearly automatic and therefore accessible to a wider range of talent. In short, the worst will be worse and the best will be better, a hallmark of successful mass media. As with other creative media, virtuosity will still play a critical role in dual reality, namely in the conception, implementation, and honing of the specific mappings between sensor data and their manifestations.

While mapping sensor data to manifestation may be at the highest level of the dual reality creative process, it not the only level by any means. Once the mappings are in place, people can still intentionally express themselves in many ways, depending on the exact nature of the mapping. The evolution of emoticons in text messages is one example of such expression using a current technology. Another is the habit of maintaining an active online presence, such as used in Internet messaging clients, by jogging the computer's mouse occasionally. In the same way, users of dual reality environments will modify their behavior so as to express themselves through the medium.

With mapping as a central pillar of the dual reality creative process, many parallels can be drawn with and lessons learned from the field of information visualization [111], especially those areas that deal specifically with visualizing information in virtual environments [112]. A complete survey of information visualization is beyond the scope of this work, but it is worth making clear the distinction between information visualization and dual reality. The most obvious difference is the inherent social dimension to dual reality; dual reality mappings exist to influence the venue and means of human interaction in the real and virtual worlds, whereas information visualization has a broader scope. Of course, the ultimate purpose of information visualization is to convey information to other people and therefore to influence them, so in a vague sense information visualization has a social dimension as well. The other important difference between information visualization and dual reality is the intent of the participants. It is the explicit intention of dual reality creators to express their personality and share pieces of their lives by mapping data from their lives into various manifestations, be it a grotesque metamorphosis of an avatar or the subtle shape shifting of a virtual space. There is a large overlap of information visualization and dual reality when it comes to manifesting data in virtual environments for the purpose of browsing, but the primary use of dual reality will be self-expression more along the lines of blogs, only much more varied and to a large extent automated.

As an example of these differences, consider a service, such as Google Maps [113] and Traffic.com [114], that visualizes traffic congestion in a large metropolitan area. Information about traffic congestion might be gathered from numerous sources, such as cell towers, arial imagery, or user input, and displayed in a variety of ways, such as on the web, in a 3D virtual environment, or text messaging. The primary use of this service is to allow participants to intelligently plan their daily commute. Although hardly social by most standards, this service does form a social feedback loop; a user of the service will change her route according to the data presented and in doing so change the nature of the data presented to the next user. However, the motivation or intent of the service is entirely devoid of self-expression, and therefore does not readily fall under the rubric of dual reality.

Chapter 6

Conclusion

"Conan, what is best in life?"

Khitan General

The greatest contribution and driver of media technology is the enabling, in steps, of consumption, communication, and creation of media by masses of people. Dual reality promises to fulfill the potential of the last of these steps by bringing together the technology of online virtual worlds and sensor networks. This dissertation lays out a philosophical framework guided by self-expression by which to begin this process and demonstrates a complete implementation of dual reality technologies. However, this is only the beginning of something that will take some time to fully realize and will likely differ in surprising ways from what is postulated here.

6.1 Contributions

As laid out in Section 1.3, this dissertation makes the following contributions:

- The framing of dual reality as a new application domain for sensor networks, virtual worlds, and media creation.

- State-of-the-art hardware and software tools for realizing dual reality interactions.
- Example dual reality environments, interactions, and scenarios.
- A new human-centric approach to sensor network design that addresses practical issues of pushing sensor networks to ubiquity, as needed to realize dual reality environments.
- Evaluation and discussion of various aspects of such dual reality systems.

6.2 Future Work

The following are only selected exerpts from a long list of possible new directions and improvements to this work, some of which are already under way.

Regarding hardware, the Plug platform's communication capabilities in particular could be improved. For example, although not as energy constrained as most sensor network platforms, the Plug platform may yet benefit from some hybrid of energy efficient and always-on MAC layers [115]. Such an approach would more closely mirror the heterogeneous nature of the platform, with small battery-powered devices roaming among a fixed powered network of Plugs. Even incorporating into the Plug an IEEE 802.11 (Wifi) radio is a possibility that shouldn't be overlooked. The robustness of the network, a necessity for long-term deployment, would benefit from a more agile channel sharing and link determination algorithm.

An obvious and significant improvement to the Plug platform would be the addition of some form of power line communication to integrate the Plugs even more tightly with their deployment environment [116, 117, 118]. Even low-bandwidth power line communication could provide a valuable side channel for network discovery, network maintenance, and even sensing – certain electrical failures could be detected and isolated by network means alone. Just as network connectivity in a wireless network reveals information about the surrounding environment, such as rough localization, so to does the network connectivity in a wired power line network. Of course, any power line communication would have to be tolerant of highly variable line noise and would have to rely on another communication channel, such as wireless, to bridge between different electrical phases or feeds, as are common in large buildings. Between standard Wifi and power line communication, the only reason for the Plugs to have the lightweight 2.4GHz radio is to communicate with power constrained sensor nodes that have no other means of communicating. Nonetheless, this is a compelling reason given the centrality of mobile sensor nodes in the currently accepted vision of ubiquitous computing.

On a more abstract level, equipping living and working spaces with a host of sensors raises many unaddressed privacy concerns. The most common concern centers on the Plug's embedded audio microphone. However, there is evidence to suggest that the motion sensors coupled with other sensing modalities are possibly more invasive [119, 120]. This is especially true given the already near-ubiquity of microphones embedded in desktop computers, laptops, cell phones, and hand held devices. Clearly, the questions of privacy involved with ubiquitous sensor networks do not yet have well-defined answers and deserve closer inspection.

Perhaps the most pressing work yet to be done with regard to the Plug platform is the automatic detection, classification, and estimation of high-level contexts and events from low-level sensor data. Without the need to conserve battery life, the Plugs are free to continuously monitor their environment and perform such calculations collaboratively within the network rather than offline.

The ideas behind the Tricorder handheld sensor network browser can be taken further as well. As a proof-of-concept, the next generation of the Nokia 770 used in the Tricorder (the Nokia N800 [121]), which includes a built-in video camera and microphone, was used to stream live audio and video into *Second Life* using the GStreamer framework [122]. This proof-of-concept was the first instance of streaming a live audio and video feed from a Nokia Internet tablet into *Second Life*. From a technical perspective, the integration of this feature with the Tricorder should be straightforward. From an application perspective, it might be part of the rich connection between the real and virtual needed, for example, to allow a person in the real world to give a meaningful tour to those in the virtual, where those in the virtual world could follow the tour guide's real-world progress as well as explore a virtual reflection of the place being toured. Used in this way, the Tricorder becomes a sensor as well.

With regard to online virtual worlds, there are broad technical challenges in making them more appealing to a wider audience, let alone more usable for dual reality applications. While *Second Life* might today be the best option for dual reality research, there is plenty of room for improvement and a variety of other virtual worlds are cropping up that may be more suitable. Aside from incremental improvements to physics and rendering engines, a better way of exchanging data between the servers running the virtual world and other servers is of particular importance to dual reality applications. Generic data sharing services, such as Swivel and the SensorBase project at CENS, might help pave the way for this [123, 124, 125].

Just as a relatively standard set of metaphors and data mappings were developed for the personal computer (e.g. a desktop containing windows), it might come about that a standard set of metaphors arises for dual reality applications. In any case, classifying and cataloging the various successful ways to map sensor data to real or virtual manifestations is sure to take a lot of experimentation.

Much of the work presented in this dissertation concerns using real-world sensor data in the virtual world. In part, this is because sensing the real world is generally simpler than actuating upon the real world (e.g. the real data pond comprising a fan and a flexible plastic sleeve). As the quality and variety of real-world actuators improves, there will be more opportunity to explore using virtual-world sensor data in the real world. The Topobo kinetic memory sculpture toys are a good example of a candidate actuation technology [126].

6.3 Final Thoughts

There may not be a natural end point to the development and evolution of dual reality, but some significant success can be claimed once the virtual world can inspire as potently as the real world; when the metaphor carried by the virtual wind does more than cause unaging trees to randomly sway and instead sparks something akin to the vast bodies of art the real wind has already ignited. Perhaps then dual reality will be a concrete way of experiencing what the French philosopher and architect Gaston Bachelard described when he wrote [127]:

> *The cleverer I am at miniaturizing the world, the better I possess it. But in doing this, it must be understood that values become condensed and enriched in miniature thinking. One must go beyond logic in order to experience what is large in what is small.*

Appendix A

Plug Hardware

This appendix details the printed circuit boards used in a Plug node. In total, four printed circuit boards were used – a high-voltage circuit board for handling all 120VAC signals (Figures A-1 and A-2), a low-voltage circuit board for handling all digital logic and microcontroller signals (Figures A-3 and A-4), a radio circuit board that attaches directly to the low-voltage circuit board (Figures A-5 and A-6), and a data log expansion module that attaches to the low-voltage circuit board through a slot in the mechanical case (Figures A-7 and A-8).

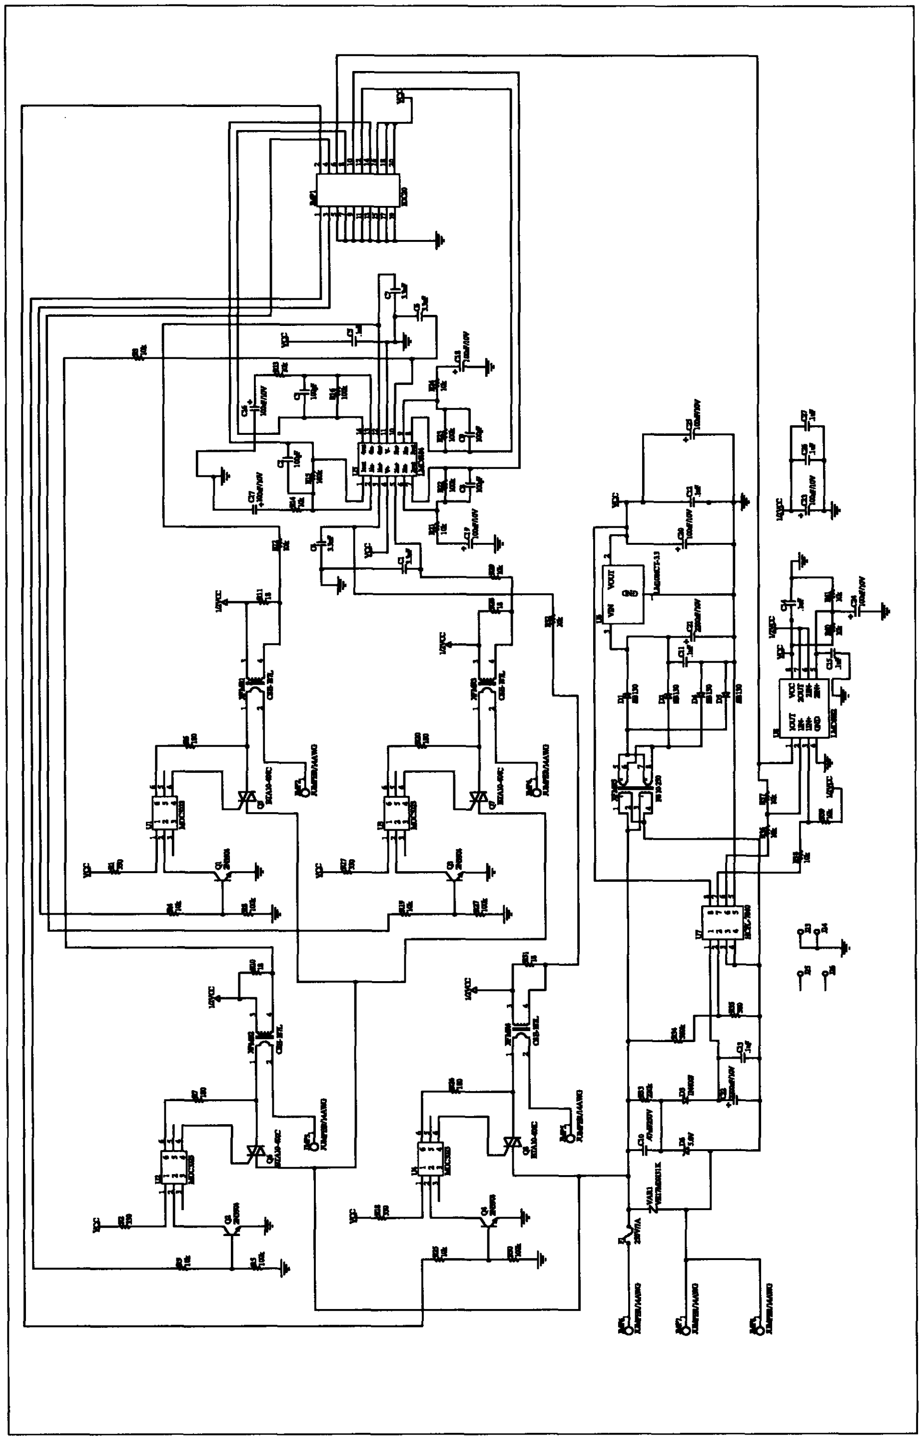

Figure A-1: The schematic for the Plug's high-voltage circuit board.

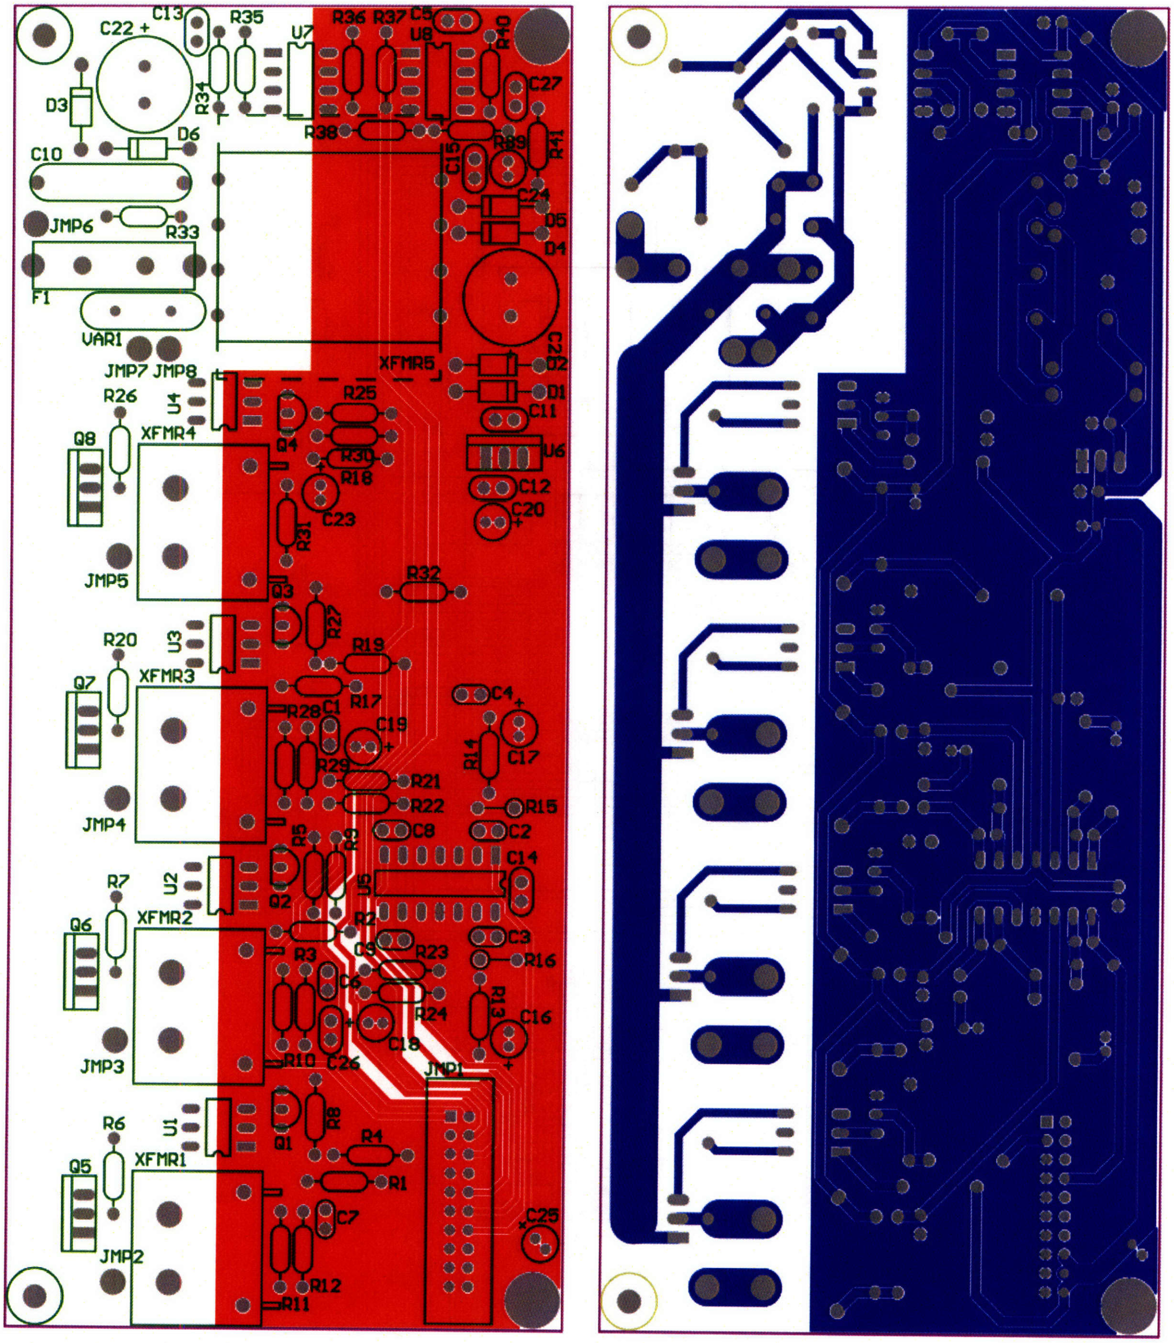

Figure A-2: The top layer (left) and bottom layer (right) views of the Plug's high-voltage circuit board.

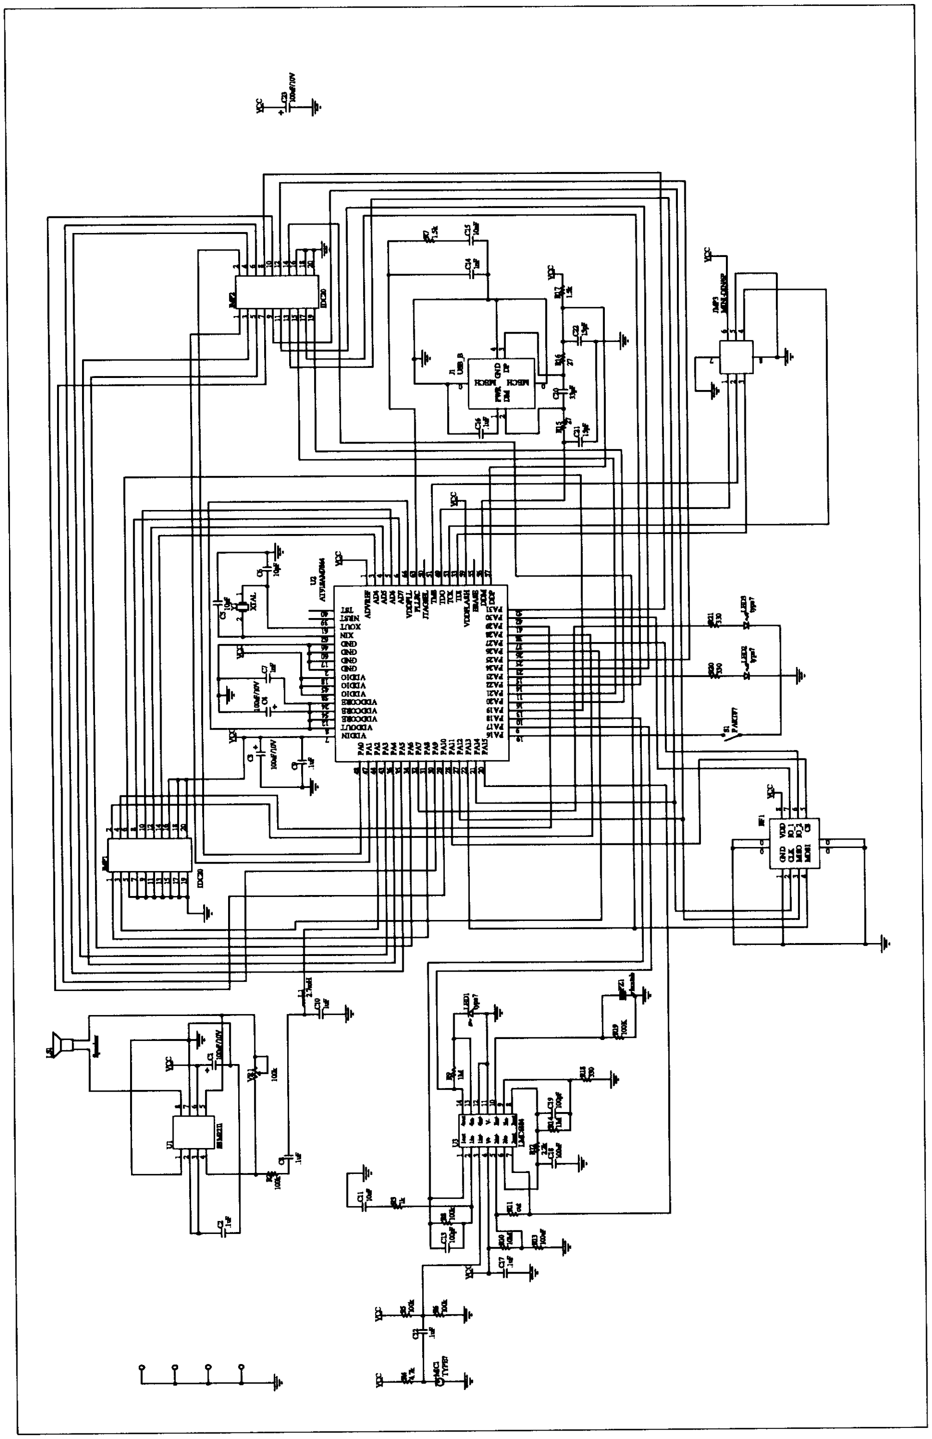

Figure A-3: The schematic for the Plug's low-voltage circuit board.

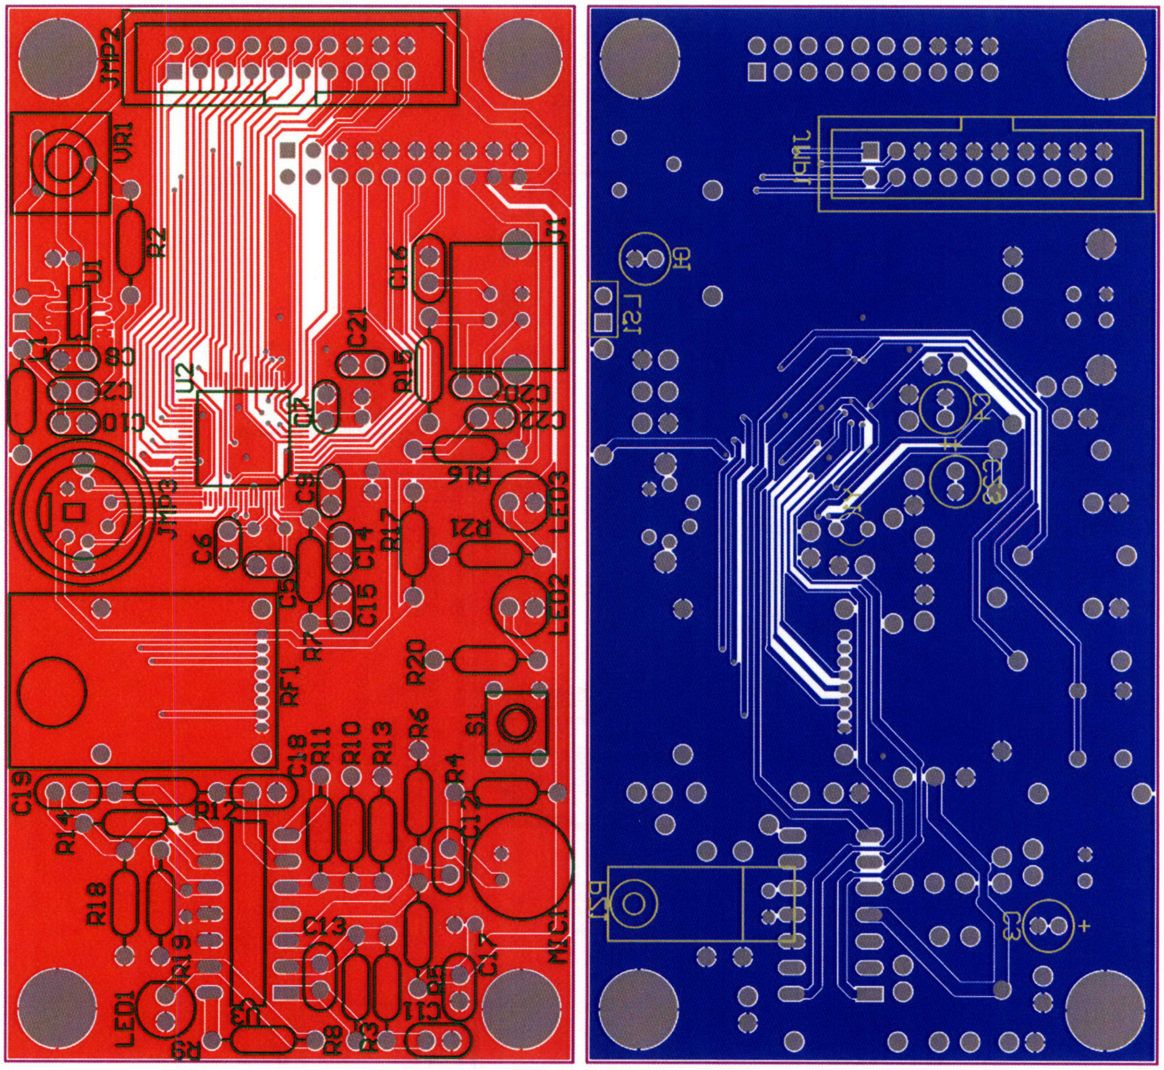

Figure A-4: The top layer (left) and bottom layer (right) views of the Plug's low-voltage circuit board.

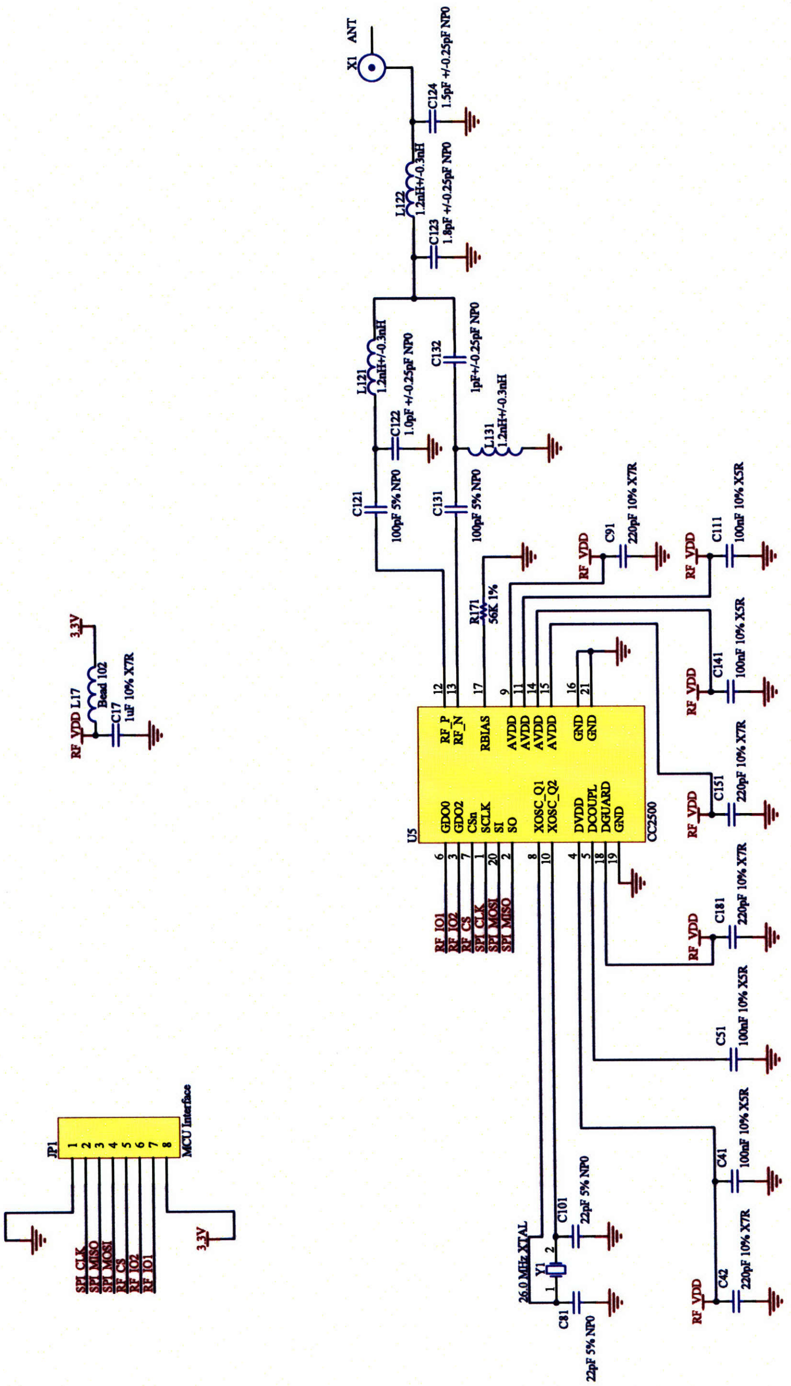

Figure A-5: The schematic for the Plug's radio circuit board.

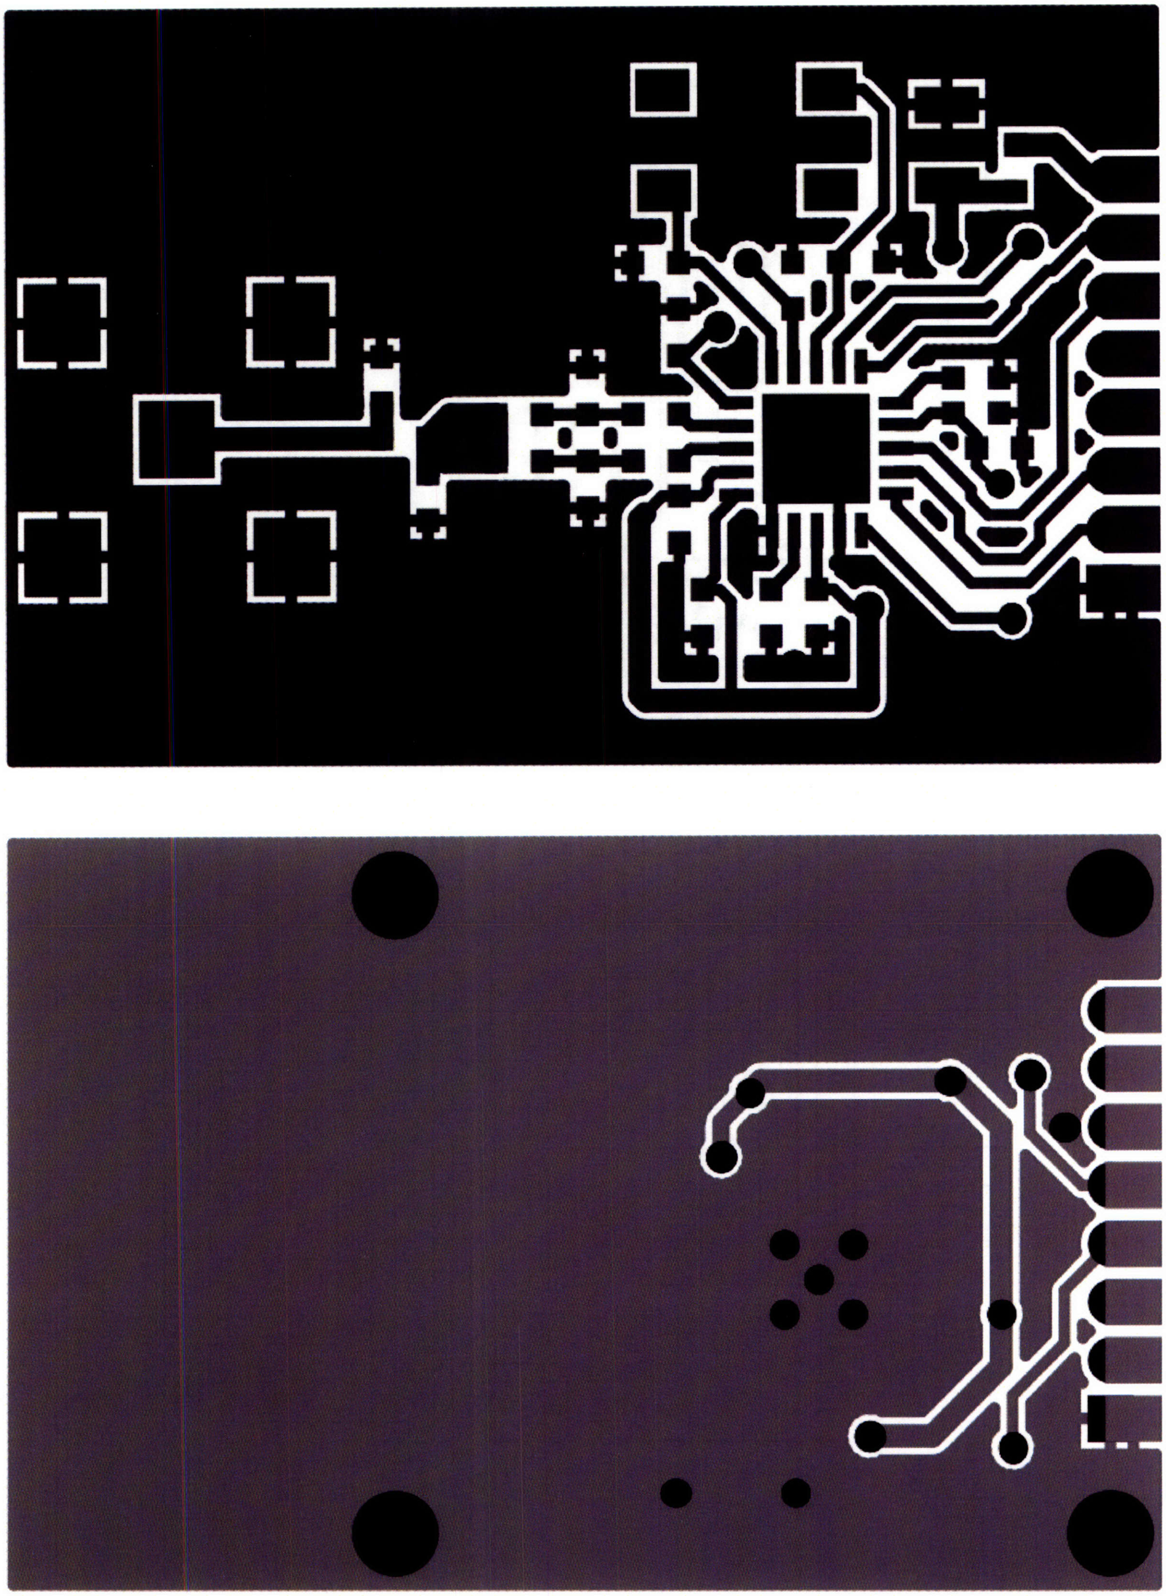

Figure A-6: The top layer (top) and bottom layer (bottom) views of the Plug's radio circuit board.

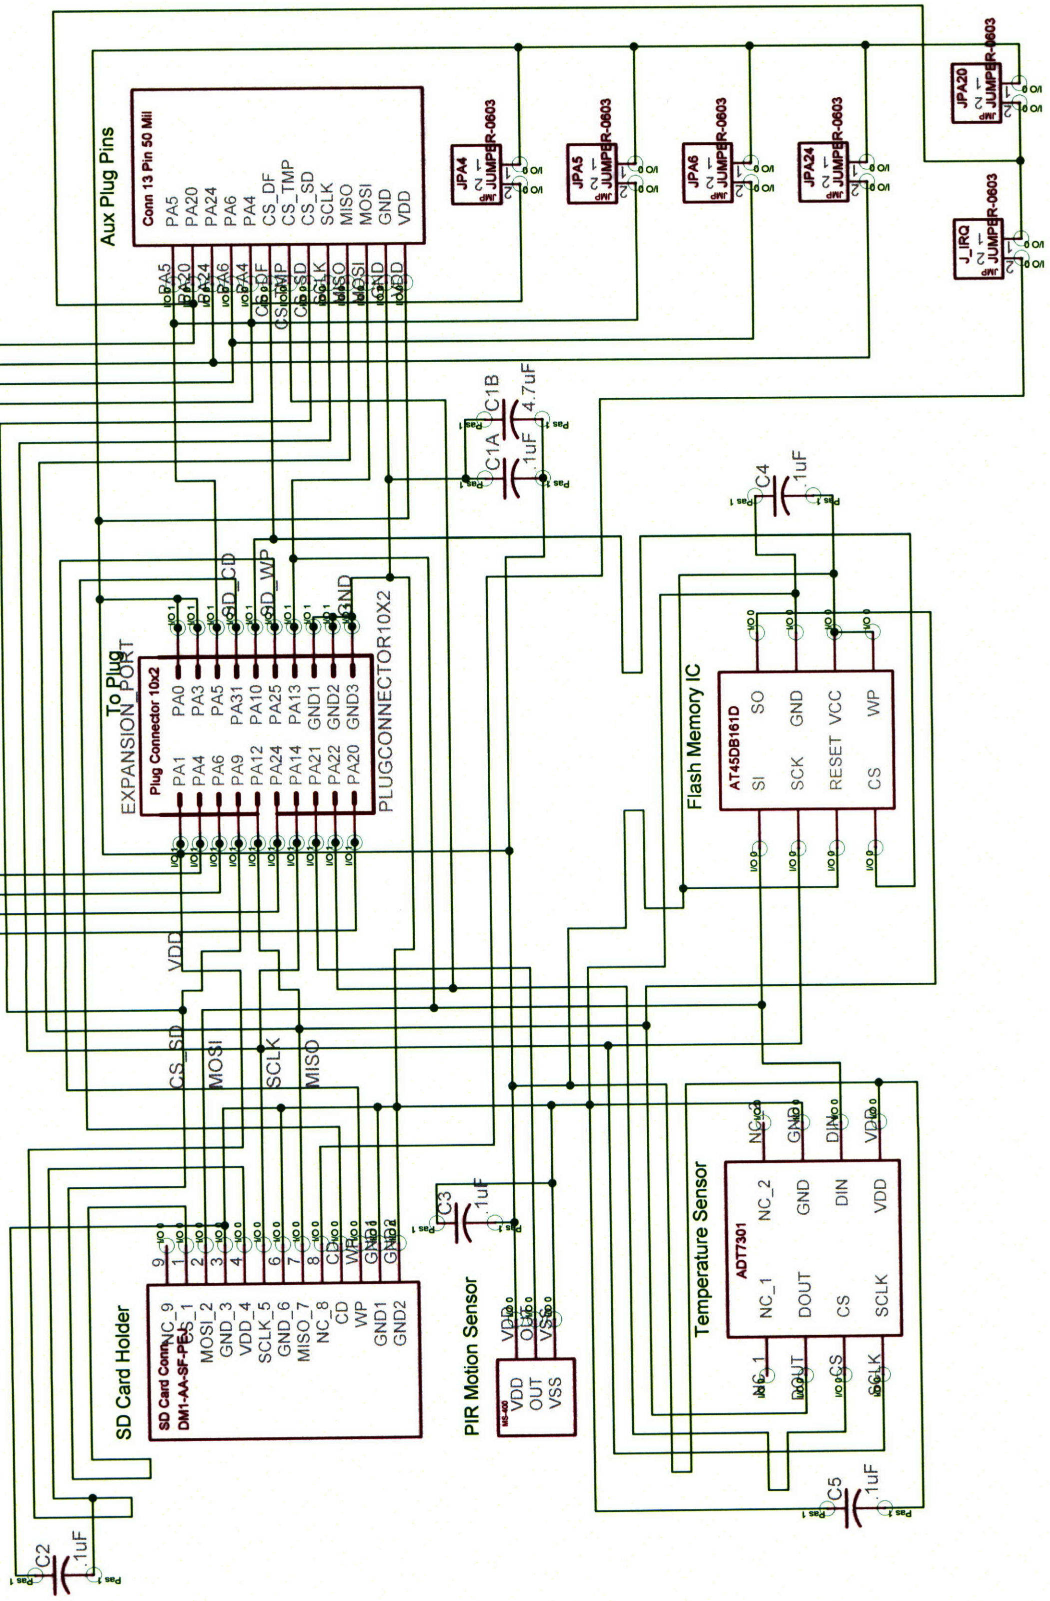

Figure A-7: The schematic for the Plug's data log expansion circuit board.

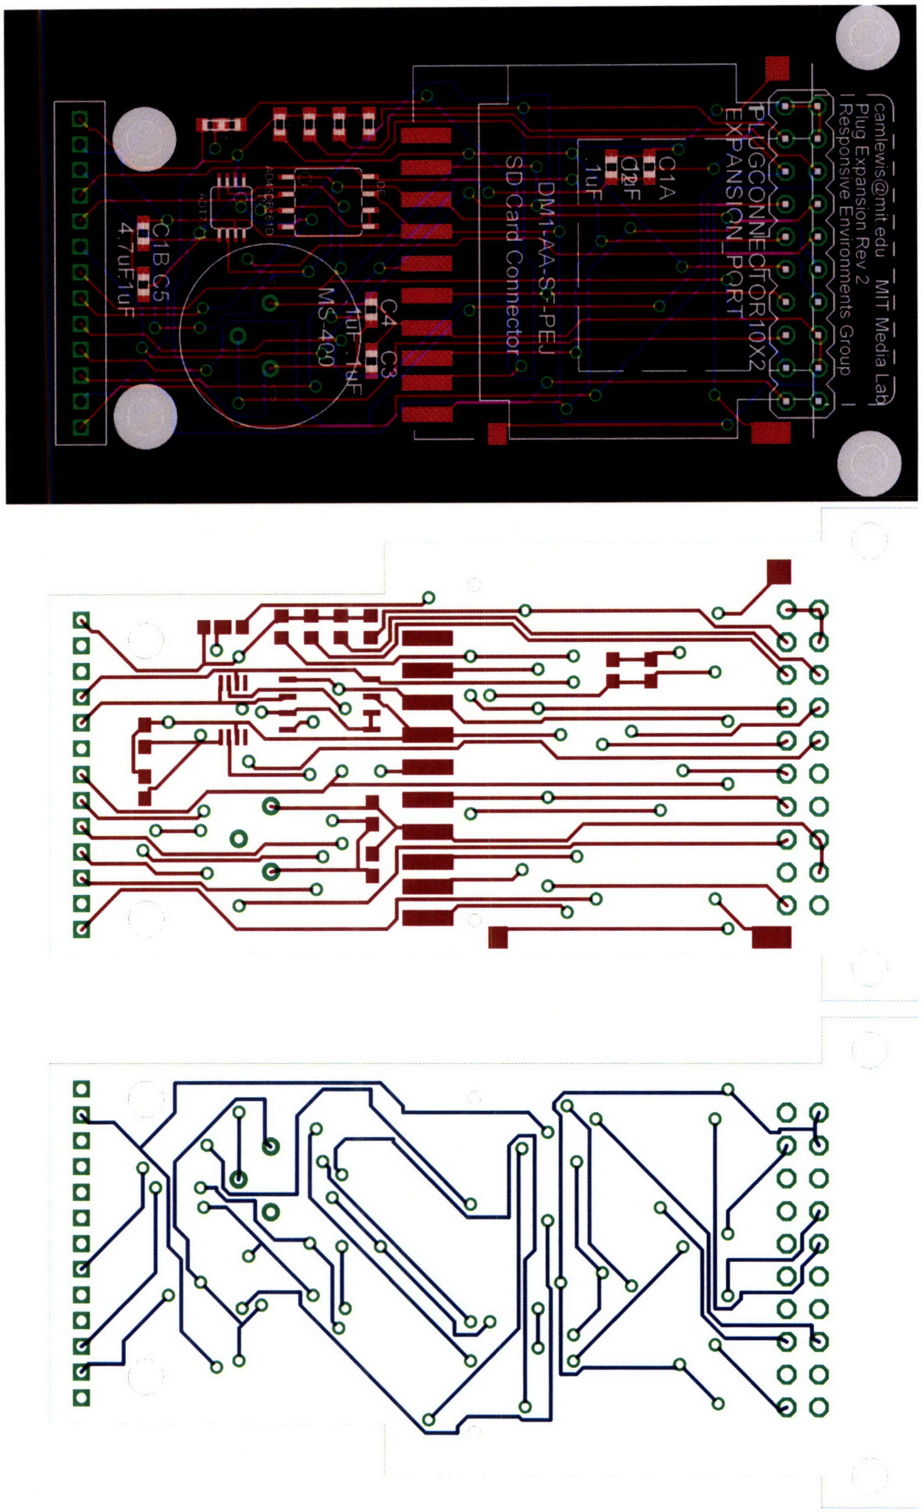

Figure A-8: The layout of the Plug's data log expansion module, including the top view of the superposition of the top, bottom, and silk screen layers (top), the top view of the top copper layer (middle), and the bottom view of the bottom copper layer (bottom).

Appendix B

Plug Firmware

The Plug firmware described in Section 4.1.1 was developed using Rowley Associates' CrossWorks for ARM version 1.5 for Linux tool chain [128]. The following is a listing of all source files used in the final implementation, all of which are available at `http://mlforge.media.mit.edu/cgi-bin/viewcvs.cgi/code/clean/?cvsroot=plugpoint`. The middleware used in conjunction with this firmware is available at `http://mlforge.media.mit.edu/cgi-bin/viewcvs.cgi/code/plugpc/?cvsroot=plugpoint`.

- adc.c, adc.h
- AT91SAM7S64.h
- button.c, button.h
- cc2500.c, cc2500.h
- clock.c, clock.h
- data_link.c, data_link.h
- data_log.h
- expansion.c, expansion.h

- filesystem.c, filesystem.h
- leds.c, leds.h
- main.c
- neighbors.c, neighbors.h
- pioa.c, pioa.h
- random.c, random.h
- route.c, route.h
- sensors.c, sensors.h
- sequence.h
- shadow_lab.c, shadow_lab.h
- sounds.c, sounds.h
- speaker.c, speaker.h
- spi.c, spi.h
- spi_dataflash.c, spi_dataflash.h
- spi_sd_card.c, spi_sd_card.h
- spi_temperature.c, spi_temperature.h
- switches.c, switches.h
- system.c, system.h
- types.c, types.h
- uid.h
- usb.c, usb.h
- usb_commands.c, usb_commands.h

Appendix C

Sensor Network Application Retasking Framework (Snarf)

An early portion of the work leading up to this dissertation consists of a framework for programming sensor networks within ubiquitous computing environments. In particular, a framework was designed and partially implemented to be:

- *distributed*: The middleware should be accessible to many simultaneous users and programmers from anywhere in the network, achieved by having each node maintain its own copy of a virtual machine that runs numerous user programs. User programs, or agents, are written for single virtual machines rather than the network as a whole and can be introduced to the network from any of the user-accessible nodes. This trades off the power of centralized programming for the flexibility of distributed programming.
- *dynamic*: The middleware should allow for incremental run-time updates to the applications it is hosting and accept applications that were unforeseen at the time the middleware was installed. This includes the ability to simultaneously run multiple programs, add new programs, and delete existing programs, all on the fly. Again, the nature of the virtual machine allows for these sort of dynamics.

- *expressive*: The middleware should allow the programmer to take full advantage of the capabilities of the sensor network, not just those needed for a specific application. The virtual machine allows for this by trading off the efficiency of an application-specific instruction set for the generality of an instruction set generic with respect to sensor network applications.

- *usable*: The middleware should provide a high-level programming language that is easy to learn and use by moderately experienced programmers. To this end, the programming language used to write source code will be modeled on Python, a high-level traditional programming language known for its simplicity and easy learning curve. That said, there should be significant deviations and additions to compensate for the particular nature of sensor network programming. For example, there should be built-in mechanisms for easily dealing with asynchronous network and sensing events.

The framework proposed to fulfill these criteria is called *Snarf* (Sensor Network Application Retasking Framework) and consists of the following components:

- Snarl (Sensor Network Application Retasking Language): a high-level programming language designed for ease-of-use for novice and advanced programmers.

- Snarc (Sensor Network Application Retasking Compiler): a compiler for translating Snarl source code into bytecode to be run on SnarfVM.

- SnarfVM (Sensor Network Application Retasking Virtual Machine): an interpreter for executing bytecode produced by Snarc. Designed to run on 8-bit microcontrollers, yet powerful enough to be useful on more capable machines.

- SnarfIDE (Sensor Network Application Retasking Framework Integrated Development Environment): a program and user interface for coordinating, controlling, and monitoring all aspects of Snarf, including uploading code to a network.

The Snarf project was inspired by and built upon PyMite, an open source project aimed at bringing a pared down version of the Python programming language to microcontrollers [129]. At the time of Snarf's inception, PyMite was thought to be defunct. However, a year or so into the Snarf project, PyMite resumed active development, at which point much of the work carried out on Snarf was merged back into PyMite and Snarf development was frozen. The two projects are similar enough to maintain close ties, but different enough to warrant separate branches.

Appendix D

Lug Hardware

This appendix details the Lug's printed circuit board schematic (Figure D-1) and layout (Figure D-2).

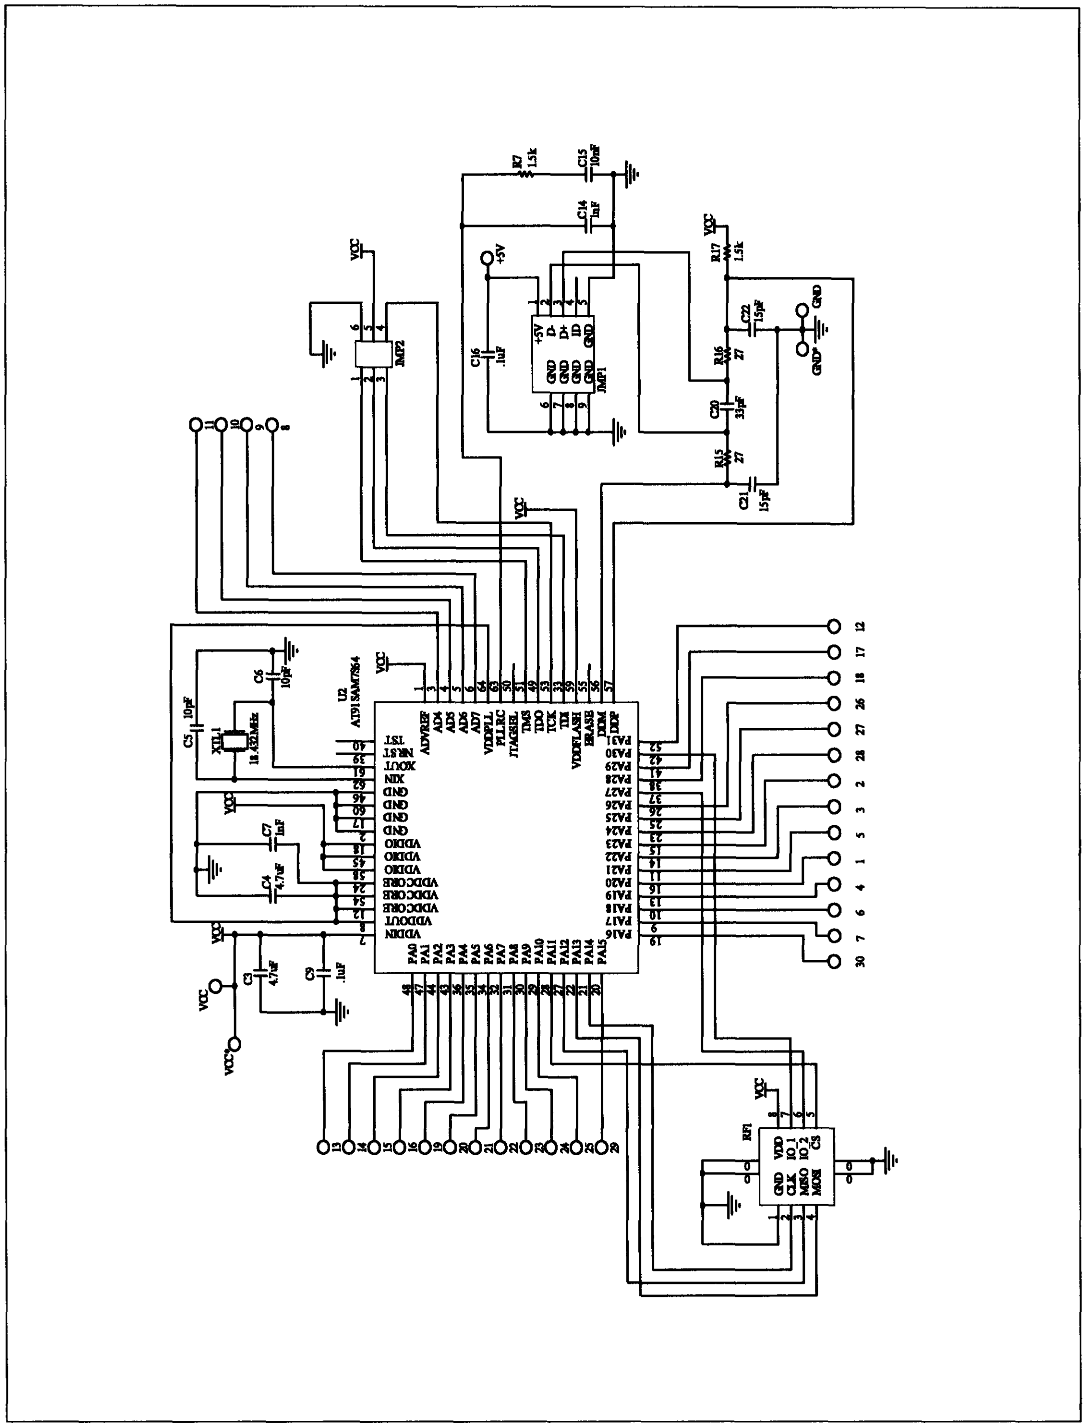

Figure D-1: The schematic for the Lug's circuit board.

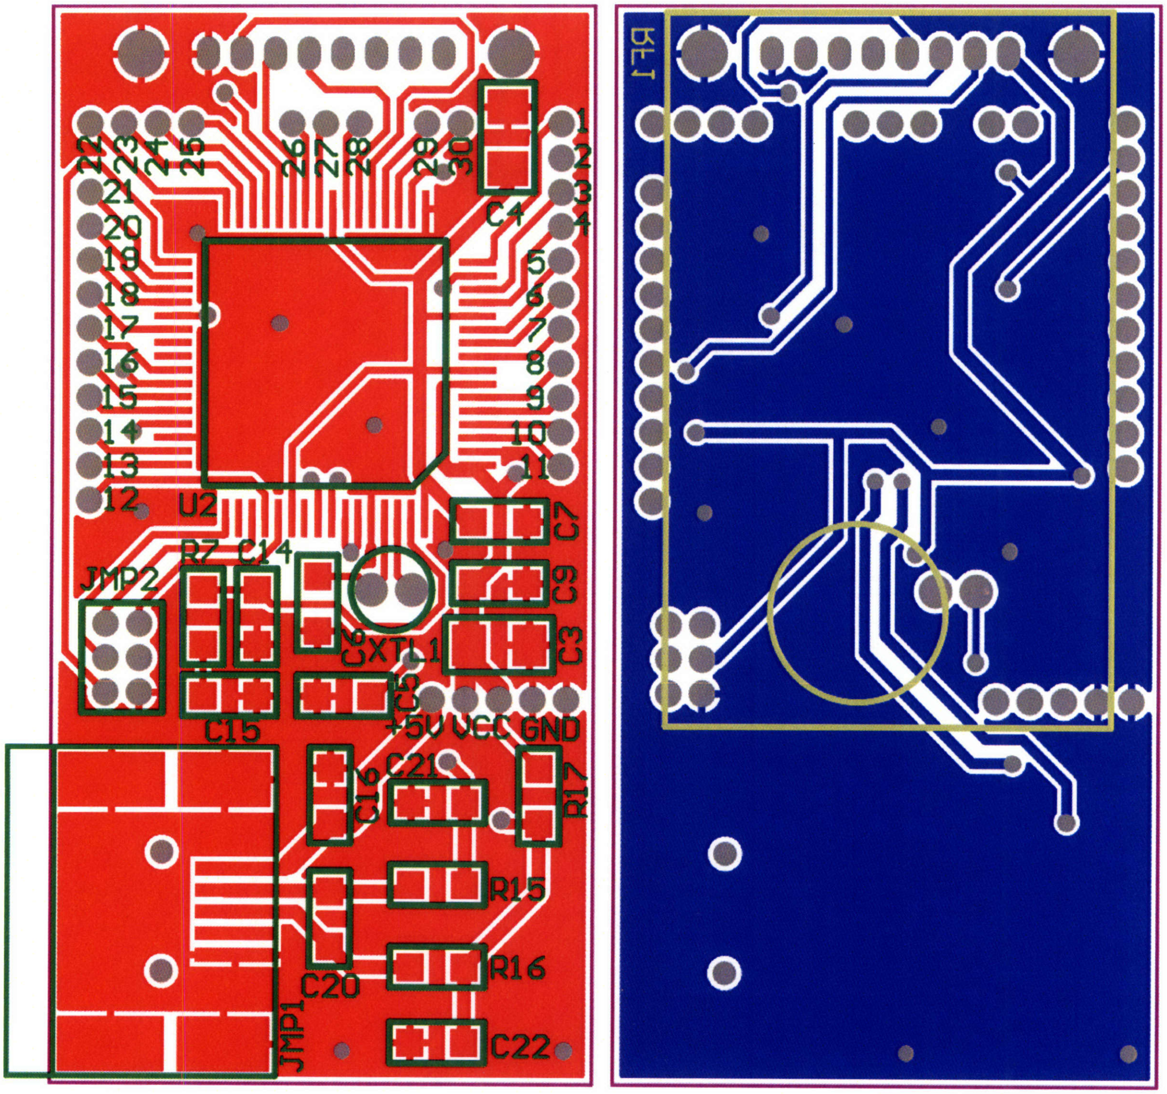

Figure D-2: The top (left) and bottom (right) views of the Lug's circuit board layout.

Appendix E

LSL Example Code

The following Linden Scripting Language (LSL) script was used in the access bracelet described in Section 4.2.2.

```
// Global Constants
string REGION_NAME = "Ruthenium";
list TELEPORT_POSITIONS = [<73,73,23>, <73,55,23>, <55,55,23>, <55,73,23>];
float PING_INTERVAL = 5.0;
integer ACCESS_CHANNEL = -302;
string TRACKER_URL = "http://thirdlife.media.mit.edu/cgi-bin/tracker.py";
string CHAT_URL = "http://thirdlife.media.mit.edu/cgi-bin/chat.py";
integer listen_handle;


// Global Variables
key wearer = NULL_KEY;
integer is_active = FALSE;


// Send a message to the tracking server.
ping()
{
  llRegionSay(ACCESS_CHANNEL, llGetOwner());
  llHTTPRequest(TRACKER_URL, [], "");
}


// Determine if the bracelet should send messages to the tracking server.
update_pinger()
{
  if ((!is_active) && (wearer == llGetOwner()) && (llGetRegionName() == REGION_NAME))
  {
    is_active = TRUE;
    ping();
    llSetTimerEvent(PING_INTERVAL);
    listen_handle = llListen(0, "", llGetOwner(), "");
    llOwnerSay("Activated!␣␣See␣the␣notecard␣included␣within␣this␣bracelet␣for␣details.");
  }
  else if (is_active)
  {
    is_active = FALSE;
    llListenRemove(listen_handle);
    llSetTimerEvent(0);
    llOwnerSay("Deactivated!␣␣See␣the␣notecard␣included␣within␣this␣bracelet␣for␣details.");
  }
}
```

```
default
{

  listen(integer channel, string name, key id, string message)
  {
    // Record everything the wearer says.
    llHTTPRequest(CHAT_URL, [HTTP_METHOD, ”POST”], message);
  }


  on_rez(integer start_param)
  {
    // Reset the state and let the attach event take care of initialization.
    is_active = FALSE;
  }

  changed(integer change)
  {
    // Respond to region change events.
    if (change & CHANGED_REGION)
    {
      update_pinger();
    }
  }

  attach(key id)
  {
    // The bracelet was just put on an avatar.
    wearer = id;
    update_pinger();
  }

  timer()
  {
    // Periodically send a message to the tracking server
    ping();
  }

  touch_start(integer total_number)
  {
    // Teleport the wearer to the region upon being touched.
    if (llDetectedKey(0) == llGetOwner())
    {
      llMapDestination(REGION_NAME,
                       llList2Vector(TELEPORT_POSITIONS,
                       (integer)llFrand(llGetListLength(TELEPORT_POSITIONS))) +
                         <llFrand(2), llFrand(2), 0>,
                       ZERO_VECTOR);
    }
  }

}
```

Bibliography

[1] William Gibson. *Neuromancer.* Ace Books, 1984.

[2] Neal Stephenson. *Snow Crash.* Bantam Books, 1992.

[3] Wikipedia. Berkeley Building — Wikipedia, The Free Encyclopedia, 2007. [Online; accessed 24-July-2007].

[4] Ambient Devices. Ambient Orb. `http://www.ambientdevices.com`, 2007.

[5] Harold E. Edgerton and James R. Killian. *Flash!: Seeing the Unseen by Ultra High-speed Photography.* Hale, Cushman & Flint, 1939.

[6] Harold E. Edgerton and James R. Killian. *Moments of Vision : The Stroboscopic Revolution in Photography.* MIT Press, 1979.

[7] Roland Gelatt. *The Fabulous Phonograph: From Tin Foil to High Fidelity.* Lippincott, 1955.

[8] Deborah Estrin, Ramesh Govindan, John S. Heidemann, and Satish Kumar. Next Century Challenges: Scalable Coordination in Sensor Networks. In *Mobile Computing and Networking*, pages 263–270, 1999.

[9] Distributed Sensor Networks. Carnegie-Mellon University workshop proceedings, December 1978.

[10] William M. Merrill, Lewis Girod, Brian Schiffer, Dustin McIntire, Guillaume Rava, Katayoun Sohrabi, Fredric Newberg, Jeremy Elson, and William Kaiser. Dynamic Networking and Smart Sensing Enable Next-Generation Landmines. *IEEE Pervasive Computing Magazine*, pages 82–89, October–December 2004.

[11] Joseph Paradiso, Joshua Lifton, and Michael Broxton. Sensate Media - Multimodal Electronic Skins as Dense Sensor Networks. *BT Technology Journal*, 22(4):32–44, October 2004.

[12] Christopher Laughman, Kwangduk Lee, Robert Cox, Steven Shaw, Steven Leeb, Les Norford, and Peter Armstrong. Power Signature Analysis. *IEEE Power & Energy Magazine*, March–April 2003.

[13] W.K. Lee, G.S.K. Fung, H.Y. Lam, F.H.Y. Chan, and Mark Lucente. Exploration on Load Signatures. In *Proceedings of the International Conference on Electrical Engineering (ICEE)*, 2004.

[14] M. Ito, R. Uda, S. Ichimura, K. Tago, T. Hoshi, and Y. Matsushita. A method of appliance detection based on features of power waveform. In *Proceedings of the International Symposium on Applications and the Internet*, pages 291–294, 2004.

[15] Mohamed El Hachemi Benbouzid. A Review of Induction Motors Signature Analysis as a Medium for Faults Detection. *IEEE Transactions on Industrial Electronics*, 47(5), October 2000.

[16] L.K. Norford, S.B. Leeb, D. Luo, and S.R. Shaw. Advanced Electrical Load Monitoring: A Wealth of Information at Low Cost. In *Proceedings of Diagnostics for Commercial Buildings: from Research to Practice*. Pacific Energy Center, June 1999.

[17] George W. Hart. Residential Energy Monitoring and Computerized Surveillance via Utility Power Flows. *IEEE Technology and Society Magazine*, June 1989.

[18] Joshua Daniel Kaufman. SeeGreen: A Tool For Real-time Distributed Monitoring of Home Electricity Consumption. Master's thesis, MIT Media Lab, May 2001.

[19] P3 International. Kill A Watt. `http://www.p3international.com`.

[20] Consumer Powerline. Home Joule. `http://www.consumerpowerline.com/homejoule/`, 2007.

[21] Spy Labs. AGS-01: GSM Transmitter Concealed in a Surge Protector. `http://www.spy-labs.com/infinity.htm`.

[22] Chip PC Technologies. JackPC. `http://www.chippc.com/thin-clients/jack-pc/`.

[23] Product Design Forums. Swivel Socket - Dynamic Surge Protector. `http://yankodesign.com/product_info.php?products_id=690`, January 2006.

[24] W. Steven Conner, John Heidemann, Lakshman Krishnamurthy, Xi Wang, and Mark Yarvis. Workplace Applications of Sensor Networks. Technical Report USC/ISI Technical Report ISI-TR-2004-591, Intel Research and Development and University of Southern California, Information Sciences Institute, 2004.

[25] Ishwinder Kaur. OpenSpace: Enhancing Social Awareness at the Workplace. Master's thesis, MIT Media Lab, May 2007.

[26] Bret Hull, Vladimir Bychkovsky, Yang Zhang, Kevin Chen, Michel Goraczko, Allen K. Miu, Eugene Shih, Hari Balakrishnan, and Samuel Madden. CarTel: A Distributed Mobile Sensor Computing System. In *4th ACM SenSys*, November 2006.

[27] Jonathan Lester, Tanzeem Choudhury, and Gaetano Borriello. A Practical Approach to Recognizing Physical Activities. In *4th International Conference on Pervasive Computing*, May 2006.

[28] Apple Inc. Nike + iPod. `http://www.apple.com/ipod/nike/`, 2006.

[29] Ubicomp: International Conference on Ubiquitous Computing. `http://www.ubicomp2007.org/`, September 2007.

[30] Mark Weiser. The Computer for the 21st Century. *Scientific American*, 265(3):94–104, 1991.

[31] Roy Want. Building a Sensor-rich World. *Pervasive Computing: Mobile and Ubiquitous Systems*, 6(2), April–June 2007.

[32] David Brin. *The Transparent Society: Will Technology Force Us To Choose Between Privacy and Freedom?* Basic Books, 1999.

[33] Robert Szewczyk, Alan Mainwaring, Joseph Polastre, John Anderson, and David Culler. An Analysis of a Large Scale Habitat Monitoring Application. In *Proceedings of the 2nd International Conference on Embedded Networked Sensor Systems (SenSys)*, pages 214–226, 2004.

[34] Phil Buonadonna, David Gay, Joseph M. Hellerstein, Wei Hong, and Samuel Madden. TASK: Sensor Network in a Box. In *European Workshop on Sensor Networks*, January 2005.

[35] M. Maroti, G. Simon, A. Ledeczi, and J. Sztipanovits. Shooter Localization in Urban Terrain. *Computer*, 37(8):60–61, August 2004.

[36] Seth Teller, Kevin Chen, and Hari Balakrishnan. Pervasive Pose-Aware Applications and Infrastructure. *IEEE Computer Graphics and Applications*, July/August 2003.

[37] J. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, and M.B. Srivastava. Participatory Sensing. In *Proceedings of the World Sensor Web Workshop at SenSys*, 2006.

[38] Marion Hermersdorf, Heli Nyholm, Jukka Salminen, Henry Tirri, Jukka Perkiö, and Ville Tuulos. Sensing in Rich Bluetooth Environments. In *Proceedings of the World Sensor Web Workshop at SenSys*, 2006.

[39] Steven Feiner, Blair Macintyre, and Dorée Seligmann. Knowledge-based Augmented Reality. *Communications of the ACM*, 36(7):53–62, 1993.

[40] Jim Youll. Wherehoo and Periscope: A Time & Place Server and Tangible Browser for the Real World. In *Extended Abstracts of the Proceedings of the Conference on Human Factors in Computing Systems (CHI)*, pages 109–110, 2001.

[41] William J. Mitchell and Federico Casalegno. Electronic Lens. `http://mobile.mit.edu/elens/`, 2006.

[42] Rainer Simon, Harald Kunczier, and Hermann Anegg. Towards Orientation-Aware Location Based Mobile Services. In *Proceedings of the 3rd Symposium on LBS and TeleCartography*, November 2005.

[43] Howard Rheingold. *The Virtual Community: Homesteading on the Electronic Frontier.* Addison-Wesley, 1993.

[44] Robin Good. Online Virtual Worlds: A Mini-Guide. `http://www.masternewmedia.org/virtual_reality/virtual-worlds/virtual-immersive-3D-worlds-guide-20071004.htm`, April 2007.

[45] Betsy Book. Virtual Worlds Review. `http://www.virtualworldsreview.com/`, February 2006.

[46] Sportvision. Virtual Yellow 1st and Ten. `http://www.sportvision.com/`, 1998.

[47] Frank S. Dean Jr., Pat Garrity, and Christopher B. Stapleton. Mixed Reality: A Tool for Integrating Live, Virtual & Constructive Domains to Support Training Transformation. In *Proceedings of the Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC)*, 2004.

[48] Electronic Arts. SimCity. `http://simcity.ea.com/`, 2007.

[49] Linden Lab. Second Life. `http://www.secondlife.com`, 2003.

[50] Cory Ondrejka. Personal communication during December 14 meeting at the MIT Media Lab, 2006.

[51] Jim Carlton. *Apple: The Inside Story of Intrigue, Egomania, and Business Blunders.* Crown Business, 1997.

[52] Linden Lab. Economic Statistics. `http://secondlife.com/whatis/economy_stats.php`, 2007.

[53] Mark Wallace. 3pointD.com: The Metaverse and 3D Web, as blogged by Mark Wallace and friends. `http://www.3pointd.com/`, 2007.

[54] Michael Rymaszewski, Wagner James Au, Mark Wallace, Catherine Winters, Cory Ondrejka, and Benjamin Batstone-Cunningham. *Second Life: The Official Guide.* Wiley, 2007.

[55] SLBuzz.com. Second Life Statistics at SLBuzz.com, 2007. Online; accessed 19-August-2007].

[56] Wikipedia. Manhattan — Wikipedia, The Free Encyclopedia, 2007. Online; accessed 19-August-2007].

[57] Saul Steinberg. View of the world from 9th avenue. Cover illustration of The New Yorker, 29 March 1976.

[58] Adrian David Cheok, Siew Wan Fong, Kok Hwee Goh, Xubo Yang, Wei Liu, Farzam Farzbiz, and Yu Li. Human Pacman: A Mobile Entertainment System with Ubiquitous Computing and Tangible Interaction over a Wide Outdoor Area. In *Fifth International Symposium on Human Computer Interaction with Mobile Devices and Services (Mobile HCI)*, pages 209–223, September 2003.

[59] PacManhattan. `http://pacmanhattan.com`, 2004.

[60] Bruce Thomas, Ben Close, John Donoghue, John Squires, Phillip De Bondi, Michael Morris, and Wayne Piekarski. ARQuake: An Outdoor/Indoor Augmented Reality First Person Application. In *Fourth International Symposium on Wearable Computers (ISWC'00)*, 2000.

[61] Gita Sukthankar. The DynaDOOM Visualization Agent: A Handheld Interface for Live Action Gaming. In *Workshop on Ubiquitous Agents on Embedded, Wearable, and Mobile Devices (Conference on Intelligent Agents & Multiagent Systems)*, July 2002.

[62] Carsten Magerkurth, Adrian David Cheok, Regan L. Mandryk, and Trond Nilsen. Pervasive Games: Bringing Computer Entertainment Back to the Real World. *ACM Computers in Entertainment*, 3(3), July 2005.

[63] Wade Roush. New Portal to Second Life: Your Phone. Technology Review, `http://www.technologyreview.com/Infotech/18195/`, 2007.

[64] Jeff MacIntyre. Sim Civics. Boston Globe: `http://www.boston.com/news/globe/ideas/articles/2005/08/07/sim_civics/`, August 2005.

[65] William Miller. Dismounted Infantry Takes the Virtual High Ground. *Military Training Technology*, 7(8), December 2002.

[66] Dean Jansen. Beyond Broadcast 2007 – The Conference Goes Virtual: Second Life. `http://www.beyondbroadcast.net/blog/?p=37`, 2006.

[67] IBM. Hursley Island in Second Life. `http://slurl.com/secondlife/Hursley/0/0/0/`, 2007.

[68] ciemaar. Real Life Control Panel for Second Life. `http://channel3b.wordpress.com/2007/01/24/real-life-control-panel-for-second-life/`, 2007.

[69] Phillip Torrone. my wearable computer – snowcrash. `http://www.flickr.com/photos/pmtorrone/sets/1710794/`, January 2006.

[70] Wade Roush. Second Earth. *Technology Review*, 110(4):38–48, July/August 2007.

[71] Gary Boone. Reality Mining: Browsing Reality with Sensor Neworks. *Sensors Magazine*, 21(9), September 2004.

[72] Turbulence. `http://www.turbulence.org/`, 2007.

[73] Ars Virtua. `http://arsvirtua.org/`, 2007.

[74] Turbulence. Mixed Realities Commissions. `http://transition.turbulence.org/comp_07/awards.html`, 2007.

[75] Gordon Bell. A Personal Digital Store. *Communications of the ACM*, 44(1):86–91, January 2001.

[76] Jim Gemmell, Gordon Bell, and Roger Lueder. MyLifeBits: A Personal Database for Everything. *Communications of the ACM*, 49(1):88–95, January 2006.

[77] MySpace. `http://www.myspace.com/`, 2007.

[78] Facebook. `http://www.facebook.com/`, 2007.

[79] Twitter. `http://twitter.com/`, 2007.

[80] Harold A. Innis. *The Bias of Communication*. University of Toronto Press, 1951.

[81] Marshall McLuhan. *Understanding Media: The Extensions of Man*. McGraw-Hill, 1964.

[82] David Cronenberg. Videodrome. Canadian Film Development Corporation (CFDC), February 1983. Tagline: First it controlled her mind, then it destroyed her body... Long live the new flesh!

[83] Kimiko Ryokai, Stefan Marti, and Hiroshi Ishii. I/O Brush: Drawing with Everyday Objects as Ink. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*. ACM, April 2004.

[84] Joshua Lifton, Mark Feldmeier, Yasuhiro Ono, Cameron Lewis, and Joseph A. Paradiso. A Platform for Ubiquitous Sensor Deployment in Occupational and Domestic Environments. In *Proceedings of the Sixth International Symposium on Information Processing in Sensor Networks (IPSN)*, pages 119–127, April 2007.

[85] Cameron Lewis. MIT 6.UAP Undergraduate Advanced Project Final Report: Prototyping and Expanding a Tool Set for a Ubiquitous Sensor Infrastructure, December 2006.

[86] Joshua Lifton, Deva Seetharam, Michael Broxton, and Joseph Paradiso. Pushpin Computing System Overview: a Platform for Distributed, Embedded, Ubiquitous Sensor Networks. In *Proceedings of the International Conference on Pervasive Computing*, pages 139–151, August 2002.

[87] Shmuel Ben-Yaakov, Moshe Shvartsas, and Stanislav Glozman. Statics and Dynamics of Fluorescent Lamps Operating at High Frequency: Modeling and Simulation. *IEEE Transactions on Industry Applications*, 38(6):1486–1492, November–December 2002.

[88] Joshua Lifton, Manas Mittal, Michael Lapinski, and Joseph A. Paradiso. Tricorder: A mobile sensor network browser. In *Proceedings of the ACM CHI 2007 Conference - Mobile Spatial Interaction Workshop*, April 2007.

[89] Nokia. Nokia 770 Internet Tablet. `http://www.nokiausa.com/770`, 2006.

[90] PNI Corporation. TCM3 3-axis compass. `http://www.pnicorp.com`.

[91] PyGame. Takes the "C++" out of "Game Development". `http://www.pygame.org`, 2007.

[92] Maemo. Open source development for Internet tablets. `http://www.maemo.org`, 2007.

[93] University of Edinburgh Centre for Speech Technology Research (CSTR). The Festival Speech Synthesis System. `http://www.cstr.ed.ac.uk/projects/festival/`, 2007.

[94] Hugo Liu and Pattie Maes. What Would They Think? A Computational Model of Attitudes. In *Proceedings of the International Conference on Intelligent User Interfaces (IUI 2004)*, January 2004.

[95] Vassar College. Sistine Chapel Re-Creation. `http://slurl.com/secondlife/Vassar/175/90/25/`, 2007.

[96] Sibley Verbeck. Founder and CEO of The Electric Sheep Company. Personal communication via e-mail, July 9, 2007.

[97] Alana Semuels. Virtual marketers have second thoughts about Second Life. `http://www.latimes.com/business/la-fi-secondlife14jul14,1,3135510.story`, July 14, 2007.

[98] IBM. Virtual Worlds: Where Business, Society, Technology, & Policy Converge. `http://www.research.ibm.com/research/press/virtualworlds_agenda.shtml`, June 15, 2007.

[99] Joshua Lifton. Media Lab Dual Reality Open House. `http://slurl.com/secondlife/Ruthenium/0/0/0/`, 2007.

[100] Drew Harry, Dietmar Offenhuber, and Orkan Telhan. Sousreality. `http://web.media.mit.edu/~dharry/sousreality/`, 2007.

[101] Barbara Barry, Ally Lee, Daniel Jang, and Jason Uh. Story Cafe. `http://storycafe.media.mit.edu/Secondlife/`, 2007.

[102] Jeff Orkin. The Restaurant Game. `http://web.media.mit.edu/~jorkin/restaurant/`, 2007.

[103] Drew Harry. Information Spaces. `http://web.media.mit.edu/~dharry/infospaces/`, 2007.

[104] Scratch. Imagine, Program, Share. `http://scratch.mit.edu/`, 2007.

[105] Drew Harry and Dietmar Offenhuber. The Projects. `http://web.media.mit.edu/~dharry/theprojects/`, 2007.

[106] Joseph Paradiso, Joshua Lifton, and Mark Feldmeier. MAS 961 Special Topics in Media Arts and Sciences: Developing Applications for Sensor Networks. `http://www.media.mit.edu/resenv/classes/MAS961/`, 2005.

[107] Daniel Olguín Olguín. Sociometric Badges: Wearable Technology for Measuring Human Behavior. Master's thesis, MIT Media Lab, May 2007.

[108] Emmanuel Munguia Tapia, Stephen S. Intille, and Kent Larson. Activity Recognition in the Home Setting Using Simple and Ubiquitous Sensors. In *Proceedings of the International Conference on Pervasive Computing*, pages 158–175, 2004.

[109] Hyemin Chung. MAS 961 Class Project: Light Analysis System. `http://web.media.mit.edu/~ence/mas961/exboard.html`, 2005.

[110] Luis Blackaller, Sanghoon Lee, and Manas Mittal. Plug Sense: MAS 622J Final Project. `http://web.media.mit.edu/~black/plugsense/cover.html`, 2006.

[111] Edward R. Tufte. *The Visual Display of Quantitative Information.* Graphics Press, 1983.

[112] Chaomei Chen. *Information Visualisation and Virtual Environments.* Springer-Verlag, 1999.

[113] Google. Google Maps. `http://maps.google.com`, 2007.

[114] Navteq. Traffic.com. `http://www.traffic.com`, 2007.

[115] Wei Ye and John Heidemann. Medium Access Control in Wireless Sensor Networks. Technical Report ISI-TR-580, USC Information Sciences Institute, October 2003.

[116] Matthias Götz, Manuel Rapp, and Klaus Dostert. Power Line Channel Characteristics and Their Effect on Communication System Design. *IEEE Communications Magazine*, 42(4):78–86, April 2004.

[117] Maxim Lobashov, Gerhard Pratl, and Thilo Sauter. Implications of Power-line Communication on Distributed Data Acquisition and Control Systems. In *Proceedings of ETFA '03 (Emerging Technologies and Factory Automation)*, volume 2, pages 607–613, September 2003.

[118] Niovi Pavlidou, A.J. Han Vinck, Javad Yazdani, and Bahram Honary. Power Line Communications: State of the Art and Future Trends. *IEEE Communications Magazine*, 41(4):34–40, April 2003.

[119] Carson J. Reynolds and Christopher R. Wren. Worse is Better for Ambient Sensing. Technical Report TR2006-005, Mitsubishi Electric Research Laboratories, March 2006.

[120] Christopher R. Wren and Srinivasa G. Rao. Self-configuring, Lightweight Sensor Networks for Ubiquitous Computing. Technical Report TR2003-24, Mitsubishi Electric Research Laboratories, October 2003.

[121] Nokia. Nokia N800 Internet Tablet. `http://www.nokiausa.com/N800`, 2007.

[122] GStreamer. Open source multimedia framework. `http://gstreamer.freedesktop.org`, 2007.

[123] Swivel. Tasty Data Goodies. `http://www.swivel.com`, 2007.

[124] CENS. SensorBase. `http://www.sensorbase.org`, 2007.

[125] Kevin Chang, Nathan Yau, Mark Hansen, and Deborah Estrin. SensorBase.org – A Centralized Repository to Slog Sensor Network Data. In *Proceedings of the International Conference on Distributed Networks(DCOSS)/EAWMS*, 2006.

[126] Hayes Raffle, Amanda Parkes, Joshua Lifton, and Hiroshi Ishii. Beyond Record and Play, Backpacks: Tangible Modulators for Kinetic Behavior. In *Proceedings of the ACM CHI 2006 conference*, April 2006.

[127] Gaston Bachelard. *The Poetics of Space*. The Orion Press, 1964. English translation of the 1958 original, La poétique de l'espace.

[128] Rowley Associates. CrossWorks for ARM. `http://www.rowley.co.uk/arm/index.htm`, 2007.

[129] Dean Hall. PyMite: A Flyweight Python Interpreter for 8-bit Architectures. `http://www.python.org/pycon/papers/pymite/`, March 2003.