particular parts of the room, presence and activity as detected using active badges or pressure sensors, and so on. The other, perhaps most important, piece of context it made use of was *the fact that it was the Reactive Room*. That is, the room was designed for meetings and presentations, and so much activity in the room could be interpreted as being appropriate to meetings and presentations. The same sorts of inferences would probably be inappropriate in other settings, such as a private office, or a home. The "meeting" context, then, also serves to disambiguate the user's goals.

The Reactive Room demonstrated the way that ubiquitous computing did not simply move out of the box on the desk and into the environment but, at the same time, also got involved in the relationship between the environment and the activities that took place there. The topic of "setting-ed" behavior will come back into focus in the next chapter; for the moment, however, we will continue to explore the development of tangible computing.

Design Trends

The systems that have been described—the vision of Ubiquitous Computing, and the Digital Desk and Reactive Room prototypes—have been firmly located in the domain of Computer Science research. However, "academic science" has by no means been the only contributor to the development of Tangible Computing. In fact, one striking aspect of the development of this line of investigation has been the contributions from the perspectives of art and design. Two pieces that have proved to be particularly inspirational to a number of researchers in this area were Durrell Bishop's Marble Answering Machine, and Natalie Jeremijenko's Live Wire.

The Marble Answering Machine was a design exercise undertaken by Bishop in the Computer-Related Design department at the Royal College of Art in London (Crampton-Smith 1995). It explored possible approaches to physical interaction for a telephone answering machine. Rather than the traditional array of lights and buttons, Bishop's answering machine has a stock of marbles. Whenever a caller leaves a message on the answering machine, it associates that message with a marble from the stock, and the marble rolls down a track to the bottom, where it sits along with the marbles representing previous messages. When the owner of the machine comes home, a glance at the track shows, easily and distinctly, how many messages are waiting—the number of marbles arrayed at the bottom of the track. To play a message, the owner picks up one of the marbles and drops it in a depression at the top of the answering machine; because each marble is associated with a particular message, it knows which message to play. Once the message has been played, the owner can decide what to do; either return the marble to the common stock for reuse (so deleting the message), or returning it to the track (saving it to play again later).

The Marble Answering Machine uses physical reality to model the virtual or electronic world. In Bishop's design, marbles act as physical proxies for digital audio messages. By introducing this equivalence, it also enriches the opportunities for interacting with the device. The problem of interacting with the virtual has been translated into interacting with the physical, and so we can rely on the natural structure of the everyday world and our casual familiarity with it. So, counting the number of messages is easy, because we can rapidly assess the visual scene; and operations such as playing messages out of order, deleting messages selectively, or storing them in a different sequence, all of which would require any number of buttons, dials, and controls on a normal digital answering machine, all become simple and straightforward because we can rely on the affordances of the everyday world.

Natalie Jeremijenko's piece "Live Wire," also sometimes known as "the Dangling String" and described by Weiser and Brown (1996), was developed and installed at Xerox PARC in 1994 and explored similar questions of the boundary between the virtual and physical worlds. Physically, Live Wire was a length of plastic "string" around eight feet long, hanging from the ceiling at the end of a corridor. Above the ceiling tiles, the wire was connected to a small stepper motor, which in turn was connected to a device on the local ethernet. Every time a data "packet" passed by on the ethernet, the stepper motor would move, and its movements would be passed on to the string. Ethernet, in its classic form, is a "shared medium" technology—all the traffic, no matter which machine sends it or which machine is to receive it, travels along the same cable.

The busier the network, the more data packets would pass by, and the more the stepper motor would move. The ethernet can carry thousands of packets per second, and so when the network was busy the motor would whir and the string would spin around at high speed, its loose end whipping against the wall nearby.

Others have followed in the footsteps of Bishop and Jeremijenko and continued to explore the design "space" around these issues of the borders between physical and virtual worlds. Feather, Scent, and Shaker (Strong and Gaver 1996) are devices for "simple intimacy." "Feather" features a feather that is gently lifted on a column of air, to indicate to its owner that, perhaps, a photograph of them has been picked up somewhere else; it is designed to convey a sense of fondness across distance. Scent, similarly, releases a pleasant, sweet smell in similar circumstances providing an awareness of distant action.

The topic of "awareness" is one that has concerned the developers of technologies for group working, who want their systems to be able to support the casual and passive awareness of group activity that coworkers achieve in a shared physical space. Strong and Gaver turn this around, though, and give us technologies for supporting shared intimacy rather than shared work. Their pieces are designed to be evocative and emotive rather than "efficient." What is particularly interesting about this group of devices is that they originate not from a technical or scientific perspective, but from a design perspective. The result of this shift in perspective is that they a reflect a very different set of concerns. It is not simply that they reflect an aesthetic component where the scientific developments are marked more by engineering concerns. That is certainly one part of it, of course; the design examples certainly do reflect a different set of principles at work. However, there is more than this.

First, the design examples discussed here reflect a concern with *communication*. What is important is not simply what they *do,* but what they *convey,* and how they convey it; and the communicative function that they carry is very much on the surface. There is an "at-a-glance readability" to these artifacts that stands in marked contrast to the "invisibility" of ubiquitous computing. Second, they reflect a holistic approach that takes full account of their physicality. The physical nature of these pieces is not simply a consequence of their design; it is fundamental to it. While it was a tenet of ubiquitous computing, for example, that the technology would move out into the world, the design pieces reflect a recognition that the technology *is* the world, and so its physicality and its presence is a deeply important part of its nature. Third, they reflect a different perspective on the role of computation, in which computation is integrated much more directly with the artifacts themselves. In the other examples, while they have aimed to distribute computation throughout the environment, there has always been a distinct "seam" between the computational and the physical worlds at the points where they meet. In these examples, however, the computational and physical worlds are much more directly connected.

The result is an approach to tangible computing that sees computation within a wider context. Ubiquitous Computing pioneers saw that, in order to support human activity, computation needs to move into the environment in which that activity unfolds. These design explorations take the next step of considering how computation is to be manifest when it moves into the physical environment, and recognizing that this move makes the physicality of computation central.

Tangible Bits

Most recently, perhaps the most prominent site for development of these ideas has been the Tangible Media group at the MIT Media Lab. A group of researchers led by Hiroshi Ishii has been exploring what they call "Tangible Bits," a program of research that incorporates aspects of both the Ubiquitous Computing program and the design perspective explored by people like Jeremijenko.

The term "Tangible Bits" reveals a direct focus on the interface between the physical and virtual worlds. The rhetoric of the computer revolution has, pretty consistently, focused on a transition from physical (the world of atoms) to the virtual (the world of bits). We talk of the future in terms of "electronic cash" to replace the paper bills and coins we carry about with us, or we speak of the "paperless office" in which paper documents have disappeared in favor of electronic documents stored on servers and displayed on screens. We envision a world in which we communicate by electronic mail and video conferencing, in which we read from "e-books," telecommute over great distances via digital communication lines, and play in virtual worlds. What these visions have in common is the triumph of the virtual over the physical. They suggest that we will overcome the inherent limitations of the everyday world (such as the need to be in the same place to see each other, or that a thousand books actually take up real shelf space) by separating the "information content" from the physical form, distilling the digital essence and decanting it into a virtual world.

The MIT Media Lab, where Ishii and his colleagues are based, is one of the most prominent proponents of this vision, especially, perhaps, in the writings of its founding director, Nicholas Negroponte. His collection of essays *Being Digital* (Negroponte 1995), explores the relationship between atoms and bits and how the development and deployment of Internet technologies is changing that relationship.

The work on Tangible Bits provides some balance to the idea that a transition from atoms to bits is inevitable and uniformally positive. It is certainly not defined in opposition to the gradual and ongoing movement of traditionally physical forms into digital media. However, it observes that while digital and physical media might be *informationally* equivalent, they are not *interactionally* equivalent. By building information artifacts based on physical manipulation, the Tangible Bits programme attempts to reinvest these distilled digital essences with some of the physical features that support natural interaction in the real world.

metaDESK, Phicons, and Tangible Geospace

Let's take an example from the work of the Tangible Bits group. The metaDESK (Ullmer and Ishii 1997) is a platform for tangible interaction. It consists of a horizontal back-projected surface that serves as the top of the physical desk itself; an "active lens," which is a small flat-panel display mounted on an arm; a "passive lens," which is transparent, also digitally instrumented; and a variety of physical objects called *phicons* (for "physical icons"). The metaDESK is shown in figure 2.4.

The functions of the various components of the metaDESK platform are best seen in terms of an application running on the desk. Tangible Geospace is a geographical information system augmented with tangible UI features and running on the metaDESK. It allows users to explore a

Figure 2.4
Interactions with geographical information on the metaDESK, using phicons, the passive lens, and the active lens. Reprinted by permission of The MIT Media Lab.

visualization of a geographical space, such as the area of Cambridge, Massachusetts, around MIT.

The geographical information, in the form of a two-dimensional map, is back-projected onto the desk, so that the user seated at the desk can see it. The user can move and orient the map using phicons. One of the phicons represents MIT's Great Dome, and when it is placed on the desk, the map is adjusted so that the position of the Great Dome corresponds to that of the phicon. As the user moves the phicon, the system adjusts the map to ensure that the phicon is always aligned with the point on the map that it represents. By moving the phicon around on the desk, the user can cause the map to move too, "scrolling" around in the geographical space. By rotating the phicon on the desk, the user can cause the map to rotate.

If a second phicon is added to the desk, say one representing the Media Lab building itself, then another degree of freedom can be constrained. The two icons, together, can be used to control the scale of the map display. If the metaDESK always ensures that the virtual Great Dome always co-occurs with the Great Dome phicon, and the virtual Media Lab always co-occurs with the Media Lab phicon, then the user can control the scale of the map by moving these two phicons closer together or further apart.

The active and passive lenses can be used to provide access to other sorts of information. In the Tangible Geospace example, the active lens is used to view a three dimensional model of the MIT Campus. The active lens is a computer display mounted on an arm over the desk. It is instrumented so that the metaDESK computer system can determine the position and orientation of the display. When this information is coordinated with the current position, scaling, and orientation of the map being displayed on the desk, the result is that the active lens can be used to control a "virtual camera" moving through the geographical space being displayed on the metaDESK. When this is combined with a three dimensional model of the campus, then the active lens can be used to give a three-dimensional viewport onto the two-dimensional map. The illusion is of "looking through" the lens and seeing a transformed view of the map underneath.

The passive lens works in a similar way, although it rests on the desk surface. The passive lens is simply a piece of transparent plastic. As it is moved around the desk, the computer system can track its current location. On the desk area directly underneath the lens, the metaDESK replaces the map with a view onto a photographic aerial record of the campus. As before, this is correlated with the current position, scaling, and orientation of the basic map, as well as the position of the lens. The effect is that it seems to the user that the lens reveals the photographic model underneath as it moves across the desk. This is similar to a user interface technique known as "magic lenses" (Bier et al. 1993), user interface components that selectively transform the content of interfaces as they are moved across the screen, although, of course, in the case of the metaDESK the lens has a physical manifestation.

The Ambient Room

Tangible interfaces such as the metaDESK explore interaction that is situated in the environment, rather than on a screen. This is even more

clearly demonstrated by another of the MIT prototypes, called the Ambient Room (Wisneski et al. 1998).

The Ambient Room is a small office cubicle that has been augmented with a variety of "ambient displays," designed to provide peripheral, background information to the occupant of the room without being overwhelming or distracting. Examples of ambient displays include projected light patterns, non-speech sounds, and objects that respond to changes in air flow.

The information that the Ambient Room conveys is typically information about activities in either physical or virtual space, such as the presence or activity of others, e-mail arriving, people logging in and out, and so forth. These can be mapped onto the displays available in the room. For instance, light patterns projected on the wall can respond to the activities of a networked computer system, conveying information about network traffic and hence activity in the virtual space; or movements in a shared project room can be mapped onto subtle sounds in the Ambient Room so that the occupant can be aware of comings and goings in the project space. Reminiscent of the Feather, Scent, and Shaker work of Strong and Gaver, these ambient displays can be used to project the actions in one space (either physical or virtual) into another; like the technologies of the Reactive Room, they can also respond to the activity of the room's occupant, providing a display that is appropriate to the context in which they are working.

It is tempting to think of the metaDESK as exploring the potential for tangible media as input technologies, and the Ambient Room as exploring their potential for output. To do so, though, would be to miss an important point, which is that, in the everyday environment, "input" and "output" are fundamentally interconnected. This is a critical feature of the tangible media explorations. They should be characterized not in terms of "input" and "output," but in terms of the coordination between phenomena; between activity in a space and the pattern of light on a wall, or between the movement of objects on the desk and the information presented there. This sort of coordination, or coupling, is fundamental to the explorations presented here; they depend upon it for the causal illusion they want to maintain.

Illuminating Light and Urp

Two other applications developed in the MIT group echo the Digital Desk in their creation of mixed physical/virtual environments for task-focused work. These are Illuminating Light and Urp, both developed principally by John Underkoffler (and illustrated in figure 2.5).

Illuminating Light (Underkoffler and Ishii 1998) is a simulation of an optics workbench, aimed particularly at students of laser holography. The interface is based on a combination of phicons and a camera/projector arrangement (which Underkoffler dubs the "I/O Bulb") similar to that of the Digital Desk. The application allows users to experiment with and explore configurations of equipment for laser holography. Real laser holography is a complex business, conducted using delicate and expensive instruments. Setting up and fine-tuning an experimental configuration can be extremely time-consuming, especially for novices. Illuminating Light allows holographers to simulate the effects of particular configurations and to explore them so as to develop a better intuitive sense for the interaction of their elements. Phicons represent physical elements such as lasers, lenses, mirrors, and beam-splitters, while the system provides a simulation of light paths through the experimental equipment, showing light emitted by the laser, redirected by mirrors, and so on. As the phicons are moved around a physical surface, the system continually updates its projection of the simulated light paths to reflect the moment-by-moment physical configuration. In addition to the simulated light beams, the system can also provide numerical descriptions of the configuration; incidence angles, distances, and so forth. In this way, users can rapidly explore a variety of configurations and develop an understanding of the consequences of different changes on the set-up.

Urp (Underkoffler and Ishii 1999) is an urban planning workbench in which physical models of buildings are combined with electronic simulations of features such as air flow, cast shadows, reflectance, and so forth. The underlying technology is similar to that of Illuminating Light but applied to a different domain. There are two sorts of phicons used in Urp. The first represent building structures. By placing these on the surface, the user can obtain a visualization of the shadows that the buildings will cast, or the wind patterns around them. Combining multiple structures allows urban planners and architects to explore the

(a)

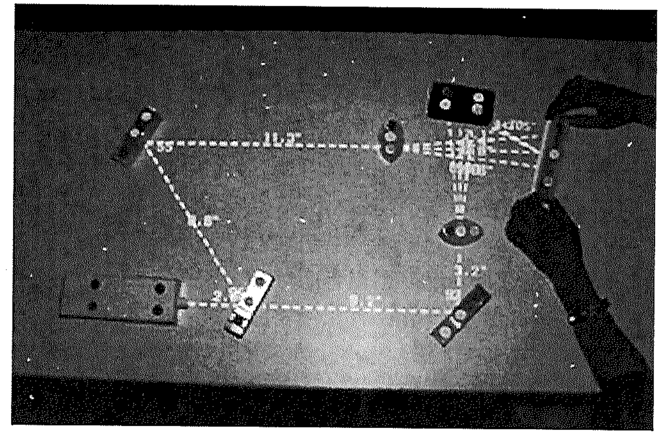

(b)

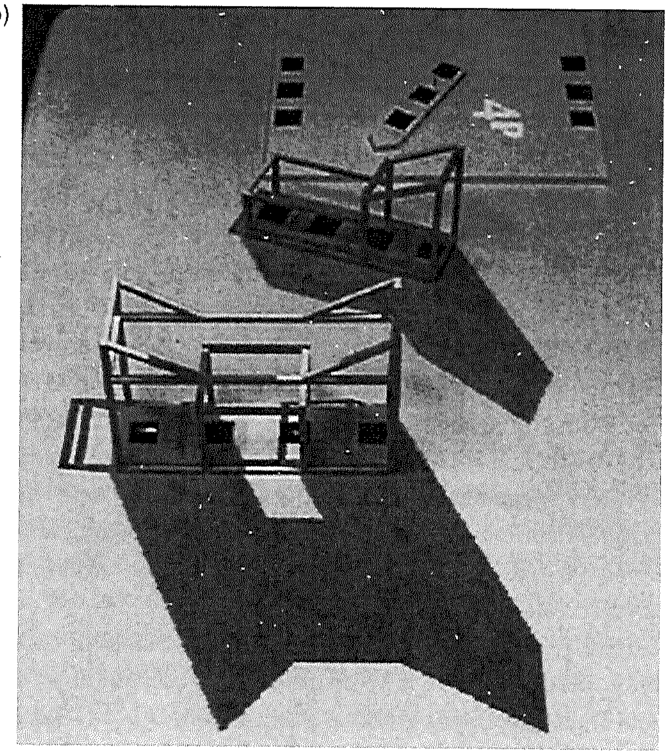

Figure 2.5
Illuminating Light (*a*) and Urp (*b*) apply tangible interaction techniques to the domains of optics and urban planning. Reprinted by permission of The MIT Media Lab.

interactions of wind, reflection, and shadow effects in an urban landscape. As with Illuminating Light, real-time tracking of the position and orientation of these phicons allows the system to update the display continuously, so that users can move the buildings around or rotate them until they find a satisfactory arrangement. The second set of phicons act as controls for the simulation. For example, a "wand" can be used to change the material of the buildings, so that the computed reflectance patterns will simulate buildings clad in brick or glass, another controls the direction of the simulated wind, while a "clock" has hands that can be moved to specify the time of day and hence the position of the sun for the shadow simulation. In this way, the simulator's controls are introduced into the same space that is the focus of the system's primary input and output.

Interacting with Tangible Computing

Tangible computing takes a wide range of forms. It might be used to address problems in highly focused and task-specific work, or in more passive awareness of activities in the real world or the electronic. It might attempt to take familiar objects and invest them with computation, or it might present us with entirely new artifacts that disclose something of the hidden world inside the software system. The bulk of this chapter has explored a range of tangible computing systems, but the survey has been far from comprehensive; indeed, I have said nothing about whole areas, such as wearable computing and context-based computing, that are clearly strongly related. My goal, however, was not to provide a catalogue of tangible computing technologies, but rather to introduce a sample of the systems that have been developed, and to begin to look for some common features of their design.

The first of these general issues that we see across a range of cases is that, in tangible computing, there is no single point of control or interaction. Traditional interactive systems have a single center of interaction, or at least a small number. Only one window has the "focus" at any given moment; the cursor is always in exactly one place, and that place defines where my actions will be carried out. Cursors and window focus insure that the system always maintains a distinguished component

within the interface, which is the current locus of interaction. To do something else, one must move the focus elsewhere. When computation moves out into the environment, as in the tangible computing approach, this is lost. Not only is there not a single point of interaction, there is not even a single device that is the object of interaction. The same action might be distributed across multiple devices, or, more accurately, achieved through the coordinated use of those artifacts. Imagine sitting at your desk to write a note. The writing comes about through the coordinated use of pen, paper, and ink, not to mention the desk itself and the chair you sit in; you might write on the page with your dominant hand while your nondominant hand is used to orient the page appropriately. These are all brought together to achieve a task; you act at multiple points at once. In the same way, ubiquitous computing distributes computation through the environment, and, at one and the same time, distributes activity across many different computational devices, which have to be coordinated in order to achieve a unified effect.

A related issue is how tangible interaction transforms the sequential nature of interaction at the interface. The single point of control that traditional interfaces adopt leads naturally to a sequential organization for interaction—one thing at a time, with each step leading inevitably to the next. This ordering is used both to manage the interface and to simplify system development. For instance, "modal" dialog boxes—ones that will stubbornly refuse to let you do anything else until you click "okay," "cancel," or whatever they need—both structure your interaction with the computer, and save the programmer from the need to handle the complexity of worrying about other actions that might transform the system's state while the dialog box is displayed. When we move from traditional models to tangible computing, sequential ordering does not hold. It is not simply that interaction with the physical world is "parallel" (a poor mapping of a computational metaphor onto real life), but that there is no way to tell quite what I might do next, because there are many different ways in which I might map my task onto the features of the environment.

These two issues are particularly challenging from a technical perspective, because they address the programming models we use to develop systems, embedded in software toolkits and applications. The third feature of tangible interaction may, however, provide some relief. This is the fact that, in tangible design, we use the physical properties of the interface to suggest its use. This is nothing new; arguably, it is what product design or other forms of physical design are all about. Kettles are designed so that we can tell how to safely pick them up; remote controls are designed to sit comfortably in the hand when oriented for correct use (at least when we're lucky). What is more, this sort of design that recognizes the interaction between the physical configuration of the environment and the activities that take place within it can also be a way to manage the sequential issues raised earlier. For instance, Gaver (1991), in his discussion of "sequential affordances" (which will be presented in more detail in chapter 4), gives the example of a door handle, which, in its normal position, lends itself naturally to turning and then, in its turned position, lends itself naturally to pulling; the whole arrangement helps "guide" one through the sequential process of opening the door through careful management of the physical configuration of the artifact. Taking this approach, designers can create artifacts that lead users through the process of using them, with each stage leading naturally to the next through the ways in which the physical configuration at each moment suggests the appropriate action to take. The relationship between physical form and possible action can give designers some purchase on the problems of unbounded parallel action.

Interacting with tangible computing opens up a new set of challenges and a new set of design problems. Our understanding of the nature of these problems is, so far, quite limited, certainly in comparison to the more traditional interactional style that characterizes most interactive systems today. The theories that govern traditional interaction have only limited applicability to this new domain. At the same time, tangible computing has been explored, largely, as a practical exercise. Most prototypes have been developed opportunistically, driven as much by the availability of sensor technology and the emergence of new control devices as by a reasoned understanding of the role of physicality in interaction. We have various clues and pointers, but there is no theory of tangible interaction. Why does tangible interaction work? Which features are important, which are merely convenient and which are simply

wrong? How does tangible computing mediate between the environment and the activity that unfolds in it?

This book is about developing answers to these questions. The interpretation that it will offer is one that is concerned not just with what kind of technology we use, or with what sorts of interactions we can engage in with that technology, but about what makes those interactions meaningful to us. From this perspective, the essence of tangible computing lies in the way in which it allows computation to be manifest for us in the everyday world; a world that is available for our interpretation, and one which is meaningful for us in the ways in which we can understand and act in it. That might seem to be quite far removed from looking at application prototypes, reactive rooms, and digital desks. The path from practice to theory will be easier to see after looking at the second aspect of embodied interaction—social computing.