

The metaDESK: Models and Prototypes for Tangible User Interfaces

Brygg Ullmer and Hiroshi Ishii

MIT Media Lab

Tangible Media Group

20 Ames St., E15-445

Cambridge, MA 02139 USA

+1 617 253 4121

{ullmer, ishii}@media.mit.edu

ABSTRACT

The metaDESK is a user interface platform demonstrating new interaction techniques we call “tangible user interfaces.” We explore the physical instantiation of interface elements from the graphical user interface paradigm, giving physical form to windows, icons, handles, menus, and controls. The design and implementation of the metaDESK display, sensor, and software architectures is discussed. A prototype application driving an interaction with geographical space, Tangible Geospace, is presented to demonstrate these concepts.

Keywords: tangible user interfaces, input devices, haptic input, augmented reality, ubiquitous computing

INTRODUCTION

The graphical user interface (GUI) has proven both a successful and durable model for human-computer interaction which has dominated the last decade of interface design. At the same time, the GUI approach falls short in many respects, particularly in embracing the rich interface modalities between people and the physical environments they inhabit. Systems exploring augmented reality and ubiquitous computing have begun to address this challenge. However, these efforts have often taken the form of exporting the GUI paradigm to more world-situated devices, falling short of much of the richness of physical-space interaction they seek to augment.

In this paper, we present research developing “Tangible User Interfaces” (TUIs) – user interfaces employing physical objects, instruments, surfaces, and spaces as physical interfaces to digital information. In particular, we present the metaDESK system (Figure 1), a graphically intensive system driven by interaction with graspable physical objects. In addition, we introduce a prototype application driving an interaction with geographical space, Tangible Geospace, to illustrate our approach.

The metaDESK effort is part of the larger Tangible Bits project [8]. The Tangible Bits vision paper introduced

the metaDESK along with two companion platforms, the transBOARD and ambientROOM. Together, these platforms explore both graspable physical objects and ambient environmental displays as means for seamlessly coupling people, digital information, and the physical environment.




Figure 1: The metaDESK system overview

The metaDESK system, shown in Figure 1, consists of several components: the *desk*, a nearly-horizontal back-projected graphical surface; the *active lens*, an arm-mounted flat-panel display; the *passive lens*, an optically transparent surface through which the desk projects; and an assortment of physical objects and instruments which are used on desk’s surface. These components are sensed by an array of optical, mechanical, and electromagnetic field sensors.

Our research with the metaDESK system focuses on the use of tangible objects – real physical entities which can be touched and grasped – as driving elements of human-computer interaction. In particular, we are interested in pushing back from the GUI into the real world, physically instantiating many of the metaphorical devices the GUI has popularized. (Figure 2, A) Simultaneously, we have attempted to push forward from the unaugmented physical world, inheriting from the richness of various historical instruments and devices often “obsoleted” by the advent of the computer. (Figure 2, B)

In addition, we more broadly explore the use of physical *affordances* [12] within TUI design. For example, our active lens is not only grounded in the metaphor of a jeweler’s magnifying lens; it also looks, acts, and is manipu-

Permission to make digital/hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication and its date appear, and notice is given that copyright is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires specific permission and/or a fee.

UIST 97 Banff, Alberta, Canada

Copyright 1997 ACM 0-89791-881-9/97/10 ..\$3.50

lated like such a device. In this way, the active lens has a certain legibility of interface in that its affordances suggest and support user’s natural expectations from the device.




Figure 2: Heritage of Tangible User Interfaces

In the following sections, we present our design approach towards making user interfaces tangible. The operating scenario of the Tangible Geospace prototype is then presented. This is followed by a description of the metaDESK implementation, including display, sensor, and software architectures. Interaction issues encountered with the prototype are then discussed, followed by future work and conclusions.

DESIGN APPROACH

The GUI “desktop” metaphor of icons, windows, handles, and controls borrows in significant part from metaphors of physical space – i.e., the metaphor of the physical desktop. We have taken the GUI desktop metaphor itself as a kind of metaphor, in our work physically instantiating the GUI window, icon, menu, handle, and control metaphors back into the real world (Figure 2, A).

Figure 3 illustrates this mapping of GUI widgetry into physical space. We give the GUI “window” device substance as a physical “lens” which may be hand-held, arm-mounted, or placed upon another surface.




Figure 3: TUI instantiations of GUI elements

We physically instantiate GUI “icons” as TUI “phicons” (physical icons) with varying levels of representational abstraction. GUI “menus” and “handles” are instantiated as TUI “trays” and “phandles” (physical handles). Trays and phandles have been introduced previously in [5]. Finally, we give physical form to GUI controls like scales and scrollbars as TUI instruments such as the rotation constraint instrument pictured in Figure 7.

In addition to physically instantiating metaphorical devices from the GUI, we have also worked to push forward from physical devices that the GUI left behind (Figure 2, B). We have found inspiration from generations of richly afforded physical instruments which have been “obsoleted” in the ascent of the personal computer.

Historical scientific instruments, drawing and design tools, and everyday objects from home and trade often display a range of rich interface affordances reflecting generations of human use and refinement – affordances potentially exploitable within user interface design. Our rotation constraint instrument of Figure 7, for instance, draws in form and function from the rich legacy of scientific instruments, while offering additional affordances for visual and haptic computer augmentation.

It is also worth being clear at the outset that we do not believe tangible user interfaces are about “replacing” graphical user interfaces. Rather, we seek to complement GUIs by embracing the richness of the physical environment and providing new opportunities for human-computer interaction in domains poorly supported by current interface approaches.

RELATED RESEARCH

A variety of research efforts have explored computationally-augmented interfaces emphasizing human interaction with the physical world. The Bricks work of Fitzmaurice, Ishii, and Buxton [5] is most directly related to the metaDESK. The Bricks “graspable user interface” research involves placing one or more bricks – abstract physical blocks tracked with 6DOF (six degrees of freedom) – onto some screen-based virtual object, b-spline control point, etc. Bricks can then be used to physically rotate, translate, or (using multiple bricks in combination) scale and deform the “attached” virtual entities by manipulating the proxying brick devices.

Bricks are an example of what we call physical handles or “phandles,” physical instantiations of the GUI “handle” device. We have gone beyond Bricks by creating a repertoire of physical devices fleshing out an approach for TUI design, including introduction of the phicon, lens, and instrument devices.

The research themes of augmented reality and ubiquitous computing are also important motivators to the TUI approach, but are marked by important differences. In particular, augmented reality research [4] has generally been directed towards visual augmentations of physical spaces through head-mounted or hand-held displays, where our work focuses on direct physical interaction with objects as elements of TUI interfaces.

The pioneering Xerox PARC work in ubiquitous computing [17] by Weiser et al. explored manners by which many computational devices could be distributed and integrated within the physical environment. A key insight of this work is the notion that computation might be spread across many

devices optimized for diverse tasks throughout the physical environment. At the same time, the original Tab, Pad, and Board devices each retained GUI-style interfaces, offering less specific insight towards integrating richer physical interface modalities.

Several other specific efforts directly address related desk- and lens-style user interfaces relating to the metaDESK. Wellner’s DigitalDesk [18] supports augmented interaction with physical paper documents on a physical desktop, identifying and augmenting these with overhead cameras and projectors. A groundbreaking example of seamlessly coupling virtual and physical worlds, DigitalDesk focused on interaction with paper and paper-based tasks, with less attention to more diverse physical interfaces and domains.

The Responsive Workbench [9] and Immersadesk [3] provide another interesting approach, each supporting interaction with a large graphical surface viewable in stereo 3D using LCD shutterglass eyepieces. However, both platforms limit interaction to a virtual environment, with no integration of physical-space interfaces beyond traditional position tracking devices.

Several research systems make use of devices similar to our active lens. The PDDM device [11] is an arm-mounted flat-panel display with force-feedback operating in the proximity of a wall projection display. The small display of the PDDM is used for “grasping” virtual objects with force feedback in a VR scene.

Another system, the NaviCam [13], uses a hand-held lens-like device for navigating physical spaces such as a bookshelf. The NaviCam follows a more conventional augmented reality approach of visual augmentation, with little consideration for the physical manipulation of augmented objects.

The Magic Lens and Toolglass research [1] provides a compelling example of widgets which graphically transform a region of a GUI desktop to provide alternative visualizations. The original Magic Lens was limited to a software GUI element manipulable with the mouse and trackball. The behavior of our passive lens device is inspired in part by the Magic Lens work.

Several other works focus on manipulation of physical objects as interface. The passive interface props work of Hinkley et al. [6] uses physical props (e.g. head viewing prop, cutting-plane selection prop) as tools for manipulating 3D models within a GUI surgical interface.

Bishop’s marble answering machine [2] makes compelling use of marbles as a physical embodiment of voice messages. Finally, Stifelman’s paper-based audio notebook [14] provides a powerful example of using an augmented physical object (i.e., a paper notebook) in an unusually natural, legible, and useful fashion.

PROTOTYPE SYSTEM

To make our TUI ideas “tangible,” we have implemented a prototype application on the metaDESK called “Tangible Geospace.” Tangible Geospace is designed around a TUI interaction with geographical space. The name partially reflects earlier geographic visualization research titled “GeoSpace” [10] by Lokuge and Ishizaki in the MIT Media Lab. A description of the Tangible Geospace operating scenario follows. A video presenting this interaction accompanies the paper.

Tangible Geospace Scenario

Several physical objects and instruments for interacting with geographical space sit in a translucent holding tray on the metaDESK’s surface.

By placing a small physical model (phicon) of MIT’s Great Dome onto the desk, a two-dimensional map of MIT appears underneath on the desk, bound to the Dome object at its location on the map. (Figure 4) The Dome phicon was constructed out of transparent machined acrylic, designed to minimize occlusion of the desk surface, to provide opportunities for visually augmenting the phicon, and to aesthetically enhance continuity between the physicality of the phicon and the virtuality of the desk-based map display.




Figure 4: Great Dome phicon in Tangible Geospace

Simultaneously, the arm-mounted active lens (Figure 5) displays a three-dimensional view of MIT with its buildings in perspective. The active lens is coupled to the Dome-based campus model, such that moving the lens navigates the 3D space in a manner consistent with an optical lens metaphor (allowing viewpoint translation, zooming, etc).

The Dome phicon acts both as a container for the digital information about MIT, as well as a physical handle (handle) for manipulating the map. By rotating or translating the Dome object across the desk’s surface, both the 2D desk-view and 3D lens-view are correspondingly transformed. The user is thus seamlessly interacting with three spaces at once – the physical space of the Dome object; the 2D graphical space of the desk’s surface; and the 3D graphical space of the active lens.



Figure 5: Active lens in Tangible Geospace

The user next takes a second physical icon from the holding tray, this time representing the Media Lab building, and places it onto the surface of the desk. Now there are two physical constraints and handles for the MIT space, allowing the user to scale or rotate the map by moving one or both objects with respect to each other. The user may grasp and manipulate both phicons simultaneously with two hands. Alternatively, two users may independently grasp one of the building phicons, cooperatively manipulating the geospace.

In this manner, there is no one locus of control, unlike the GUI's use of the mouse. Rather, the interaction is constrained by the physics of the physical environment, supporting multiple pathways of single- and multi-user interaction likely unrealizable with the mouse-based paradigm. This interaction is similar to an interaction within the Bricks work [5], but differs in the object-semantics attributable to the phicons.

With the geospace active on the metaDESK, we have implemented several physical instruments for further manipulating and viewing the space. First, we have created an instrument called the "passive lens," (Figure 6) a wood-framed transparent surface that functions as an independent display when augmented by the back-projected desk. Since passive lens devices are passive transparent surfaces, many variously afforded lenses might be used simultaneously with no additional active display resources (i.e., additional computer-driven screens).

Using the passive lens, the user may interact with a secondary overlay view of the MIT campus. We used an aerial orthographic photograph as the inline view, providing an interesting representational contrast with the hand-rendered map and polygonal 3D model of the desk and active-lens surfaces. The passive lens also conceptually supported other overlays views consistent with physical instantiation of the Magic Lens metaphor [1].



Figure 6: Passive lens, with Great Dome phicon

When manipulating Dome and Media Lab phicons on the desk to scale and rotate the geospace, the simultaneous rotation of both phicons poses an ambiguous interaction. The software resolution of this ambiguity is discussed in the Interaction Issues section. As an alternative to the two-phicon scaling/rotation interaction, we implemented a rotation constraint instrument made of two cylinders mechanically coupled by a sliding bar (Figure 7). This instrument allows scaling and rotation manipulation of the two geospace control points, while preventing the ambiguous two-phicon rotation by the intrinsic mechanical constraint.



Figure 7: Rotation constraint instrument

IMPLEMENTATION

The metaDESK hardware architecture is illustrated in Figure 8. The largest component is the desk itself, a back-projected near-horizontal graphical surface used to display 2D geographical information within the Tangible Geospace prototype. Above the desk, an arm-mounted flat-panel display serves as the "active lens" used to display 3D geographical information. In addition, several passive physical objects are manipulated by the user on the surface of the desk. Sensing in the system is performed by a computer-vision system inside the desk unit, along with magnetic-field position sensors and electrical-contact sensors. Three networked computers are used to coordinate the system as a whole.

Display Architecture

The desk back-projection unit is based on the Vision-Maker™ product from Input Technologies. It uses a three-tube projector to display computer graphics via two internal mirrors onto a plexiglass diffuser surface, which forms the near-horizontal display surface of the desk. The display resolution is 1280x1024 pixels. While a completely horizontal display surface is desirable for supporting physical objects without slippage, the VisionMaker has a minimum display angle of 12 degrees from the horizontal. We cover the display surface with a clear plastic film that minimizes object slippage.

Mounted alongside the desk is the active lens, an arm-mounted flat panel display. The active lens is supported with the arm-mount of a jeweler's magnifying lens. The optical lens was removed and replaced with a specially mounted 25-cm 640x480 pixel LCD TFT color flat panel display. The display-mount has three degrees of freedom (DOF), while the arm support itself has another three DOF. An Ascension Flock of Birds™ 6DOF magnetic-field position sensor was attached to the flat-panel display for tracking its spatial position and orientation.

The desk and active lens displays are driven by Intel Pentium Pro and SGI Indigo2 computers using Intergraph Intense3D-T and SGI Extreme 3D graphics accelerators. Graphics display was managed by an [incr Tcl]-based extension called 3wish [16], which among other things provided a platform-independent scripting interface to the SGI and TGS Open Inventor 3D graphics toolkits.

The active lens is used to display a navigable 3D model of MIT campus buildings (~30K polygons), while the desk displays a 2D image of MIT campus. The 2D map was displayed as a texture-mapped polygon; while only manipulated as a 2D image, at 1140x500 pixels or 570K square pixels, realtime rotation and scaling required hardware texture acceleration.




Figure 8: metaDESK System hardware architecture

The 3D active lens view is from a virtual camera positioned orthogonal to the physical active lens surface. Because it is

desirable to move the active lens scene-camera both closer to and further from the MIT campus scene than allowed by the arm assembly's physical constraints, the distance of the virtual camera from the active lens surface follows an exponential curve with empirically-derived coefficients. Registration between the active lens and desk displays is discussed in the Interaction Issues section.

The passive lens is made of a 12cm diameter, 1cm thick circle of fiber-optic cluster material, ringed with a wooden frame and handle. Early iterations of the passive lens were tested both with a plexiglass "lens" and with an empty frame without interior surface. However, the intention with the passive lens was to give the illusion of an independent display surface, i.e. a separate screen. Neither the plexi-glass nor empty-frame approaches sustained this illusion.

By raising the image of the back-projected display to the upper surface of the passive lens, the fiber-optic cluster material succeeded in visually simulating an independent display surface. The wooden frame assisted this illusion by masking minor alignment errors between the back-projected passive-lens screen image and the fiber-cluster viewing portal.

During movement of the passive lens across the desk, the severity of these errors is a function of tracking technology and graphics frame rate. Originally computer-vision was used to track the passive-lens, but later a Flock of Birds sensor was used to provide faster, more precise graphics updates.

The passive lens display was generated from a second 2D map of MIT campus, texture-mapped onto a near-circular polygon on the metaDESK. The texture coordinates of this polygon were dynamically calculated to accommodate the position and scale of the phicon-manipulated campus map.

Sensor Architecture

Part of our design aesthetic with the metaDESK was to imagine every object in the physical environment as a potential container or handle for digital information, and thus a potential interface on the metaDESK. Consequently, we were especially interested in using passive, minimally tagged physical objects as TUI controls whenever possible. Towards these ends, the Media Lab and Great Dome phicons and the rotation-constraint instruments were designed to be computationally passive (free from active electronics).

Tracking of these passive interface objects is achieved with computer vision. The desk is augmented with two cameras mounted inside its chassis aimed at the back-projected diffuser display surface from underneath. This is illustrated in Figure 8. This camera geometry allows objects to be monitored as a largely 2D vision problem free from hand and body object-occlusions. In addition, this approach realizes a modest user-privacy gain in that objects more than ~10cm above the desk surface are invisible to the cameras because of the diffuser coating.

One of the two internal desk cameras is a computer-controlled pan-tilt-zoom Canon VC-C1, used for initial tests at visible-light object tracking and identification. Objects on the surface of the desk were illuminated with pixels from the back-projected desk display, in a sense transforming the display+camera geometry into a flatbed scanner.

This proved unsatisfactory for general object tracking because of interference between graphical output and camera input. To allow physical objects to be tracked, the second camera is used for computer vision in the infrared optical regime. We illuminate objects on the desk's surface with two security-camera IR LED-arrays mounted within the desk, and monitor the resulting scene with a monochrome video camera outfitted with an infrared filter.

This approach cleanly filtered the projected computer graphics from camera view because of the minimal infrared component of the projected graphics. In addition, controlled object illumination is generated by the invisible infrared lamps, with the added bonus that fluorescent room lighting produces minimal interfering infrared emissions.

The software of the infrared vision system has two layers: a "tag-track" low-level vision layer, and an upper layer which filters noise-induced tags, identifies objects, and tracks objects from frame to frame. The tag-track vision layer was implemented with background-subtraction vision software by Thad Starner of the MIT Media Lab Vision and Modeling group, executed on an SGI Indigo2 R4400-250MHz + Galileo video digitizer at seven frames per second. This software provides for each vision frame a list of unidentified "blobs" extracted from the scene.

Our selection of transparent acrylic Media Lab and Great Dome phicons, while interesting given their visual continuity with the desk display, was challenging because the objects were (not surprisingly) nearly invisible to both infrared and visible-light cameras. We addressed this by backing the phicons with "hot mirrors," material which is reflective to infrared but transparent to visible-light. This approach allowed objects to be tracked in the infrared while retaining visible transparency.

Objects are alternatively identified with a resistor tag electrically identifiable by an SGI-based LEGO Dacta Control Lab™ attached to compartments of the metaDESK's tray; or by assigning objects to fixed tray compartments which are monitored for contents with electrical switches or the in-desk vision system. By recording when these (labeled) objects leave the tray and applying this label to the next unlabeled object appearing in the infrared vision scene, we are able to successfully identify and track physical objects on the desk, maintaining a seven frame-per-second tracking rate.

Software Architecture

The metaDESK system is designed to support the operation of tangible user interfaces – interfaces using physical objects as interfaces to digital information. We developed a

software architecture to sustain this model of physical objects as interface, modeling each object as having certain "capabilities" for sensing its physical state, displaying computed outputs, and communicating with other objects.

The metaDESK demonstrates the use of objects as interface with devices including the active lens, passive lens, and phicons. These TUI elements illustrate diverse incarnations of sensing and display capabilities. The active lens contains a real sensor and real display. The passive lens incorporates a real sensor and a "virtual display." Finally, our phicons in a sense employ "virtual sensors" and virtual displays. This is illustrated in Table 1.

As this example implies, the technical realization of TUI elements can become complex. In implementing Tangible Geospace on the metaDESK, we faced the software architecture challenge of cleanly realizing our desired implementational metaphor while managing the complexity of a working system.

TUI objects	sensor	display
active lens	real (Flocks of Birds)	real (flat panel LCD)
passive lens	real (Flocks of Birds)	virtual (desk-augmented fiber-cluster surface)
phicons	virtual (computer vision)	virtual (desk-augmented display)

Table 1: Real and virtual object sensors/displays

We approached the metaDESK software architecture in two fashions. First, we implemented a platform-independent scripting meta-language, 3wish [16], with extended display, sensor, and distributed computing support. Secondly, we designed a software architecture called proxy-distributed or "proxdist" computation, which provides mechanisms for proxying sensor and display capabilities on behalf of passive physical objects. [15]

The metaDESK software is implemented with 3wish [16], a set of physical sensor and actuator, 3D graphics, and distributed computing libraries built with the "[incr Tcl]" object-oriented extensions to the Tcl scripting language. 3wish currently operates on SGI Irix and Intel-based Windows NT platforms. Cross-platform graphics are implemented using the TGS port of Open Inventor, while Tcl's shared object loading facilities are used for accessing platform-dependent sensor and computer-vision libraries. 3wish's Tcl based graphics implementation was valuable in allowing interface code to be executed transparently across multiple platforms.

Secondly, the "proxdist computation" approach allows us to provide proxied sensing and display capabilities for Tangible Geospace's varied interface objects. This approach is illustrated in Figures 9 and 10. Tangible Geospace is built around seven interface objects: the active and passive lenses, desk, tray, Dome and Media Lab phicons, and rotation-constraint instrument. While each object acts as an input device and four serve as displays (the lenses, desk, and tray), the underlying sensing and display "capabilities"

are normally dispersed across the hardware of the metaDESK system.

Figure 9 illustrates the proxdist architecture for the metaDESK. The metaDESK’s physical architecture is divided into a series of physical objects, sensors, displays, and computers, each represented as layers on the diagram’s left.

On the right side of the diagram, the software architecture is displayed. The software architecture includes client/server layers for networked communication with hardware sensors and displays. In addition, sensor/display proxy layers render transparent to the individual object proxies which physical technologies are used to realize the virtual sensing/display capabilities of each interface object. Finally, the namespace client/server pair is used to manage naming abstraction and to coordinate distributed systems resources.

The intention of the proxdist architecture is to provide an API for each physical interface object which cleanly abstracts the object’s virtual sensing and display capabilities. Proxdist computation makes transparent whether sensors, displays, and supporting computers are local to a TUI interface object, or proxied on behalf of the object by the surrounding environment.

The net effect of this design can be found by noting several aspects of Figure 9’s proxdist diagram. First, each physical world interface object is paralleled by a digital world object proxy. Next, the interconnection between physical interface objects and their augmenting sensors/displays is quite complex, reflecting the intricacies of interfacing with the physical world. Simultaneously, the interconnections between the interface object proxies and the sensor/display proxies remains simple, abstracting away both underlying sensor/display technologies and networked machine dependencies.




Figure 9: Proxdist metaDESK architecture

This layered hiding of physical-world and machine-dependency complexities from the digital-world object proxies is a key feature of proxdist architectures, for it allows even completely passive objects to be treated as

having local sensing, display, communication, and computation capabilities.

To provide a concrete example, we will consider the case of the passive lens on the metaDESK. The passive lens device is tracked with a physical sensor, a Flock of Birds device, and derives its “virtual” display through back-projection by the metaDESK. The physical sensors/displays are addressed by the names

```
tmg:metadesk::sensor:flock
```

and

```
tmg:metadesk::display:desk
```

These names, representing (in the Flock case) zone “Tangible Media Group,” encapsulation “metaDESK,” “sensor” capability of type “flock”, are associated with [incr Tcl]-based sensor/display API’s, and resolved to server host/port TCP/IP addresses through the namespace server.

The Tangible Geospace application does not directly reference the above sensor/display clients. Instead, references are made to sensor and display proxies named

```
tmg:metadesk::sensorproxy:plens
```

and

```
tmg:metadesk::displayproxy:plens
```

The sensor proxy abstracts whether computer vision, the Flock of Birds, or some alternate technology is used for passive lens position tracking. In addition, the Flock sensor is physically mounted at a different position than the functional center of the passive lens device. The sensor proxy returns the functional position, again providing insulation from the complexities of physical-world implementation.

Similarly, the display proxy abstracts the display device to which passive lens imagery is rendered. This (ideally) makes it invisible to the application whether the passive lens device is a fully independent display like the active lens, or a device deriving its functionality from environmental displays as with the passive lens.




Figure 10: Tangible Geospace software architecture

The above-described sensor client/server support has been fully implemented, and simple sensor proxies have been coded for the Tangible Geospace prototype. Sensor clients communicate with servers using a non-blocking “request value/provide value” protocol, with sensors objects caching update values locally.

Sensor clients optionally interpolate server value updates. This is used in the case of computer vision to smooth slow updates, at the cost of slightly increased latency. More sophisticated future clients will hopefully employ dead reckoning or predictive filtering techniques.

The display server/client/proxy layers are more complex than for sensors. Sensors can be thought of as open-loop data sources, streaming sensor readings through servers to sensor clients and proxies. On the other hand, the metaDESK’s graphical displays are closed-loop data sinks which input scene geometries and slave updates to the completion of each graphics frame rendering.

Therefore in our implementation, messages to display proxies include 3wish code bound for server-side execution. This is considerably more complex than the sensor proxy case. Instead, the display server/client/proxy hierarchy is currently approximated by executing Tangible Geospace display proxies locally on the desk and active lens display computers. Thus, while the sensor server/client/proxy hierarchy is fully implemented, completion of the display server/client/proxy hierarchy remains to future work.

INTERACTION ISSUES

Embodiment of multiple display viewpoints

One issue we explored in the metaDESK implementation is the complementarity between multiple embodied scene representations and viewpoints. Navigation of immersive 3D scenes is a difficult interface problem. The embodiment of multiple display viewpoints on the desk, active lens, and passive lens devices provided an interesting alternative navigational approach. When attempting to reach a new 3D viewpoint with the active lens, the user may locate the target viewpoint and orientation with a glance to the 2D desk map. The user may then rapidly move the active lens to the new viewpoint, relying on kinesthetic cues and the 2D scene instead of continuous navigation of the 3D landscape.

It is worth noting that especially with the active lens, scene registration between lens, desk, and user point-of-view is an issue. The active lens displays a 3D scene from a camera viewpoint orthogonal to the surface of the lens’ physical display, not as computed from the vantage of the user’s eye position. In contrast, the Responsive Workbench [9] and Immersadesk [3] derive camera viewpoints by tracking the head-position of a single controlling user. However, this approach is usually encumbering, and generally fails to gracefully support multi-user viewing and interaction.

In choosing to align camera viewpoint with the active lens’ orientation, we adopted an approach similar to that of digital cameras with LCD backfaces designed to decouple eye-position from camera-viewpoint. From the standpoint of this alternate metaphor, consistency of the active lens viewpoint is achieved, at the same time intrinsically allowing multiple users to share the lens view (assuming sufficient flat panel angle-of-view).

Phicon level of abstraction

In the Design Approach section, we introduced the notion of TUI physical icons or “phicons” as a parallel to GUI icons. The Tangible Geospace scenario provided examples of this interface device with the Great Dome and Media Lab landmark phicons. These phicons are relatively literal in that their physical form implies a specific geographical association. At the same time, one can easily imagine other more abstract landmark phicons. While Tangible Geospace currently realizes only literal phicons, it is interesting to outline other possibilities for exploration in future work.

Houde and Salomon explore the progression from abstract to specific GUI icon representations in [7], along with a discussion finding inspiration in the affordances of physical objects. In a parallel fashion, Figure 11 portrays a progression from generic phicons perhaps equivalent to the Bricks physical handles of [5] to literal phicons like the Great Dome and Media Lab. A fourth progression, that of “actualities,” includes a class of objects like family heirlooms which have some set of prior associations caught up in their physical form.




Figure 11: Comparison of data icons with phicons

Interface consistency and ambiguity resolution

The Tangible Geospace scenario mentions an issue of ambiguity arising from user rotation of the Great Dome or Media Lab phicons in the two-phicon case. The position information alone of two phicons exactly constrains the configuration of a rigid map. Given the additional rotation constraint, should the map be warped, the offending inputs ignored, or some other response taken? Our implementation ignores the offending rotation, but other user expectations might be reasonable. Moreover, the situation becomes even more complex with a three-phicon configuration, as illustrated in Figure 12.

Given this range of reasonable interpretations for a given phicon configuration, it is reasonable to question how the system should resolve this ambiguity. In coming to terms with ambiguity, we have attempted to leverage the use of context, constraints, and closure of interaction. The two-phicon scaling/rotation interaction in Tangible Geospace provides an example of the use of context. The conjoining

of these phicons defines a geographical context, such that the space may be manipulated by moving these objects with respect to each other.




Figure 12: Three-phicon alternatives – (a) warp to fit; (b) flag and ignore outlier; (c) best fit with error display; (d) discrete per-phicon views

While the landmark phicons partially constrain the map transformation, there is ambiguity in interpreting the rotation of one or both objects. The rotation constraint instrument (Figure 7) provides an interesting example of resolving this ambiguity by constraining the transformations expressible with the instrument’s mechanical structure.

Still, there are cases when neither context nor constraint uniquely disambiguates an interaction. Here, we fall back to the closure of interaction between the user and interface, visually highlighting the locus of ambiguity and allowing for further contextualization, constraint, or re-expression. Case B of Figure 12 illustrates this approach, with the third phicon visually highlighted by the desk as “uninterpretable” without further context or constraint.

FUTURE WORK

The metaDESK’s Tangible Geospace prototype is intended as a working proof-of-concept both motivating and illustrating notions of tangible user interfaces. It is designed to tangibly demonstrate and embody a repertoire of new interaction techniques, and is not engineered as a targeted end-user application.

In future work, we look forward to developing this research in a variety of respects. We are interested in exploring phicons with different levels of abstraction and persistence of digital association, and developing means for providing “closure” with users in the presence of ambiguous or inconsistent physical events. Expanded research in the context of Tangible Geospace would also benefit from targeting upon a particular user community, as well as obtaining at least qualitative interface feedback from such users.

Beyond Tangible Geospace, we are developing other applications on the metaDESK which use phicons, lenses, and instruments in non-geographical contexts. We are interested in new applications using combined 2D and 3D graphical surfaces for information without native visual or physical form. We are also eager to explore interfaces which manage multiple applications on the metaDESK. We suspect this will require going beyond graphical partitions of the metaDESK workspace.

At the same time, our concern is less with Tangible Geospace or even the metaDESK per se, and more with the broader notion of TUIs as physical interfaces to digital information. In this context, developing evocative and intuitive models and metaphors for coupling physical and digital spaces is a primary concern. While the GUI desktop metaphor provides an interesting and productive first step, it is only one of many possible TUI models. The metaDESK’s combination of lenses hint towards an interesting “optical” metaphor [8], which along with other models is a focus of ongoing work.

CONCLUSION

We have introduced the metaDESK system, a user interface platform supporting physical interaction with digital information through the manipulation of physical objects, instruments, and surfaces. The metaDESK gives physical form to graphical user interface (GUI) devices, physically instantiating icons, windows, menus, handles, and controls as phicons (physical icons), lenses, trays, phandles (physical handles), and instruments. This use of physical objects as interfaces to digital information form the basis for tangible user interfaces (TUIs).

The metaDESK platform uses tangible objects and instruments manipulated upon a near-horizontal display surface internally monitored with infrared computer vision. This interface surface is complemented with the arm-mounted “active lens” flat-panel display, as well as the fiber-optic-bundle “passive lens” device which acts as an independent display through augmentation by the back-projected desk. Use of the metaDESK is demonstrated with the Tangible Geospace prototype application, a tangible user interface driving interaction with geographical space.

We believe the metaDESK system embodies a novel and useful interface environment for mixed physical- and digital-space interaction, as well as an approach to user interface with much broader applicability. By bringing the tangibility of the physical world to computing, we see new opportunities for human-computer interaction beyond GUI.

ACKNOWLEDGMENTS

Many people have contributed to the development of the metaDESK system. In particular, we would like to thank undergraduate research assistants Minpont Chien, Dylan Glas, and Chris Lopes for work on the Tangible Geospace software, physical prototyping, and computer vision software, respectively; Thad Starner, Ali Azarbayejani, et al. of the Vision and Modeling group for help with the visible-light and IR computer vision hardware and software; Bill Verplank of Interval Research for suggestion of the rotation constraint instrument; Michael Owu of the MIT Campus Planning office for a 3D digital model of the MIT campus; MassGIS and MIT for orthographic images of MIT and surrounds; and the graduate students of the Tangible Media group, among others, for research support and feedback on the draft. We also thank the Media Lab Things That Think consortium for support of this work.

REFERENCES

1. Bier, E., Stone, M., Pier, K. et al. (1993). Toolglass and Magic Lenses: The See-Through Interface. *Computer Graphics (SIGGRAPH'93 Proceedings)*, pp. 73-80.
2. Crampton Smith, G. (1995). The Hand That Rocks the Cradle. *I.D.*, May/June 1995, pp. 60-65.
3. Czernuszenko, Marek. (1996). Immersadesk web page. <http://www.eecs.uic.edu/~mzczernus/Immersadesk/-immersadesk.html>.
4. Feiner, S., MacIntyre, B., Haupt, M., and Solomon. (1993). Windows on the World: 2D Windows for 3D Augmented Reality. *Proc. of UIST'93*, pp. 145-155.
5. Fitzmaurice, G., Ishii, H., and Buxton, W. (1995). Bricks: Laying the Foundations for Graspable User Interfaces. *Proc. of CHI'95*, pp. 442-449.
6. Hinckley, K., Pausch, R., Goble, J., and Kassel, N., (1994). Passive Real-World Interface Props for Neurosurgical Visualization. *Proc. of CHI '94*, pp. 452-458.
7. Houde, S., and Salomon, G. (1993). Working Towards Rich & Flexible File Representations. *Proc. of INTERCHI'93, Adjunct Proc.*, pp. 9-10.
8. Ishii, H., and Ullmer, B. (1997). Tangible Bits: Towards Seamless Interfaces between People, Bits, and Atoms. *Proc. of CHI'97*, pp. 234-241.
9. Krueger, W., and Frohlich, B. (1994). The Responsive Workbench. *IEEE Computer Graphics and Applications*, 14(3), pp. 12-15.
10. Lokuge, I., and Ishizaki, S. (1995). GeoSpace: An Interactive Visualization System for Exploring Complex Info. Spaces. *Proc. of CHI'95*, pp. 409-414.
11. Noma, H., Miyasato, T., and Kishino, F. (1996). A Palmtop Display for Dextrous Manipulation with Haptic Sensation. *Proc. of CHI'96*, pp. 126-133.
12. Norman, D. (1988). *The Psychology of Everyday Things*. New York: Basic Books.
13. Rekimoto, J., and Nagao, K. (1995). The World through the Computer: Computer Augmented Interaction with Real World Environments. *Proc. of UIST'95*, pp. 29-36.
14. Stifelman, L. (1996). Augmenting Real-World Objects: A Paper-Based Audio Notebook. *Proc. Of CHI'96 Companion*, pp. 199-200.
15. Ullmer, B. (1996). Behavioral Realizations of Proxy-Distributed Computation. <http://www.media.mit.edu/~ullmer/courses/agents/paper1.html>, March 1996.
16. Ullmer, B. (1997). 3wish: Distributed [incr Tcl] Extensions for Physical-World Interfaces. To appear in *Proc. of Tcl/Tk'97*.
17. Weiser, M. (1991). The Computer for the 21st Century. In *Scientific American*, 265(3), pp. 94-104.
18. Wellner, P. (1993). Interacting with paper on the Digital Desk. In *Comm. of the ACM*. 36(7), pp. 86-96.