Information Mobility Between Virtual and Physical Domains

Garth Shoemaker
Department of Computer Science
University of British Columbia
garths@cs.ubc.ca

ABSTRACT
Humans are ideally suited to interaction with physical information artifacts. These capabilities should be supported by computing systems, and users should be able to move seamlessly between interacting with virtual artifacts and equivalent physical artifacts. Unfortunately, existing computing systems are centered almost solely around virtual information interaction, with interaction metaphors developed specifically for computing systems. It is important that future computing systems evolve to properly support data mobility between physical and virtual domains, and as a result, natural human interactions. We perform a survey of existing research relevant to data mobility, identify key problems which need addressing, and speculate on future research that may help alleviate the problems identified.

Author Keywords
information mobility, human-computer interaction, physical interaction, collaboration

INTRODUCTION
We live in a society where computing technology has become ubiquitous. Every modern office has personal computers to support most job functions. Specialized computer tools are used in engineering and mechanical fields. Information of all kinds is flowing continuously over the Internet from site to site to support everyday tasks in and out of the workplace. Email, PowerPoint files, instant messages, Word documents, graphs, and diagrams all play critical roles in modern business functions.

However, despite the near universal adoption of computing technology, the predicted paperless office has failed to materialize [32]. In fact, paper use per office worker has continued to increase without pause [17].

The dichotomy of pervasive access to computing and the persistence of paper as a medium of choice raises several questions. First, why has paper withstood the test of time? Or, from another perspective, why have computers failed to replace paper? Second, assuming we will continue to work in environments where both computers and paper (and other physical objects) play critical roles in our everyday interactions, how can we best support human interaction with information existing in both domains? What are the problems surrounding the reality of working in both virtual and physical domains? Which of those problems have been solved, and which remain unsolved?

This paper will focus on the second set of questions. We will start by attempting to define the concepts fundamental to interaction with information in modern workplaces. Central to this will be the division (sometimes fuzzy) between the virtual and physical domains. We will continue with a review of the literature dealing with interaction with information in both the physical and virtual domains, focusing on the movement of information between these two domains. Following this, we will identify some of the key open questions in the field. And, finally, we will posit what we consider to be likely avenues of future research.

Domains of Information Interaction
Throughout the paper we will be dealing with the concept of “information.” This is a nebulous and potentially far-reaching concept. Information can exist in a human generated and easily comprehensible form, such as words on a piece of paper, or it can be inherent in our surroundings, such as the growth rings on a tree indicating age. Information can also be unknown, such as the existence or lack of alien life. While we will touch on different forms of information, we are mostly interested in dealing with information which is human generated, is used in every day interactions, and is explicitly stored in either virtual or physical form.

In order to characterize transitions of information between the virtual and physical domains we must be able to distinguish between the two. This can be quite challenging at times, especially with recent work on augmented reality, where information may reside simultaneously in physical and virtual forms. The use of haptic computer input devices, where manipulation of digital data is mediated by a physical interface, can also blur boundaries. We will nevertheless attempt to set out definitions which allow us to proceed with our review of the field.

The Virtual Domain
Information in the virtual domain is that which exists as bits in a computing system. It is generally made visible to users via some sort of projection device, and made editable via an input device such as a mouse or keyboard.

Some of the tell-tale properties of information in the virtual domain are as follows:

1. Exists as data stored in computer memory and/or on computer disk.

2. Is made visible to a user only through some intermediary, such as a computer monitor or projector.
3. Can be modified by a user only through some intermediary, such as a keyboard or a mouse.
4. A single instance may be presented in a variety of ways (representation vs. presentation)
5. A single instance may be presented independently in a number of locations.
6. A single instance may be replicated with ease.

These properties foster certain interaction patterns on the part of the users. The ability to present information in different ways makes it easy for users to explore information from different perspectives. For example, a user could view some gathered statistical data as a bar graph, or a scatterplot. As another example, a user of a mapping application could examine elevation as contour lines in 2D, or as an extruded 3D image. The dynamic nature of virtual data also leads to other powerful possibilities, including the ability to clone data, and transport it through a network over large distances, essentially instantly.

The drawback of information stored in the virtual domain relates mostly to how users interact with it. First, users are limited by the input devices made available, usually a keyboard and a mouse. These input devices were designed to be general, they can be used in any number of contexts, however that generality is also a drawback. In any specific situation where interaction with specific data would benefit from specialized interactions, these interactions must be mediated through the same general purpose input devices. The second major drawback of interacting with virtual information relates to how this information is output. Most users rely on computer monitors which have a diagonal size on the order of 20 inches, far smaller than the work surfaces we are accustomed to using, such as tables, and walls. This problem, known as the screen real estate problem, leads to many difficulties in information interaction in the virtual domain.

The Physical Domain

Information in the physical domain is generally that which has some real-world form. This information may exist as abstract notation (e.g. words on paper), or as explicit representation (e.g. a scale model), but regardless, it has some tangible form which may be interacted with directly by a user. Most of the examples in this paper dealing with the physical domain will involve paper specifically, as that is the physical medium most often used for storing and manipulating information, but we will not limit ourselves entirely to paper.

Some of the tell-tale properties of information in the physical domain are as follows:

1. Exists in some tangible form.
2. Each instance is unique and independent.
3. Can be modified directly by a user, using any tool of choice.

The physical nature of the information leads directly to many benefits. People are able to employ their spatial awareness and motor skills, developed over millenia of evolution, in order to guide their interactions. This lends physical information artifacts a particular advantage over computers in many respects. As evidence to support this, it has been found that it is often easier for people to read documents printed on paper, rather than on a computer monitor [24]. Another advantage to information in physical form is that a person can improvise by adapting any aspect of their environment to be a tool operating on that information. For example, a person is guaranteed to be able to use a pair of scissors on a piece of paper, or place a weight on the paper to hold it still. There is no need to explicitly develop support for the interaction between objects. This is in contrast to computing environments, where all interactions between virtual entities (e.g. applications, files, protocols) must be anticipated and explicitly programmed.

It should be noted that not all forms of information in the physical domain are easy to interact with. For example, words chiseled into a large stone don't support either easy modification or transportation. Of course, in this situation, the person who carved the stone likely meant it to be immobile and persistent. Humans have developed ways of storing information in the physical domain that match well with the intended role that information plays. Words on paper are easily transported, edited, and disposed of. Words carved into stone are none of those things.

USE PATTERNS

When looking at information mobility between the virtual and physical domains, it is necessary to consider more than just the capabilities of relevant technologies. What is more important, and what should be the primary motivation behind development of supporting technologies, is the usage patterns of the information. Users of information have particular needs when interacting with information, and it is those needs which trigger the movement of information between domains. in this section we will consider first the environments in which people generally use information, followed by the additional consideration multiple people collaborating, and finally we will discuss several specific examples of information use in different industries.

Environment

Considering the office workplace, we must consider the entirety of the environment, as the flow of information among people and between domains is shaped by that environment. There has been extensive work examining office work environments, initially in the fields of sociology and ethnography [1], and more recently in the context of man-machine interfaces [6, 35, 39]. All of these works reach a common conclusion that our environment very fundamentally shapes our actions. First, the space in which we work governs how we can organize information in that space. For example, a desk allows us to lay out papers in stacks, whereas a cork board allows us to pin paper, but not stack it. Second, the structure of our environment defines how we interpret ownership of that space. A desk is usually interpreted as being a very personal things, and as a result we expect our organization of information to not be interfered with by others. In contrast, a whiteboard is often considered to be a communal resource, resulting in protocols being developed for sharing space, preserving old content (e.g. PLO for please leave on), and work interruption.

Perhaps the most personal and important aspect of a person's work environment is the desk. Researchers have examined how desks are organized, and how the desks themselves serve as repositories of information in the physical domain [19]. After having observed workers in an office, Malone classified desk users into two categories, pilers and filers. Filers keep their documents well organized, while pilers simply stack things up, and often end up digging through their piles looking for something they've lost. These two fundamentally different approaches to information management should be considered to be of significance, regardless of whether one is considering desk organization, organization of offices in generally, or indeed organization of information in the virtual domain.

Now that computers are ubiquitous, the desk of almost every office worker has been augmented in a sense with another work surface, the computer monitor. This monitor is meant to work in conjunction with the desk, however the two are disjoint and isolated from one another. Also, while the physical interactions involving the desks can take advantage of the large surface area of the desk, interaction in the virtual domain is often limited to the very small areas of our computer monitors. Monitor size has been increasing quite slowly over the last several decades, but there is evidence that we may be on the cusp of a major advancement, where any individual person will have a virtual workspace as large as their physical workspace. Researchers at Microsoft are investigating the impact such a change will have on our work patterns and our interactions with computers, and how those interactions will need to adapt [26]. The general consensus, ironically, is that the standard GUI desktop metaphor will have to be modified significantly when our digital work surfaces become the size of desks, but there is no clear picture of what form the new interactions will take.

Figure 1. A multiprojector display system as shown in [26]

Collaboration

It is also critical to recognize the role of information in collaborative tasks [8]. Specialized workspaces such as large tables and white boards, which are well suited to collaboration, are often used. Researchers have investigated the behaviour of groups collaborating towards a shared goal [38]. It was found that these groups often prefer to use material tools such as pen and paper, even when more "sophisticated" tools such as computers are available. It was hypothesized that this is due to several advantageous factors in favour of the material tools, including size, public location, and physical affordances. The unfortunate fact is that computers are simply not well designed for collaboration, especially co-located collaboration. Typical computers have only one keyboard and mouse, precluding the possibility of multi-user simultaneous interaction. Furthermore, computer monitors are ill-suited to being the center of attention in collaborative efforts. It is difficult for two or more people to gather around a computer monitor.

Towards the goal of developing an understanding of what makes work surfaces such as tabletops and white boards effective, these workspaces have been researched. Researchers have investigated the unique characteristics of tabletops, identifying key factors influencing collaboration, including seating, engagement, orientation, and personal space [15]. Further work into the space aspects of working around a tabletop found that collaborators working around a tabletop tend to operate with self-enforced territoriality rules, clustering information artifacts based on the notion of ownership. Artifacts that were shared tended to be placed in a central location, whereas artifacts considered personal were kept close to the perceived owner [31]. The situation is different with whiteboards, however, as they do not lend themselves easily for organizing paper and other information artifacts. Rather, whiteboards are used more often in conjunction with pens as persistent display areas [7]. The capabilities identified in these papers all benefit collaboration in one way or another, but none of them exist in standard computing systems. it is necessary to extend computing interfaces to embrace these capabilities.

Task Specific Information Use

in order to develop a understanding of how people use information, it is useful to examine the information needs of certain specialist information workers. Special domains put particular demands on users, which results in those users displaying very particular patterns of information usage. We will see that there are multiple areas in which specialist information users require information in both virtual and physical domains, and would benefit from a streamlining of information flow between those domains.

One example of an environment where information management is of the utmost importance is the hopspital. Researchers have examined the information needs of medical doctors working in a traditional non-computerized office, and discovered that only 30 percent of doctors' information needs were being met in a timely manner [5]. This is due to the difficulty in filing and accessing large amounts of paper-based information, and is one of the reasons for a push to modernize and computerize the field of medicine. An example of the successful introduction of a digital information management tool into a medical practice was recently discussed [2]. The system described augmented a paper-based scheduling approach with an computerized awareness system employing large wall-mounted displays, cameras, and personal tracking technology. This provided staff resonsible for scheduling with an awareness of doctor locations and surgery progression. This tool successfully improved the logistical efficiency of the hospital, and was retained past the expected trial period due to the reluctance of the hospital to relinquish the system. In contrast, a more negative result was obtained on the introduction of another computer system with the intended goal of improving the efficiency of nursing staff switch-over during shift changes [21]. What was found was that while the system apparently eliminated the need for an oral briefing of the arriving nurses by the leaving nurses, in fact the formal oral briefing was merely replaced by an informal briefing. The computer system, while serving well as a repository of information, failed to replace the need for a face-to-face communication of important details. The lesson drawn from these two examples is that replacing paper-based information management systems with computerized methods can provide benefits, but only when the needs and behaviours of participants are fully understood. Without a deep understanding of how information is used, it is likely that there will be unexpected and perhaps negative results to the introduction of a computerized system.

Another domain with very interesting information use patterns is civil engineering and construction. In this domain there are many interested parties with different skills and backgrounds who must collaborate in a complex and interdependent manner to achieve a shared goal. There have been many efforts to improve the construction process through the use of computers, but these have largely failed. Researchers have observed that this is largely due to the nature of the task [**?**]. The interdependence of the multiple roles in the construction process makes it difficult to introduce a system in a piece meal manner. This would result in an information barrier between the traditional paper-based approaches, and the new computerized approaches. The solution would appear to be to replace the entire process in one fell swoop, but it is understandable that the industry would be resistant to this, due to the importance of maintaining schedules, and a minimal tolerance for error. The conclusion is that due to the nature of information flow in the construction industry that making inroads with new technologies will be difficult.

As a last task-specific area of consideration, we will use airline pilots. Modern airplanes are some of the most computerized work environments, yet paper remains central to the task of flying a plane. Airline pilots use paper for planning flights, coordinating actions during flights, and updating information as flights progress [23]. Yet, despite the use of both paper and computerized information, there appears to be little need for the migration of information between physical and virtual domains. The nature of the virtual and physical information is different, with the physically based information being persistent (e.g. flight plan) , and the virtually managed information being constantly varying (e.g. altitude measurements). It seems that the airline industry has managed to identify and leverage the benefits of both the virtual and physical domains.

Figure 2. Airpline pilots using both paper-based and virtual information, as shown in [23]

SUPPORTING TECHNOLOGIES

A number of technologies have been developed which in some manner help support the transfer of information between the virtual and physical domains. Here we provide an overview of these technologies, highlighting their strengths and weaknesses.

Moving From Virtual to Physical

The primary technology for moving information from the virtual domain to the physical domain is the printer. The printer has evolved over several decades, but its basic role remains unchanged. It takes digital data and outputs it on sheets of paper. There are specialized variations on the printer theme, for example plotters are used to output high-quality diagrammatic information on very large sheets of paper. Until recently, however, printers were limited to 2D output. There is a major shortcoming with existing printing technology, in that it is generally an awkward process to generate a printout. In office environments especially, printers are often fairly distant from a user's computer. This requires the user to issue the print command, walk to the printer, evaluate the state of the printout, collect the printout (if it has completed successfully), and return to their computer. This is a significant interruption in a user's natural workflow.

Some recent work has stretched the capabilities of printers into the third dimension. The field is known as "rapid prototyping," [4] not to be confused with the identically named technique for developing computer software interfaces. While this technology is significant for bringing printing into the third dimension, it is hampered in it's flexibility by the fact that special techniques are required for outputting different materials. As an example, specific approaches had to be developed for printing both metal objects [25] and ceramic objects [11]. These techniques are also somewhat limited in that production of the 3D object is quite time consuming. In general, rapid prototyping has been used in situations where turnaround time for production of the printed artifact could acceptably be hours.

We see that that while there is support for moving information from the virtual domain to the physical domain, that the process required to do so acts as a significant barrier to the user. It is not possible with existing technology for a user to move seamlessly from the virtual to the physical.

Moving From Physical to Virtual

There are several options for moving information from the physical world to the virtual world. The most basic example is the simple human-centered act of reading from a piece of paper and transcribing that information on a keyboard. This process is of course time consuming and error-prone, which suggests a need for more advanced techniques.

Digital flatbed scanners provide the next level of sophistication in converting physical information to digital. Such a scanner accepts 2D data, usually in the form of paper, and converts this information into a digital bitmap representation. While the bitmap format is appropriate for images, it is not ideal for text, vector, or other kinds of data. It is desirable for this data to be in its “natural” form such that it can continue to be manipulated, searched, and otherwise interacted with in the virtual world. Luckily, algorithms have been developed which are capable of extracting text from images [13]. More general feature extraction algorithms have been developed as well, but an entirely robust approach for extracting 2D objects from bitmaps remains elusive [9, 36, 28].

The problem of moving 3D information from the physical to the virtual has led to other technological advances. A 2D representation of a 3D scene can be captured digitally using either a still digital camera or a digital video camera. Simply producing a bitmap representation of the scene may be enough for some purposes, but for more advanced tasks it is necessary to interpret the scene, and extract information regarding the geometry of its contents. This problem is known as “image based rendering,” and is a large field with several distinct sub-areas [33]. The problem becomes even more complex if one desires to capture the geometry of the 3D scene in real time, as discussed by Rusinkiewicz et al [27].

A slightly simpler case of general capture of 3D scenes in real time is dealt with by the technique known as motion capture. Motion capture employs controlled environments with objects of interest tagged with special markers which are tracked by sensors [20]. Tracking the positions of the sensors in real time allows the system to draw conclusions regarding the positions of the objects attached to the sensors. Motion capture and related technologies are only marginally relevant to our work, and will not be elaborated upon further.

It is evident that the problem of moving information from the physical world to the virtual is a difficult one. There are a variety of solutions with different approaches for different scenarios. In the situations we’re most concerned with, the use of everyday information most often stored on paper, the 2D scanning related technologies are obviously the most relevant. Focussing on these, it is apparent that we run into the same limitations as those we identified with printing. it is a significant distraction for a user to scan in documents, and it is not possible for a user to do so directly in the work context.

Mixed Domain Technologies

Many researchers have recognized that there are severe limitations to the traditional supporting technologies such as printers and scanners. These researchers have explored technological solutions which ease the transfer of information between the physical and virtual domains. In several cases, these technologies have taken the approach of blurring the distinction between the physical and virtual domains. If a technology is able to do this successfully then a user can leverage the strengths of both physical and virtual interactions without being concerned with explicitly moving between the two.

An early example of blurring the boundary between the physical and virtual is the XLibris system, developed at Xerox [29]. The designers of this system took the approach of implementing a fully computerized interface that reproduced many of the beneficial affordances of paper. Aspects of physical interaction that they preserved include the appropriate paper sheet aspect ratio, the ability to interact with a pen-like stylus, and the ability to make natural annotations. This system could be considered an early prototype of the now popular tablet PC platform. The approach is powerful in many ways, but has drawbacks. The fact remains that a tablet PC device remains large and bulky compared to sheets of paper, and does not share many of the capabilities that we take for granted with paper, including easy portability, and compatibility with tools such as paperclips, scissors, and standard ink pens.

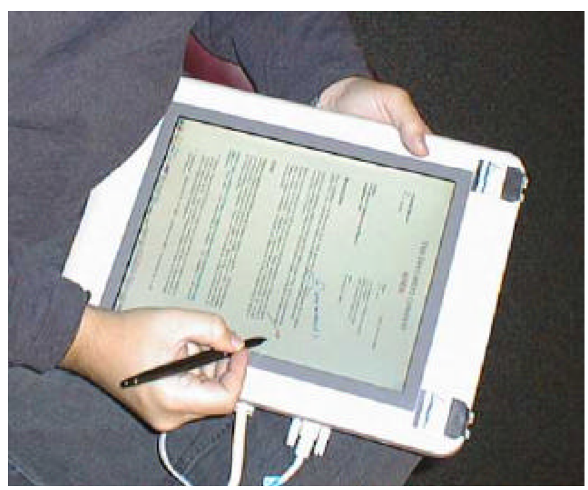

Figure 3. A document editing interface with paper-like affordances, as shown in [29]

An approach which diverges from that of the XLibris system is to enhance the capabilities of actual paper, rather than making a computer system behave more like physical information artifacts. Such an approach was pioneered with early work done at Xerox PARC which explored the use of paper as a computer interface [14]. The researchers recognized that paper has certain advantages over computers, including ease of use, portability, and cost. Their system, called XAX, allows users to write information on specially formatted machine readable printed forms, which are then faxed or scanned into the system. The system employs image processing algorithms to interpret the bitmaps of the inputted forms, and performs appropriate operations. The operations performed are defined by a user in an action description language. While the authors state that the system was used in a production environment for some time, it is clear that there are significant limitations to this approach. Most notably, the user is limited to highly structured input on the pre-printed forms. Natural annotations are not supported. Furthermore, there is no support for immediate feedback as the paper form is filled out.

As mentioned, one of the major restrictions of the XAX system is the requirement that forms be specially defined by a user, with appropriate action defined in an action description language. This limits the use of paper as an interface to preplanned regulated scenarios. This limitation was addressed by a similar but more advanced approach, known as Paper Augmented Digital Documents (PADDs) [10]. The PADDs approach allows users to use natural pen-based annotations on paper printouts, which are then scanned into a computer and interpreted into digital form. The PADDs approach is made possible through the use of an Anoto infrared digital pattern superimposed on all printed documents. The pattern at any point on any page is unique to that location. A special pen, similar to the MEMO-PEN [22], is used for annotations, and scans the digital pattern as it writes. The sequence of pen strokes is stored in the pen, which is later downloaded to a computer. The computer uses the collection of strokes and the identification of pages edited to merge the strokes into the digital versions of the documents.

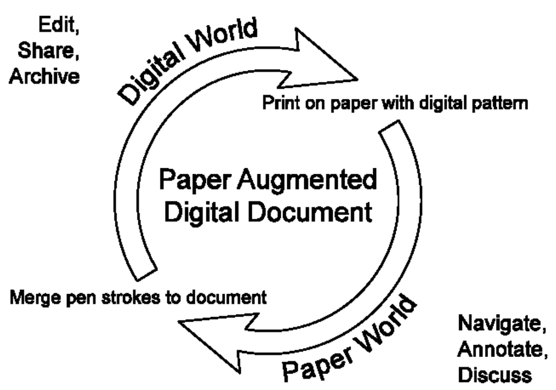

Figure 4. Cycle of information transitioning between domains, as shown in [10]

Further work on a similar paper-based interface has been implemented by Liao et al [16]. The authors describe a system dubbed PapierCraft, based on the PADD architecture, which provides a richer set of editing tools to the user of the paper interface. This implementation doesn't appear to take a significantly different approach to what was described by Guimbretiere, but it seems that the system is more mature, robust, and feature-rich.

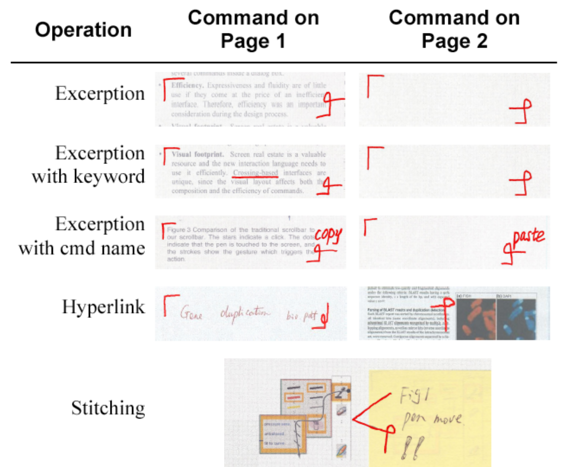

Figure 5. Pen-based interactions with paper interface, as shown in [16]

A third system that employs paper interfaces activated by smart pens was described by Signer et al [34]. This system again employs the Anoto pattern enabled paper and pen for registering pen actions with locations. In this case, however, the researchers developed the system for a specific real world scenario, namely tourist brochures and maps for a festival. In such a scenario the users are not dealing with an editing task, where annotations are recorded and fed into a computer system. Instead, the users use the pen to request real-time up-to-date festival information by pointing at related areas of the brochure or map. For example, the user touches the pen to an event in the festival brochure in order to discover if there are tickets left. This feedback obviously can't be provided via the paper, as it is static. Instead, the designers provided a bluetooth headset, through which audio generated by a text-to-speech application is fed. This provides a fully interactive application, something which had not previously been realized. The main limitation to this approach is, of course, that feedback is only available through sound.

More traditional augmented reality approaches have also been used for mixing information in the virtual and physical domains. The Ariel project has looked at the problem of architects tracking the progress of construction projects [18]. It was observed that many small changes are made to engineering drawings as a project is built. Architects spend a great deal of time updating their paper drawings to properly reflect the reality of a construction project, as compared to the original plan, but these changes are often not transferred to the central digital versions. The effort to do so is considered to be too large. The Ariel project enhanced the engineering drawings of the architects by superimposing a projected virtual interface to the online engineering drawings. This was intended to ease the process of updating the online drawings.

Unfortunately the paper does not discuss the use of the actual system in an engineering environment. We suspect that such a system would run into problems in practice, where it would be difficult to set up the necessary projectors and computers in a fluid work environment such as a construction site.

More recent work has taken the augmented reality approach to more sophisticated levels. Researchers at MIT developed a table-top architectural simulation system which was designed to leverage the computational abilities of computers as well as the advantages of tangible interaction [3]. The system allows users, in particular architects and planners, to place models representing buildings on a top-projected work surface. The user also places objects who's orientations and positions represent variables such as sun position, wind direction, and wind intensity. A camera recognizes the locations of all objects, and a computer runs simulations calculating shadow locations, reflections from windows, and wind flow around and between buildings. The simulation results are then projected back onto the work surface. Users are free to move models around, with the simulation adjusting in real time. The paper describes a truly elegant information interface, yet the system discussed is somewhat limited, being only able to handle a limited number of different objects at a high level of abstraction. It is unlikely that the system could easily be generalized to situations where a complex construction must be planned at a high level of detail.

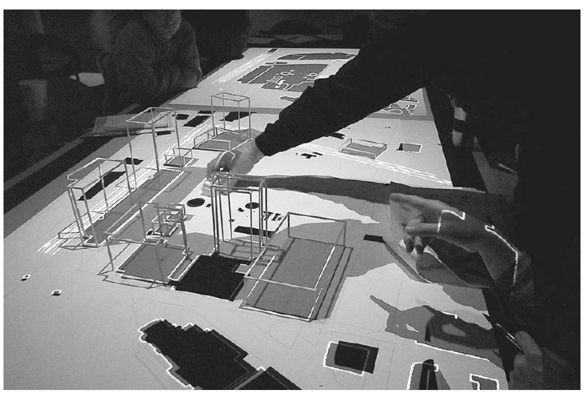

Figure 6. A mixed virtual and physical urban planning simulator, as shown in [3]

We have discussed several approaches which blur the boundaries between the virtual and physical domains. These approaches fall into three distinct categories. First, there are computing systems which are augmented with desirable characteristics of physical objects. This can improve interaction with these computers, but as we saw they are still limited by the requirements that computing hardware be fit into the device, making them bulky and inflexible. The second approach is to augment actual paper with computing abilities. This increases the power of working with paper, but as we saw there are hurdles when it comes to synchronization with information in the virtual domain. Paper simply isn't capable of providing real-time feedback. Finally, we saw a true mixed domain system. This was perhaps the most elegant solution of all, but is hampered by being so dependent on fixed projectors, cameras, and computing resources.

UNSOLVED PROBLEMS

After considering the motivation for supporting information mobility between the virtual and physical domains, and examining the current state of technology supporting such mobility, we have identified some unsolved problems which need to be addressed in order for the field of study to advance.

Lightweight Movement Between Physical and Virtual

The status quo for moving information between the physical and virtual domains has hardly changed in the last several decades. Printers and scanners have advanced in terms of resolution, colour accuracy, and speed, but they have not changed significantly in terms of how they fit into users' workflow. It is still necessary at minimum to remove one's attention from the workspace, operate the dedicated scanning or printing device, and fiddle with the physical (paper) artifacts.

Some efforts have been made to address this problem through the use of mixed domain augmented reality systems, but they fall short in terms of being a realistic solution. The systems developed so far are too dependent on specialized hardware to be practical in everyday environments, and are very application specific. There is no general lightweight solution for moving information between the virtual and physical domains.

Real-Time Physical-Virtual Synchronization

All of the approaches discussed that involve information which is stored both on a computer and in a physical form share one very significant disadvantage. This disadvantage is that while the information exists in the physical domain, there is no connection maintained between that information and the information in the virtual domain. The information in the physical domain is effectively isolated from the virtual domain. This can lead to severe problems. For example, it may be desirable to maintain synchronization of all copies of that information. This is easily done if the work is performed solely in the virtual domain. Modern collaboration tools allow for shared editing of documents, spreadsheets, and any number of other information artifacts. The isolation of information in the physical domain makes it impossible to perform any kind of real-time synchronization between copies and results in conflicts between different copies should any edits be performed on any copy.

RESEARCH OPPORTUNITIES

In considering the unsolved problems as identified in the previous suggestion, we have identified a few likely candidate ideas for advancing the current state of research in the field.

Information Mobility in the Workspace Context

We have discussed how the effort required to explicitly move information between the virtual and physical domain is significant. There are possibilities for future work that center on reducing this overhead.

The first possible solution to this problem depends on the observation that existing technologies (printers and scanners) require that users shift outside of their natural workspace in order to initiate the transfer of information. For example, moving information from a piece of paper into a computer requires that a user move from the table to a scanner. Similarly, moving information from the computer to paper requires that a user shift focus from the computer screen to a printer. Nobody centers their work efforts around printers or scanner beds. We suggest that a significant amount of improvement could be realized if all information transfers were initiated and realized in the context of the natural work environs. We offer two specific examples of how this could be realized, on for movement from the physical to the virtual, and one from the virtual to the physical.

In Situ Mobility From Physical to Virtual

It is apparent that the typical physical work area for viewing and interacting with virtual information is the computer monitor. This is usually a desk mounted unit, but with increasing frequency it is a large wall-mounted or table-top system. In order to provide in situ mobility from the physical to virtual domain we must use the monitor itself as a mechanism for inputting information. This is a departure from the monitor's role as purely an output device, but it could be considered a natural progression to collapse the roles of monitor and scanner as complementary output and input devices into a single device which fulfills both roles.

The benefit derived from using the monitor as an input device is derived from the fact that is can allow a user to input information directly into the actual work context. In fact, we suggest that there should be a direct 1:1 correspondence between the location of the object being input, and the location to which the virtual copy is placed in the virtual domain. In simpler terms, we suggest that a user should be able to simply hold up a document or an object to the screen, upon which a virtual copy of that object is created in placed in the virtual domain in a location corresponding to where the physical object was held.

The interaction as defined leads to two main benefits over the status quo of using a scanner. First, several steps are removed from the process. The user must simply place an information artifact against the screen, and the information capture occurs automatically. Second, the information is automatically placed in the intended location in the virtual domain.

It is obviously necessary to develop new display technologies to support the interactions described. The monitor must be able to capture information placed on it's surface. A webcam, such as the cameras included in the new Apple laptops is not adequate, as such a device can't capture an object placed directly on the display surface. We require that the display surface itself be able to capture as well as display information. There are a few possible implementations of this. The first option requires that the display surface be built from a semi-transparent material. Information is projected on it from within, but a camera located inside the monitor also captures the image of anything placed against the front of the monitor. The monitor switches very quickly back and forth between the role of capturing and displaying information. Interleaving these operations would result in a seemingly continuous display of information, obviating the possibility of reduced image quality. A second approach involves the placement of many camera input "pixels" embedded between the display output pixels on the display surface. Each input pixel would be capable of capturing visible information in a very small field of view. Combining information from all input "pixels," the system would be able to produce an image of the entire scene in front of the display surface. This approach is the subject of a patent filed by Apple Computer [37], but it is unknown of any actual working system has been developed.

In order to clarify how our solution as presented might work, it is helpful to consider a few illustrative scenarios.

Scenario 1: Inputting a text document. First consider the scenario of a user having a piece of paper with text on it, which needs to be input into the computer. The user takes the piece of paper and places it flat against the monitor, with the text facing the monitor. The monitor is able to recognize that an object has been placed against it, and captures an image of what that object is. The computer then processes the image and recognizes that the image contains a page of text. The system can then create a new editable document which is positioned in the workspace at the location that the real document was placed.

Scenario 2: Inputting artwork. A user of a graphics package may want to import the image of a real-world object. Take as an example an apple. With our system the user would hold the apple up to the monitor, oriented in such a fashion that the side of the apple facing the monitor is the desired view. The monitor again recognizes the object, and captures an image. If the user holds the apple against a document in a paint program, the image of the apple is inserted into the document in the position at which it was held. The user is then free to manipulate the apple image in the art application. An obvious shortcoming of this approach is that the image captured is 2D. This is acceptable for 2D art packages, but is not adequate for capturing 3D models for 3D applications.

Scenario 2: Online purchase. In this scenario we are concerned with a payment involving a credit card and a website. This differs from the previous two scenarios in that the visual format of the physical object is irrelevant. Only the data content is significant, namely the credit card number, expiry date, and owner's name. After the user has chosen the items to purchase, the website shows a page with the list of items, the price, and a blank spot the shape and size of a credit card. The user then places a credit card against the blank spot. The computer processes the image, extracting the pertinent information, and transmits it of the network to complete the purchase. In this scenario the direct input saved the user the effort of manually inputting the credit card information, and also avoided the possibility of transcription errors. It is also useful to note that the action of placing the card against the monitor is useful as a metaphor, as it is similar to the action of swiping a card through a reader.

In Situ Mobility From Virtual to Physical

Allowing for the seamless transition of information from the virtual to the physical is significantly more challenging than in the opposite direction, which we just described. While the capture technology required to move from physical to virtual exists, the major challenge for moving from virtual to physical is the production of an actual physical object. The suggested interaction we will describe here is what we picture as being ideal from a user's standpoint, unfortunately the technology to make it possible does not yet exist, nor is there any indication that it will be available in the near future. Our approach is very much a "what-if" scenario.

A major concern in streamlining the process of transitioning from the virtual to the physical is again operating on information in context. We want to minimize operations that break a user's workflow by removing the user from the work environment. There is a major difficulty, however, in producing physical information artifacts via a computer display. Ideally, the user should be able to pull a physical information artifact directly from the display. For example, a user could view a word document, and peel a copy of the document off the display. This document would then be a physical object which the user could treat as any other piece of paper. This is obviously a fanciful idea, however there is a faint hope that this might be possible at some point in the future. One way that this idea may be realizable would be to have a display surface of two layers, a lower layer for displaying output as with a normal display, and a surface layer which is able to produce physical paper-like documents. There are several requirements that this surface layer would have to fill:

1. Transparent, so that the lower layer is visible to the user.
2. Writable, so that persistent images can be written.
3. Shapeable, so that a physical sheet of a certain size can be produced.
4. Replenishable, so that multiple documents can be generated.

It is difficult to speculate on what kind of surface might fulfill these requirements. The material would have to be normally transparent, as indicated, but also as mentioned it would be necessary for it to be written to before the user peels it off. Furthermore, it would have to be shapeable, meaning an arbitrarily shaped sheet would have to be cut out from the surface as a whole in order for the user to peel it off. The requirement that it be replenisheable is a particularly difficult one to support. A user needs to be able to peel an unlimited number of documents from the surface of the display. One possibility for filling this requirement is that the surface layer be generated by a fluid which flows into place and hardens into the material which ultimately becomes the physical artifact. As a new artifact is peeled from the display, the empty spot is replenished by the fluid, which hardens in place.

Enabling Synchronization between Physical and Virtual Domains

As discussed previously, one of the advantages of working in the virtual domain is the support for management of multiple copies of data. Data can be easily copied and transported over large geographic distances. Different copies can be worked on independently by multiple users. If desired, at some point in time the multiple copies can be synchronized, as is done with version control systems. Another approach is to allow geographically distributed users to collaborate on a single copy of information. This is done with collaborative authoring tools.

It would be hugely desirable for the multiple copy management capabilities of the virtual domain to extend to the physical domain. This is problematic, of course, as physical objects have no computing power, let alone communications capabilities. We believe that the solution is to blur the distinction between virtual and physical beyond what has been done already.

Synchronization between objects in the physical domain and between the physical and virtual domains can be achieved by adopting information artifacts which appear physical to the user, but which have capabilities drawn from the virtual domain. Taking a paper-like artifact as an example, such an artifact would be thin and light like paper, be editable with a pen or pencil like paper, yet would have dynamic content and communications capabilities like a computing device. All the desirable affordances of a physical object would be present, with the additional desirable capabilities of information management in the virtual domain. The perfect technological candidate for such an approach is so called "digital paper," in development by a number of companies.

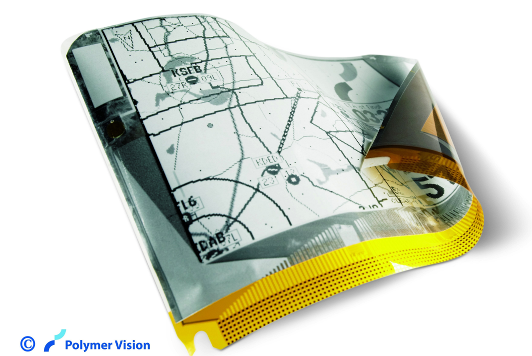

Figure 7. PolymerVision flexible displays, as shown at www.polymervision.com

Aspects of interaction with digital paper have been discussed, notably by Holman et al [12]. In their paper they describe interaction techniques based on gestures such as rubbing, flipping, and pointing, for performing the corresponding actions of copying between documents, scrolling through pages, and clicking on content. What is missing from this discussion is the possibility of multiple users working simultaneously on shared content, each with separate pieces of digital paper. We suggest that Holman's language of digital paper interaction be extended to allow for the coordination of remote collaboration, including simultaneous editing, data locking, and version control. The details of how this would be realized is left for future work.

It will likely be a long time before digital paper matures to the point that it is able to act as a replacement for normal paper. It will likely be even longer before digital paper is able to incorporate the computation and communication capabilities necessary to support proper version control and collaboration. It appears, however, that the research community is at least progressing down the correct path of investigation which will eventually make this a reality.

CONCLUSIONS

We have discussed the issue of information mobility between the virtual and physical domains. We began by attempting to define what we consider to be information, and what qualifies as being either the physical or virtual domains. We then discussed information use patterns, and concluded that the way people use information should be the main driver and motivator behind the development of technologies which support information management and interaction. Following this, a survey of technologies which support information mobility was performed. We identified significant work which fell into three categories: that which enables information flow from the virtual to the physical, that which enables information flow from the physical to the virtual, and that which blurs the boundaries between the two domains. It was found that while work in the third category is the most promising in terms of future growth, it is also limited by practical concerns surrounding deployment. After the discussion of past work, we identified two major areas where there is work remaining to be done. These are supporting lightweight movement of information between the physical and virtual domains, and the real-time synchronization of physical objects with virtual information. Finally, we described two hypothetical interaction models and related technologies which may advance the state of research in these two areas of interest.

REFERENCES

1. I. Altman. *The Environment and Social Behavior*. Brooks/Cole Publishing Company, 1975.
2. Jakob E. Bardram, Thomas R. Hansen, and Mads Soegaard. Awaremedia - a shared interactive display supporting social, temporal, and spatial awareness in surgery. In *Proceedings of CSCW 2006*, pages 109–118, 2006.
3. Eran Ben-Joseph, Hiroshi Ishii, John Underkoffler, Ben Piper, and Luke Yeung. Urban simulation and the luminous planning table: Bridging the gap between the digital and the tangible. *Journal of Planning in Education and Research*, 21:195–202, 2001.
4. Marshall Burns. *Automated Fabrication*. Prentice Hall, 1993.
5. D. G. Covell, G. C. Uman, and P. R. Manning. Information needs in office practice: are they being met? *Annals of Internal Medicine*, 103(4):596–9, 1985.
6. Anind K. Dey and Gregory D. Abowd. Towards a Better Understanding of Context and Context-Awareness. Technical report, GVU Technical Report GIT-GVU-99-22. College of Computing, Georgia Institute of Technology, 1999.
7. P. Dillenbourg and D. Traum. Does a shared screen make a shared solution. In *Proceedings of the Computer Support for Collaborative Learning (CSCL) 1999 Conference*, pages 127–135, 1999.
8. Paul Dourish and Victoria Bellotti. *Awareness and coordination in shared workspaces*. ACM Press New York, NY, USA, 1992.
9. W. E. L. Grimson. *Object recognition by computer: the role of geometric constraints*. MIT Press Cambridge, MA, USA, 1991.
10. Francois Guimbretiere. Paper augmented digital documents. In *Proceedings of UIST '03*, pages 51–60, New York, NY, USA, 2003. ACM Press.
11. J.W. Halloran. Freeform fabrication of ceramics. *British Ceramics Transactions*, 98(6):299–303, 1999.
12. David Holman, Roel Vertegaal, Mark Altosaar, Nikolaus Troje, and Derek Johns. Paper windows: interaction techniques for digital paper. In *CHI '05: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 591–599, New York, NY, USA, 2005. ACM Press.
13. S. Impedovo, L. Ottaviano, and S. Occhinegro. Optical character recognition–a survey. *INT. J. PATTERN RECOG. ARTIF. INTELL.*, 5(1):1–24, 1991.
14. Walter Johnson, Herbert Jellinek, Jr. Leigh Klotz, Ramana Rao, and Stuart K. Card. Bridging the paper and electronic worlds: the paper user interface. In *CHI '93: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 507–512, New York, NY, USA, 1993. ACM Press.
15. Russell Kruger, Sheelagh Carpendale, and Saul Greenberg. Collaborating over physical and electronic tables. In *Extended Abstracts of CSCW '02*, pages 139–140, 2002.
16. Chunyuan Liao, Francois Guimbretiere, and Ken Hinckley. Papiercraft: A command system for interactive paper. In *Proceedings of UIST '05*, pages 241–244. ACM Press, 2005.
17. Ziming Liu and David G. Stork. Is paperless really more? *Communications of the ACM*, 43(11):94–97, 2000.
18. Wendy E. Mackay, D. S. Pagani, L. Faber, B. Inwood, P. Launiainen, L. Brenta, and V. Pouzol. Ariel: Augmenting paper engineering drawings. In *Proceedings of CHI 1995*, pages 421–422, 1995.

19. T.W. Malone. How do people organize their desks?: Implications for the design of office information systems. *ACM Transactions on Information Systems (TOIS)*, 1(1):99–112, 1983.

20. T.B. Moeslund and E. Granum. A survey of computer vision-based human motion capture. *Computer Vision and Image Understanding*, 81(3):231–268, 2001.

21. Glenn Munkvold, Gunnar Ellingsen, and Hege Koksvik. Formalizing work - reallocating redundancy. In *Proceedings of CSCW 2006*, pages 59–68, 2006.

22. Shinji Nabeshima, Shinichirou Yamamoto, Kiyoshi Agusa, and Toshio Taguchi. Memo-pen: A new input device. In *Proceedings of CHI 1995*, pages 256–257, 1995.

23. Saeko Nomura, Edwin Hutchins, and Barbara E. Holder. The uses of paper in commercial airline flight operations. In *Proceedings of CSCW 2006*, pages 249–258, 2006.

24. K. O'hara and A. Sellen. A comparison of reading paper and on-line documents. In *Proceedings of CHI 1997*, pages 335–342, 1997.

25. F. Ribeiro. 3d printing with metals. *Computing & Control Engineering Journal*, 9(1):31–38, 1998.

26. G. Robertson, M. Czerwinski, P. Baudisch, B. Meyers, D. Robbins, G. Smith, and D. Tan. The large-display user experience. *Computer Graphics and Applications, IEEE*, 25(4):44–51, 2005.

27. S. Rusinkiewicz, O. Hall-Holt, and M. Levoy. Real-time 3D model acquisition. *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pages 438–446, 2002.

28. K. Saeed, M. Tabdzki, and M. Adamski. A New Approach for Object-Feature Extract and Recognition. *9th International Conference on Advanced Computer Systems-ACS*, 2:389–397, 2002.

29. B.N. Schilit, G. Golvchinsky, and M.N. Price. Beyond paper: Supporting active reading with free form digital ink annotations. In *Proceedings of CHI 1998*, pages 249–256, 1998.

30. Marcus Schreyer, Timo Hartmann, and Martin Fischer. Supporting project meetings with concurrent interoperability in a construction information workspace. *Information Technology in Construction*, 10:153–167, 2005.

31. Stacey D. Scott, M. Sheelagh T. Carpendale, and Kori M. Inkpen. Territoriality in collaborative tabletop workspaces. In *Proceedings of the 2004 ACM conference on Computer supported cooperative work*, pages 294–303. ACM Press New York, NY, USA, 2004.

32. A.J. Sellen and R.H.R. Harper. *The Myth of the Paperless Office*. MIT Press, 2003.

33. H.Y. Shum and S.B. Kang. A Review of Image-Based Rendering Techniques. *IEEE/SPIE Visual Communications and Image Processing (VCIP)*, 213, 2000.

34. Beat Signer, Moira C. Norrie, Michael Grossniklaus, Rudi Belotti, Corsin Decurtins, and Nadir Weibel. Paper-based mobile access to databases. In *SIGMOD '06: Proceedings of the 2006 ACM SIGMOD international conference on Management of data*, pages 763–765, New York, NY, USA, 2006. ACM Press.

35. Lucy A. Suchman. Office Procedure as Practical Action: Models of Work and System Design. *ACM Transactions on Office Information Systems*, 1(4):320–328, 1983.

36. Q.M. Tieng and WW Boles. Recognition of 2D Object Contours Using the Wavelet Transform Zero-Crossing Representation. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 19(8):910–916, 1997.

37. Michael Uy. Integrated sensing display. Patent, 2006.

38. Steve Whittaker, , and Heinrich Schwarz. Back to the future: Pen and paper technology supports complex group coordination. In *Proceedings of CHI 1995*, pages 495–502. ACM Press, 1995.

39. Steve Whittaker, Erik Geelhoed, and Elizabeth Robinson. Shared workspaces: how do they work and when are they useful? *International Journal of Man-Machine Studies*, 39(5):813–842, 1993.