The Challenges and Opportunities of Integrating the Physical World and Networked Systems

John Barton, Tim Kindberg
Internet and Mobile Systems Laboratory
HP Laboratories Palo Alto
HPL-2001-18
January 24th , 2001*

E-mail: john_barton@hpl.hp.com, timothy@hpl.hp.com

nomadic computing, pervasive computing, context-aware computing, World Wide Web

This paper presents a vision of how to link entities in the real world with the networking infrastructure to create an integrated mobile computing environment. We describe the challenges and opportunities implied by that linkage, and outline the web-based approach that we have adopted towards them.

* Internal Accession Date Only
Approved for External Publication

The challenges and opportunities of integrating the physical world and networked systems

John Barton & Tim Kindberg

Internet and Mobile Systems Laboratory
Hewlett-Packard Laboratories
1501 Page Mill Road
Palo Alto, CA 94304, USA
john_barton@hpl.hp.com, timothy@hpl.hp.com

Abstract

This paper presents a vision of how to link entities in the real world with the networking infrastructure to create an integrated mobile computing environment. We describe the challenges and opportunities implied by that linkage, and outline the web-based approach that we have adopted towards them.

1 Introduction

‘Ubiquitous computing’, ‘nomadic computing’, ‘pervasive computing’, ‘context-aware computing’ and ‘augmentation of the real world’ have been topics of research since the early 1990s [37, 21, 30, 39]. Whatever their differences, those topics are related by a desire to merge the physical world with the virtual world of electronic services and applications. While the visions of these efforts have not yet been realised, we can see from these Mobicom proceedings that wireless networking and technologies for portable and embeddable devices are maturing. Applied at the intersection of the physical and networking worlds, these maturing technologies can bring the next generation of mobile computing visions out of the lab and into users’ everyday lives.

We believe that a focus on integration of the physical world and the virtual world in the network can provide a leap forward in the level of our debate about mobile computing and in the level of functionality that we offer. We haven’t yet been able to deliver a true mobile computing solution with broad appeal and substantial value. Currently, the user’s experience of mobile computing consists largely of being able to read their email or browse the web—to carry out otherwise desktop-bound activities—from their laptop, personal digital assistant (PDA) or mobile phone. And even some of those mundane activities are frequently hampered by the need first to make complicated configuration settings that depend upon where the user happens to be.

Simply extending the traditional, desktop-centric infrastructure will not meet the requirements of nomadic people in their daily situations. People are naturally nomadic. People move from place to place and use the objects around them as they follow their interests and day-to-day activities. Current mobile computing systems actually work to shield applications from nomadicity and environmental changes. To exploit the full potential of mobile computing we need to break out of this shield so that we can leverage the power of network infrastructure in real world applications. Adding network-enhanced electronic functionality on top of the real, everyday world dramatically expands the scope for mobile computing applications that matter to users.

Networked sensors in the hands of users

The challenge that this paper describes is how to deliver new types of mobile computing services and applications that will enhance a much broader spectrum of users’ activities than traditionally desktop-based ones such as reading email and web browsing. The key technology that we advocate to accelerate the adoption of mobile computing technology is a sensor-enhanced, mobile-aware, portable Web client realised in a variety of portable devices.

By ‘sensor’ we mean any device that captures data or content from the real world. A simple example is a bar-code reader on a PDA. Other types of sensor that can be conveniently carried or integrated with portable and wearable devices include infrared receivers, RFID tag readers, GPS and other positional sensors, cameras and portable document scanners. The devices could be PDAs, phones, cameras, camcorders, key fobs and so on. They are digital tools enhanced with detectors for the physical world and connected by wireless networks to the virtual world. The virtual world is enhanced with applications and services sensitive to the data available in the physical world. Those applications and services interact with the user through a device that they carry or through networked devices in the user’s environment. The user’s sensor-enhanced device becomes their control for network services pertaining to their physical world.

This scenario is now possible because of a combination of achievements:

- portable and embeddable sensors for capturing values from physical objects and actuators for manipulating physical objects;
- Cheap tagging technologies including barcodes and RFID;
- portable and wearable computing devices for use in processing those sensed values and for offering an interface to the user;
- wireless connectivity to infrastructure services and to nearby devices in the physical environment.

Sensors have been integrated in specialised process control and medical applications for some time. Now they are becoming available in portable devices at the same time that wireless networking and Internet services are maturing.

This combination puts the power of network infrastructure literally in the hands of users for mundane activities. From the users' perspective, they are engaging in straightforward actions such as pointing and taking readings as part of their everyday activities. The results are not straightforward: suddenly parts of the physical world are integrated with their applications. The distinction between the physical and virtual world becomes blurred–as it should be. What mobile computing facilities are needed to seize this opportunity? That is our challenge.

The light switch

To convey the way that a mobile networking application could differ from its desktop-bound cousins, we introduce an example that is extremely simple: turning on a light bulb. We want to turn on this light bulb using pre-network technology, network technology, and finally mobile computing technology.

To make this concrete we've built a simple demonstration. We work in an open-plan environment of partitioned workspaces called 'cubes'. Cubes are lit from above, but the cubes do not have individual light switches.

To switch on the overhead lights using 'pre-network' technology, we would have to walk down the hallway to the cluster of switches, identify the correct one for our cube, and actuate the switch.

Fortunately, our building is equipped with the technological advance of web-enabled light switching. We use our PC to navigate to a web page for our floor's lighting system and click on our cube's location on the floor's plan. Then the lights come on.

This network-era alternative to a physical switch and a walk down the hall illustrates two characteristics typical of problems at the boundary of computing and the physical world: (1) the need to actuate something in the user's physical environment–the light bulb–that does not have a convenient physical interface–a light switch–and (2) the poor match of a PC as an alternative interface.

The PC is in the wrong place: we want to switch on the lights as we enter the cube. Using the PC takes effort: even if it is already booted, one has to sit down, find the browser window and select the link to the cube's light control.

Of course one could argue that the 'ideal' solution already exists: a light switch at the cube entrance! However, reconfiguring cubes equipped with light circuits is prohibitively expensive.

There is a network approach that is both more flexible than light circuits and more convenient for end users: place a 'soft' switch at the cube entrance. This looks like a light switch but it is a network component that emulates the PC's browser signal. Home-automation systems like X10 [20] use this approach. A soft switch has a physical interface to actuate, can be positioned wherever we like, and directly names the light bulb for our cube. The cube can be reconfigured because the switch need only be on the network, not on the specific light circuit.

The mobile computing solution has similar properties without the need for the switch or the wires connecting that switch to the network. We use a physical but symbolic representation of a light switch: a barcode next to a 'light' icon on a piece of paper pinned to the cube wall where one would expect a physical switch to be. On entering, we use a PDA with an integrated laser barcode scanner and wireless network connection to scan the barcode. The sensed identifier is converted by a network lookup into a URL; an HTTP GET operation on that URL illuminates the cube.

The mobile computing solution uses an iconic physical interface sensed by a handheld device and mapped by network software to a name for a physical action. This physical interface is both conveniently located for the user and extremely flexible for the cube provider. Of course we could not today justify a PDA and a wireless network just for cube users to adjust their lights. But we envisage a world where such applications of wireless networking are so numerous that PDAs and wireless networks are as commonplace as light switches and electric circuits are today.

Outline of the paper

Section 2 gives an overview of the growing body of research into linking the physical and virtual worlds. Section 3 outlines CoolTown's web-based approach, which, we believe, will help to make the physical/virtual linkage useful to mobile computing users. Section 4 presents some of the challenges that must be tackled to bring about this potential leap forward for mobile computing. Section 5 concludes with a discussion of related perspectives.

2 Physical-Virtual Links

In this section we describe how researchers have utilised the components mentioned above—sensors, actuators, portable and wearable devices—to produce various types of linkage between the physical and the virtual. We aim to show the opportunities for constructing new mobile computing systems, and to provide a basis for analysing the remaining research challenges.

2.1 Linkage mechanisms

All of the systems we shall discuss involve physical entities, virtual entities, and network-based linkage mechanisms between them. Before describing the types of linkage that appear at the application level in those systems, we first define our terms and describe the linkage mechanisms that are available.

Physical entity. A physical entity is any person, place or thing. We include entities that have electronic functionality of their own (e.g. a printer) and those that don't (e.g. a book, a sheet of ordinary paper or–in most cases!–a person).

Virtual entity. A virtual entity is a resource or service whose functionality is linked to a physical entity–for example, a web page describing the entity or a Java object that controls it.

Linkage mechanism. A physical entity may be linked to several virtual entities, and *vice versa.* Various types of linkage mechanism can be used to effect relationships between physical and virtual entities. The linkage may be causal, so that the state of the physical entity affects the state of the virtual entity or *vice versa*. But the simplest case, which we shall describe more fully, is where we use a sensor to identify a physical entity, and thereby invoke an operation on the corresponding virtual entity. For example, a user may do this to obtain information about an object in front of them, or a smart shipping container may sense and thus record its contents.

Identification

We can identify a physical entity by reading a tag attached to it, by recognising it using image-processing techniques, or by looking it up from its position.

Tags. If we can get up close enough to the entity, then reading an identification tag attached to it is often an efficient method. There are several tag technologies [36]:

Infrared tags. The tag may be a low-cost, low-power infrared 'active badge' worn by a person [35], or a 'beacon' attached to a thing or place [18]. These devices emit the identifier over IrDA, for reception by infrared-equipped devices including PDAs.

Optically sensed tags. These include standard barcodes, already found on many everyday items, and symbols specially designed for easy capture with a digital camera, including 'cybercodes' [28] and 'glyphs' [15]. Barcodes can be printed and inexpensive readers are becoming available.

RFID tags. Radio frequency (RF) identification tags can be read at a distance and, since they operate by induction, require no power source of their own. A recent advance is the ability to 'print' an RFID antenna onto surfaces using conductive ink [25].

Contact tags. iButtons [7] are read by electrical contact with their casing. Like RFID tags, they do not require their own power.

These tagging technologies have relative advantages and disadvantages in terms of cost and suitability for different physical environments. For example, sometimes reading at a distance is desirable (IR, RFID, glyph recognition); in other situations it is preferable that the user should bring the reader up close to make a definite and unambiguous identification, such as by scanning a barcode.

Computer vision. Stereo computer vision–object recognition–may be used instead of tagging [4]. This has the advantage of eliminating the logistics of tagging but it has the disadvantage of requiring relatively powerful computing resources. Considerable work is needed before this method could be used routinely.

Positioning. In cases where objects move rarely or not at all, or are automatically tracked, a third means of identification is to use a positioning sensor to determine the entity's coordinates, and so look them up in a database. By adding an electronic compass, we can also identify remote objects by pointing at them. The GPS is a widely available positioning system, but only outdoors. Short-range RF triangulation may be deployed indoors or out. We can incorporate ultrasound techniques for more fine-grained positioning, down to a few centimetres [14, 27].

2.2 Applications of virtual-physical links

Once we can establish virtual/physical links, the physical and virtual entities may play a variety of roles in augmenting their counterparts across the physical-virtual divide. We group the research efforts to explore this augmentation, to highlight the potential for leveraging this work in applications of mobile computing.

Physical browsing. Many systems take everyday physical entities as they stand, regardless of whether they have electronic functionality of their own, and augment them by systematically correlating digital documents with them. We'll call this *physical* *browsing*: users designate entities that interest them, and thereby obtain documents ('pages') about them.

The DigitalDesk project at Xerox EuroPARC [38] used ceiling-mounted cameras and image recognition to identify objects on an ordinary desk. When the user places a piece of paper on the desk, DigitalDesk makes the corresponding electronic document available for manipulation.

Advances in mobile networking and tagging enabled extension of this idea beyond the physical desktop. Want et al. at Xerox PARC augmented books and documents by attaching RFID tags and presenting the electronic version to users who scanned them with handheld devices [34]. The Guide system [6] presents tourists with web pages about the sites they visit. In CoolTown, users can find out about the place they are in by reading an infrared beacon with their handheld device [19]. A CoolTown shopper who scans the barcode on a can of beans can select pages that the store provides about the product, or pages maintained by interest groups—for example, organic- or diabetes-related pages.

Physical objects as content 'repositories'. We can associate objects with content so that users may transfer the content to one another or move it from place to place by passing the corresponding object around or carrying it. The object does not have to store the content physically: the content can remain in the network. The object serves as a mnemonic to the user and, by being identifiable, as the name or location of the content to the system. For example, the object may be a watch or small building block in the case of the I-Land project [32.]. Or it may be anything with an attached 'post-it' note bearing a barcode, such as CyberCode stickers [28] and WebStickers [16]. Users run a 'physical binding' application to bind their chosen content with the object; they run a physical browsing application to retrieve the content from the object.

Copy-and-paste in the real world. Some objects (such as scanners and voice recorders) are sources of content. Others (printers and audio players) are content sinks. By adding 'clipboard' objects that act as temporary repositories of content or pointers to content, we can implement copy-and-paste across locations. Thus a user can be presented with an image at an amusement park, bind the image to their watch and, when they get home, wave their watch in front of the printer in order to print it there, or in front of their PC to insert it into a document. DigitalDesk [38] and 'InfoStick' [22] implement variants on this idea.

Objects as communication points. By associating barcodes on objects such as walls, books or cans of beans with an electronic bulletin board, we enable users who encounter the same object (or instances of the same class of object) to communicate with one another. In CoolTown we call this 'virtual graffiti'. A more serious application is equipment maintenance. The person who finds a fault in, for example, an outlying equipment item in an industrial plant, leaves a voice message there so that the person who later comes along and repairs it can access the message conveniently [26].

Objects as 'physical icons'. Physical objects can be bound to actions, and thus play a role very similar to that of icons and other widgets on our PC desktops, but across the network. In the light-switch example of Section 1, the barcodes on the wall play such a role: they are bound to the 'lights on' / 'lights off' actions. Masui and Siio [24] describe how the interfaces to the Hi-Fi, washing machine and other devices in the home can be laid out as barcodes on the pages of a notebook that the user carries with them, so as to be able to control those devices remotely. Ullmer and Ishii [33] describe a computer-enhanced 3D-modelling system in which architects turn the hands of a physical clock to observe the shadows cast by model buildings at different times of the day.

Physical objects may also participate as 'input' to application functions. For example, in CoolTown, virtual entities corresponding to physical entities within a physical place are bound into the virtual place's registry [5]. CoolTown users can bind virtual entities into the registry by sensing the corresponding physical object to select its virtual counterpart [18]. Another example is where the creator of a web page wishes to select the source of an image to be placed in the page. To do so, they may sense the source device, such as a camera or a scanner.

Objects as physical representations of state. A particularly salient way for systems to communicate the state of electronic processes is to change the state of a physical object. For example, a certain Silicon Valley research centre maintains a fountain whose height corresponds to the value of the company's share price. Note that in this example the virtual entity is actively controlling the physical entity, whereas in the previous examples the virtual entities were all passive services invoked by the user.

Mixed reality. Some systems use image and sound projection to enhance physical objects. For example, Billinghurst and Kato [2] describe how users wearing head-mounted displays are shown superimposed images as they view physical objects. The DigitalDesk and the Augmented Surface [29] both employ ceiling-mounted projectors so that users see a combination of physical documents and projected electronic documents on their work surface.

Smart environments. A smart environment is based on a physical environment such as a house (e.g. the Aware Home [17]) or room (e.g. the iRoom [10]) equipped with sensors and actuators. What makes the environment smart is that services and applications in the infrastructure process the sensor readings, in order to trigger events, to adapt their behaviour or to control the physical state of the room. To follow our light switch example, a smart cube could be configured to sense the presence of a human and switch on the lights automatically. The Aware Home has ‘smart floors’ to recognise the occupants from their footsteps. This and other sensing techniques are used in applications such as helping users find lost objects.

In the Aware Home, services and applications persist in the infrastructure while the physical contents change. An interesting variation is where the virtual contents also change as humans and other physical contents come and go. That occurs in CoolTown, as entities are bound into and unbound from a place's registry. The virtual entities in the registry of a CoolTown place may be active, and discover one another through the registry. For example, the virtual object corresponding to a user may discover the telephone and any other means of communication that exist within the place, and make that information available so that the user can be contacted. Note that this is different from ‘service discovery’ systems such as SLP [13] and SSDP [11]: entities such as telephones that are not themselves on the information network participate through their virtual counterparts.

3 CoolTown: the real-world wide web

The CoolTown project [5, 18, 19] creates linkages between the physical and virtual worlds in the form of the ‘real-world wide web’: an integration of the Web with physical entities. We realise the linkages as web *links*. The real-world Web differs from the conventional web only in that users can find links by sensing the physical world, as well as by browsing web pages. It utilises the same HTTP and URI standards.

CoolTown incorporates a method called ‘eSquirt’ for collecting links—URLs—from infrared beacons attached to walls, printers, radios, pictures and other physical objects. Users obtain services by browsing the links or by ‘squirting’ them into infrared receivers on appliances such as printers, projectors, Internet radios [23] or PCs.

Infrared enabled us to leverage many off-the-shelf devices. But the choice of using URLs turned out to be more significant. It meant that any type of web-accessible content or service could be associated with physical entities and accessed from any type of HTTP-capable device [3]. Increasingly many handheld devices are equipped with browsers. Correspondingly, increasingly many appliances in the environment, such as printers, are equipped with an HTTP server that enables them to be controlled from browsers. And many appliances can participate in eSquirt, either because they have an HTTP client of their own that can fetch and process the content bound to URLs provided to them, or because the URLs can be passed to a machine deeper in the infrastructure that drives them.

We have broadened our use of sensing mechanisms to encompass barcodes, RFID tags and iButtons. That massively increases the set of physical entities to which we can bind virtual services—to include, for example, printed documents and cans of food. We have also enabled the mapping from the physical to the virtual to be one-to-many, so as to accommodate the differing preferences and activities of nomadic users. Software components called resolvers turn the identifiers into URLs. Users select resolution services by navigating on the Web.

The real-world wide web that we have begun to construct takes mobile computing out of the laboratory. It makes it possible for users to engage simultaneously in mobile computing and their familiar physical world. It leverages the wide deployment of web software and services, and it offers users a natural step forward: from browsing cyberspace to browsing the physical world.

4 The mobile computing challenges

Whatever approach is taken to realise the linkages, many barriers must fall to integrate the network with the physical world and realise the potential of mobile computing.

Usability challenges

Putting sensor-enhanced, portable, wirelessly connected clients into the hands of users raise several usability issues:

Switching between the physical and the virtual. The cognitive and physical effort of sensing itself (for example, reading barcodes), and of viewing and interacting with virtual resources and services, threatens to distract us and interfere with our physical activities. Can we design interfaces that enhance the physical world with the virtual and allow users to switch conveniently between them?

Appropriate Devices Currently PDAs have demanding interfaces due to their small screen and cumbersome two-handed input modes. Ease of use and appropriateness to the user’s task should lead to a diversity of devices that mirrors the diversity of people and their work. Reduced cost through large volume leads towards a single device, a ‘perfect’ PC/phone combination [12]. But the ‘one size fits all’ device tends to have low efficacy for any particular task. Can we find ways to lower the cost of diversity in devices?

Common use models. The nomadic user expects to find new virtual resources and services as they change their physical location. However, to place the user into completely unfamiliar user-interface territory every time they enter a new physical area would be unacceptable. Moreover, where networking is integrated with the physical world, the 'user interface' becomes complicated by the effect of physical actions and circumstances on virtual behaviour. Can we develop a common use model for access to virtual services from the immediate physical world?

Digital furniture. The digital future has portable devices but also digital 'furniture': networked electronic tools specialised for simple roles in tasks that involve multiple computers and services. Just as public places such as cafés and amusement parks currently provide tables and chairs, we expect them increasingly to provide printers and kiosks where users can browse information. Can we provide an approach to mobile computing that integrates handheld devices with digital furniture—as we are attempting with the CoolTown eSquirt facility?

Personalization. Within common use models, the nomadic user requires a degree of personalization in the data and applications that they use and in the response that they receive from their current environment. For example, a user can search for information on a can of beans by scanning its bar code. Diabetics, vegetarians, shoppers looking for bargains, and store employees would seek different information. Can we enable customization while managing ambiguity?

Configuration. An intrinsic problem of a 'personal' device lies in the complexity and investment inherent in personalization. To maintain the sense of a physical-virtual link for nomadic users, we require rapid client reconfiguration as the user enters new physical areas. Hardware and software installation, configuration, and maintenance take time and expertise. How can we obtain the effects of personalization without its costs?

System challenges

Bridging the physical and virtual worlds requires overcoming issues common to other applications of mobile computing. For some issues, the challenges have the same character others have noted [9]. Configuring a portable device's network as its nomadic owner moves or as they start up their device in a new area, as well as keeping devices small and light weight while adding sensors and better networking are critical issues for success in physical/virtual linkage, but this is equally true for many other applications. Other issues have a different twist once the physical/virtual link enters.

Discovery. Nomadic clients operate initially ignorant of the local resources and services in an area to which the user has moved. How can we discover the address of a peer or server we have not previously encountered but we are physically near? How do we define 'near' and 'physical area' for services? Rarely does a subnet correspond to what the user thinks of as a separate place. For example, two adjacent meeting rooms will often be connected to the same subnet. We cannot rely on subnet multicast alone.

Registration. The other side of the discovery problem is provisioning and registration [18]. Maintenance of the physical-virtual links will require network software, computation, and storage resources analogous to but much larger than today's Domain Name System. Can we build truly pervasive linkage services that we can count on routinely throughout the global Internet? Even with such services, configuration of the physical/virtual links has to be routine and even automatic to be useful. How will users deal with stale links? How can we synchronize virtual representations with changes in the real world? How can users, who understand the underlying real-world semantics, best be involved in building and maintaining the links?

Security and privacy. Adding physical/virtual linkage doesn't create new difficulties for wireless network security, but it does increase the pressure to push wireless systems beyond enterprise systems. We need our devices to work in businesses other than our own; in shops and schools we may only visit rarely. Our clients will be moving out of our firewalled fortresses and off our proprietary telecommunications networks. To use devices with confidence and to be allowed to use local resources we encounter will require new protection for both sides.

Mobile devices can be even more personal than personal computers: they can go everywhere we go. Privacy will thus become an even more important issue for users whose physical location and other attributes may become known. This aspect of security is mostly unknown territory. We need to prevent inappropriate use of information about the use of particular physical-virtual links by particular devices if these devices can be traced to a particular person. Thus anonymous or pseudonymous traversal of physical links will be needed. And digital furniture such as kiosks that the nomadic user employs temporarily and then leaves should be 'amnesic', so that users can be assured of leaving no personal information behind [31]. Potential users of future mobile computing infrastructure will hesitate until these issues are more clearly understood.

Naming. Sensor support raises new issues that research must tackle. For example how do we give names—identifiers that can be sensed—to trillions of physical entities? A CoolTown infrared beacon provides a direct URL of a resource to a client, but the beacon is an active device requiring a battery and hardware. We also use identifiers looked up to obtain a URL that is a function of contextual factors such as the user or location. The lookup gives us the flexibility of a level of indirection and allows compact identifiers that can be implemented in inexpensive linear barcodes and RFID tags. However, this lookup requires a highly scalable name service.

Name space management is also an issue. Some name spaces such as Universal Product Codes (UPC) are rigorously managed. But a scalable naming system should be able to assign unique identifiers to trillions of entity instances (UPC codes identify classes). It should be possible for ordinary individuals to mint new identifiers without cost. And we must be able to incorporate identifiers belonging to heterogeneous name spaces such as UPC and iButton identifiers. Is the Web naming system (URIs) prepared for such an onslaught?

5 Conclusion

We have highlighted the opportunities and challenges in a mobile computing system based upon links between the physical world of electronic and non-electronic entities and the virtual world of applications and services. We have surveyed a variety of research results that help build a foundation for establishing the links. We have advocated the real-world wide web, which users access with task-focused, sensor-enhanced, wireless devices.

Our view of the importance of this arena for mobile computing overlaps two recent Mobicom challenges. Esler et al. [8], in describing the Portolano research program, called for more research on multiple user interfaces, horizontal services, and computation in the infrastructure, all critical ingredients for building links between the physical and virtual worlds. Our focus here is on a systematic means for selecting services by integration of physical objects with the user interface. This simplifies discovery and allows us to invoke services without requiring agent assistance. Our adoption of Web technology leverages its advantages as a distributed computing system that puts the user in the semantic driving seat.

Banavar et al. [1] described an application model for pervasive computing. We subscribe to their vision statement in detail. Our approach to realizing this vision differs fundamentally. By augmenting web clients with sensors and input devices, we allow services to emerge that leverage this capability. Compared to monolithic application development, we sacrifice control and certainty: the service developer and the client developer will probably never meet. In return we gain the ability to adapt to rapidly evolving technologies, a critical element for mobile computing, as it has been for e-commerce.

These and the other efforts we have described highlight the emergence of a new arena for mobile computing. Sensor-enhanced wireless devices integrate the physical world with network services and allow users into this new arena.

References

[1] G. Banavar, J. Beck, E. Gluzberg, J. Munson, J. Sussman and D. Zukowski. “An Application Model for Pervasive Computing”. *Proc. MOBICOM 2000: the 6th Annual Int’l Conf. on Mobile Computing and Networking* August 6 - 11, 2000, Boston, MA USA, pp. 266-274

[2] M. Billinghurst, and H. Kato, “Collaborative Mixed Reality”. In *Proc. of Int’l Symp. on Mixed Reality (ISMR '99).* pp. 261-284.

[3] G. Borriello, R. Want: “Embedded Computation Meets the World Wide Web”. *CACM* 43(5) pp.59 -66

[4] B. Brumitt, B. Meyers, J. Krumm, A. Kern and S. Shafer. “EasyLiving: Technologies for Intelligent Environments”. *Proc. of the 2nd Int’l Symp. on Handheld and Ubiquitous Computing*. Bristol, UK 25-27 September 2000, pp. 12-29.

[5] D. Caswell and P. Debaty. “Creating Web Representations for Places”. *Proc. 2nd Int’l Symp. on Handheld and Ubiquitous Computing 2000*, p 114-126.

[6] K. Cheverst, N. Davies, K. Mitchell and A. Friday. “Experiences of Developing and Deploying a Context-Aware Tourist Guide: The GUIDE Project”, *Proc. MOBICOM 2000: the 6th Annual Int’l Conf. on Mobile Computing and Networking* August 6 - 11, 2000, Boston, MA USA, pp. 20-31.

[7] Dallas Semiconductor. http://www.ibutton.com.

[8] M. Esler, J. Hightower, T. Anderson, G. Borriello. “Next Century Challenges: Data-Centric Networking for Invisible Computing”. *Proc. 5th Annual ACM/IEEE Int’l Conf. on Mobile Computing and Networking,* August 15-19, 1999, Seattle, Washington, USA, pp. 256-262

[9] G. Forman and J. Zahorjan, “The Challenges of Mobile Computing”, *Mobility:Processes, Computers, and Agents.* ed D. Milojicic, F. Douglis, and R. Wheeler. ACM 1999, pp. 270-285.

[10] A. Fox, B. Johanson, P. Hanrahan, and T. Winograd, “Integrating Information Appliances into an Interactive Workspace”, *IEEE Computer Graphics & Applications,* 20(3), May/June 2000

[11] Y. Goland, T. Cai, P. Leach, Ye Gu. “Simple Service Discovery Protocol/1.0” *http://www.upnp.org/ draft_cai_ssdp_v1_03.txt.*

[12] J. Guttag, “Communications Chameleons”. *Scientific American*, Aug 1999.

[13] E. Guttman, “Service Location Protocol: Automatic Discovery of IP Network Services”. *IEEE Internet Computing*, Jul.-Aug. 1999. pp. 71-80.

[14] A. Harter, A. Hopper, P. Steggles, A. Ward, P. Webster. “The Anatomy of a Context-Aware Application”. *Proc. 5th Annual ACM/IEEE Int’l Conf. on Mobile Computing and Networking*, Seattle, Washington, USA, August 1999, pp. 59-68.

[15] D.L. Hecht, “Embedded Data Glyph Technology for Hardcopy Digital Documents”, *Proc. Society of Photo-Optical Instrumentation Engineers Symp. on Electronic Imaging, Science and Technology*, San Jose, Calif., 1994, Vol. 2171, pp. 341-352.

[16] L. E. Holmquist, J. Redström, P. Ljungstrand: “Token-Based Access to Digital Information”. *Proc.1st Int’l Symp. on Handheld and Ubiquitous Computing 1999*: pp. 234-245.

[17] C. Kidd., G. Abowd, C. Atkeson, I. Essa, B. MacIntyre, E. Mynatt, T. Starner “The Aware Home: A Living Laboratory for Ubiquitous Computing Research”. *In the Proceedings of the Second International Workshop on Cooperative Buildings - CoBuild'99*.

[18] T. Kindberg and J. Barton. “A Web-based Nomadic Computing System. To appear in *Computer Networks, Special Edition on Pervasive Computing,* Elsevier. HP Labs Technical Report #2000-110.

[19] T. Kindberg, J. Barton, J. Morgan, G. Becker, D. Caswell, P. Debaty, G. Gopal, M. Frid, V. Krishnan, H. Morris, J. Schettino, B. Serra, and M. Spasojevic. “People, Places, Things: Web Presence for the Real World”. *Proc. 3rd Annual Wireless and Mobile Computer Systems and Applications*, Monterey CA, USA, Dec. 2000. p 19.

[20] P. Kingery. “Digital X10”. *http://www.hometoys .com/htinews/feb99/articles/kingery/kingery13.htm*

[21] Kleinrock, L. “Nomadicity: Anytime, Anywhere in a Disconnected World”. *Mobile Networks and Applications* 1(4), 1997, pp. 351-357.

[22] N. Kohtake, J. Rekimoto, Y. Anzai: “InfoStick: An Interaction Device for Inter-Appliance Computing”. *Proc. 1st Int’l Symp. on Handheld and Ubiquitous Computing 200*0, pp.246-258.

[23] V. Krishnan, and G. Chang, Customized Internet radio. *Proc. Ninth International World Wide Web Conference*, Amsterdam, 2000.

[24] T. Masui and I. Siio. Real-World Graphical User Interfaces. *Proc. 2nd Int’l Symp. on Handheld and Ubiquitous Computing 2000,* Bristol, UK 25-27 September 2000, pp. 72-84.

[25] Motorola. *http://www.motorola.com/GSS/SSTG/ smartcard/white_papers/BiStatix_Whitepaper.pdf.*

[26] J. Nilsson J., T. Sokoler, T. Binder, N. Wetcke: “Beyond the Control Room - Mobile Devices for Spatially Distributed Interaction on Industrial Process Plants,” *Proc. 2nd Int’l Symp. on Handheld and Ubiquitous Computing 200*0, pp 30-45.

[27] N. Priyantha, A. Chakraborty and H. Balakrishnan. The Cricket Location-Support System. *Proc. MOBICOM 2000: the 6th Annual Int’l Conf. on Mobile Computing and Networking*. pp 32-43.

[28] J. Rekimoto and Y. Ayatsuka, “CyberCode: Designing Augmented Reality Environments with Visual Tags”, *Proc. of Designing Augmented Reality Environments 2000,* ACM.

[29] J. Rekimoto, M. Saitoh: “Augmented Surfaces: A Spatially Continuous Work Space for Hybrid Computing Environments”. *Proc. 1999 Conference on Human Factors in Computing Systems (CHI '99),* Pittsburgh, PA, May 15-20, 1999. pp. 378-385.

[30] D. Salber, A. K. Dey, and G. D. Abowd, “The Context Toolkit: Aiding the Development of Context-Enabled Applications. In *Proc. 1999 Conference on Human Factors in Computing Systems (CHI '99),* Pittsburgh, PA, May 15-20, 1999, pp. 434-441.

[31] F. Stajano, R. J. Anderson: “The Resurrecting Duckling: Security Issues for Ad-hoc Wireless Networks”. *Lecture Notes in Computer Science, Vol. 1796*, Springer, 2000 pp. 172-194.

[32] N.A. Streitz, J. Geißler, T. Holmer, S. Konomi, C. Müller-Tomfelde, W. Reischl, P. Rexroth, P. Seitz, and R. Steinmetz, “i-LAND: An interactive Landscape for Creativity and Innovation” *ACM Conference on Human Factors in Computing Systems (CHI '99),* Pittsburgh, PA, U.S.A., May 15-20, 1999, pp. 120-127.

[33] B. Ullmer and H. Ishii. “Emerging Frameworks for Tangible User Interfaces”. *IBM Systems Journal* Vol. 39, Nos. 3 & 4, pp. 915-931.

[34] R. Want, K. P. Fishkin, A. Gujar, B.. Harrison: “Bridging Physical and Virtual Worlds with Electronic Tags. In *Proc. 1999 Conference on Human Factors in Computing Systems (CHI '99),* p. 370-377.

[35] R. Want, A. Hopper, V. Falcao, J. Gibbons: “The Active Badge Location System”. *ACM Transaction on Information Systems* 10(1): (1992) pp. 91-102.

[36] R. Want, D. M. Russell. “Ubiquitous Electronic Tagging”. *Distributed Systems Online, IEEE.* http://www.computer.org/dsonline/articles/ds2wan.htm

[37] M. Weiser, “Some Computer Science Issues in Ubiquitous Computing”. *Communications of the ACM*, 36(7), 1993, pp. 74-84.

[38] P. Wellner: “Interacting with Paper on the Digital Desk”. CACM 36(7)1993, pp.86-96.

[39] P. Wellner, W. Mackay, and R. Gold, “Computer Augmented Environments: Back to the Real World”, *Comm. of the ACM*, Vol. 36, No. 7, July 1993, pp. 24 - 26.