Paper Augmented Digital Documents

François Guimbretière
Department of Computer Science
Human-Computer Interaction Lab
University of Maryland,
College Park, MD, 20742
(301) 405 7952
francois@cs.umd.edu

ABSTRACT

Paper Augmented Digital Documents (PADDs) are digital documents that can be manipulated either on a computer screen or on paper. PADDs, and the infrastructure supporting them, can be seen as a bridge between the digital and the paper worlds. As digital documents, PADDs are easy to edit, distribute and archive; as paper documents, PADDs are easy to navigate, annotate and well accepted in social settings. The chimeric nature of PADDs make them well suited for many tasks such as proofreading, editing, and annotation of large format document like blueprints.

We are presenting an architecture which supports the seamless manipulation of PADDs using today's technologies and reports on the lessons we learned while implementing the first PADD system.

Keywords: Paper Augmented Digital Document, PADD, Anoto, Paper based user interface, Digital pen.

INTRODUCTION

For several decades, experts have predicted that the advent of more powerful and compact computers will result in the creation of paperless offices. Yet, as pointed out by Sellen et al. in "The Myth of the Paperless Office" [28], the consumption of paper is on the rise, and with few exceptions, office work still relies heavily on paper. Sellen et al. provided a careful analysis of the reason of this state of affairs, pointing out the wide gap between paper affordances, such as ease of navigation and annotation, high information density display, and digital document affordances, such as ease of distribution, archival and search.

Many systems have been proposed as solutions to narrow this gap. Some, such as the DigitalDesk [30], and Ariel [18], proposed bringing digital resources to paper. Others, such as Xax [13], Intelligent Paper [7], Audio Notebook [29] and Anoto [4], used paper as an input medium. Others still, such as the Freestyle system [15] or XLibris [27], explored how paper affordances could be provided on tablet computers, such as the recent Tablet PC.

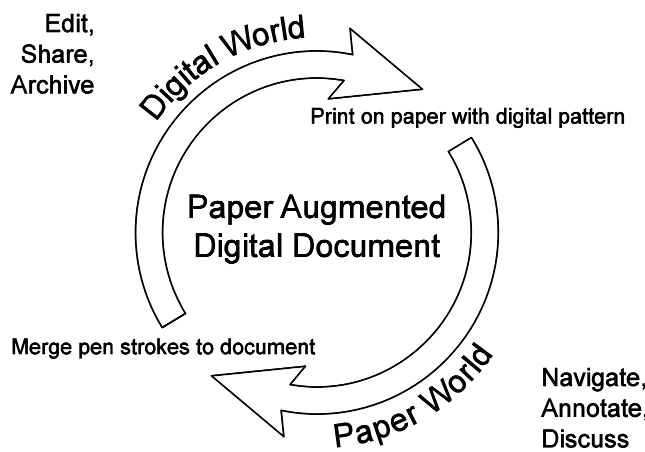

Figure 1: Life cycle of a Paper Augmented Digital Document. PADDs are digital documents, which can be manipulated either in the digital world or in the paper world. They provide affordances of both digital-based and paper-based documents.

The work presented in this paper explores a fourth track, which has been given little attention in the past: *cohabitation*. In this approach, the digital world and the paper world are treated on an equal footage: paper and computers are simply two different ways to interact with *Paper Augmented Digital Documents* (PADDs) during their life cycle (Figure 1). While in the digital realm, PADDs offer all the digital affordances, but require the use of a computer to access them. While in the paper realm, PADDs can only record marks performed on them using a digital pen, but they offer all the affordances of paper because they do not require the use of a nearby computer. At any time, the input gathered on paper can be merged with the original document to be processed as a new editing cycle starts.

We believe that the cohabitation paradigm supported by PADDs will prove very powerful since its basic cycle reflects the transient role of paper in the few successful paperless companies, such as DanTech in [28]. It is well adapted to many activities that currently rely heavily on paper, such as proofreading, editing of drafts, and annotating large format documents, like blueprints.

In this paper, we present an architecture to support PADDs, discuss practical requirements for each elements of the architecture, and identify several important design trade-off for systems bridging the gap between the digital and paper world. We also report on our experience while implementing the first prototype system based on an extension of the Adobe Acrobat application [3] and the io pen, an Anoto [4] compatible pen build by Logitech [16].

MOTIVATIONS

This work was inspired by a salient breakdown in the everyday life of a majority of knowledge workers [28]: while a majority of users prefer to review a document in a printed form, annotations, once made on paper, are time consuming to incorporate back into the original digital document.

This breakdown reflects the tension between the set of affordances provided, on one hand, by paper and, on the other hand, by digital documents. Following the analysis by Sellen [28], we believe that the hegemony of paper as a prime medium for review is rooted in its three fundamental affordances:

- Paper is light, flexible and easy to annotate. These characteristics make it the perfect medium to carry pieces of information and annotate them even in difficult environments, such as a construction site or public transportation where tight space, awkward settings and poor lighting conditions may be expected.
- Paper is easy to navigate using tactile input, making it possible to read and navigate at the same time.
- Paper can provide large, inexpensive, high-resolution display surfaces either by using large sheets, such as the ones typical in engineering and architecture, or by creating a dynamic display of smaller pieces, as when arranging several document on one's desk.

Yet, when the review is over and its outcome is to be processed, the limitations of paper become all too obvious:

- Paper is a static medium that cannot be modified, re-layout, searched or indexed.
- Paper is expensive to duplicate and distribute.
- Paper is expensive to archive.

To address this breakdown, we propose *Paper Augmented Digital Documents,* digital documents which one can manipulate either electronically or on paper. Because of their chimeric nature, PADDs can seamlessly provide the affordances of both paper and digital documents. They are easy to create from digital content, to distribute, and to archive. They are also easy to annotate, navigate, and assemble in large, inexpensive high-resolution displays. Of course, depending on the current form, a very different set of features will be available. Yet, because the PADD infrastructure enables the users to move from one form to the other, one can pick the most appropriate form without the worry of how one's current work will become part of the final document.

Typical uses

While PADDs can have many uses, we are presenting here the two use scenarios which drove our original design.

Proofreading

For our first scenario, we are using the very common task of proofreading, or commenting on a document. After creating or receiving an electronic copy of a PADD, the reviewer will first print the document on paper. Then, she will start her work using a digital pen to annotate the printout. When she is done, she will simply synchronize her pen with her personal computer and see her notes integrated inside the original document. Her annotations are now part of the digital document and flow with the text. As a result, they are always at her locus of attention, relieving her from dividing her attention back and forth between the annotated copy and the computer screen. Furthermore, her word processor may help her by processing the most common proofreading marks, allowing her to focus on more complex editing tasks.

Collecting "as built" design

Our second target task involves annotations performed on large building blueprints used to record the "as built" design. During this operation, contractors in the field will record on printed blueprints the difference between the expected design and the finished project. These annotations are sent back to the office for the creation of new blueprints. Using PADDs for this task, the drawings are printed as usual but annotated with a digital pen. Back at the office, the strokes are merged directly with the original drawings. As in proofreading, this presents the advantage of providing the corrections at the locus of attention without the need to switch back and forth between the screen and the oversized printout. In addition, if the digital pen strokes are combined with a voice recording, such as in the Audio notebook system [29], or video stream, such as in the RECALL system [31], this system can further reduce the ambiguity of the marks, an ongoing problem in this task.

It is important to note that while digital solutions for our first target task are starting to emerge, such as the XLibris system [27] on the software side, and the Tablet PC [22] on the hardware side, the second task presents technological challenges, such as providing a large (A0) high resolution display with sufficient autonomy for a day of work and the reliability to accept the environment of an unfinished

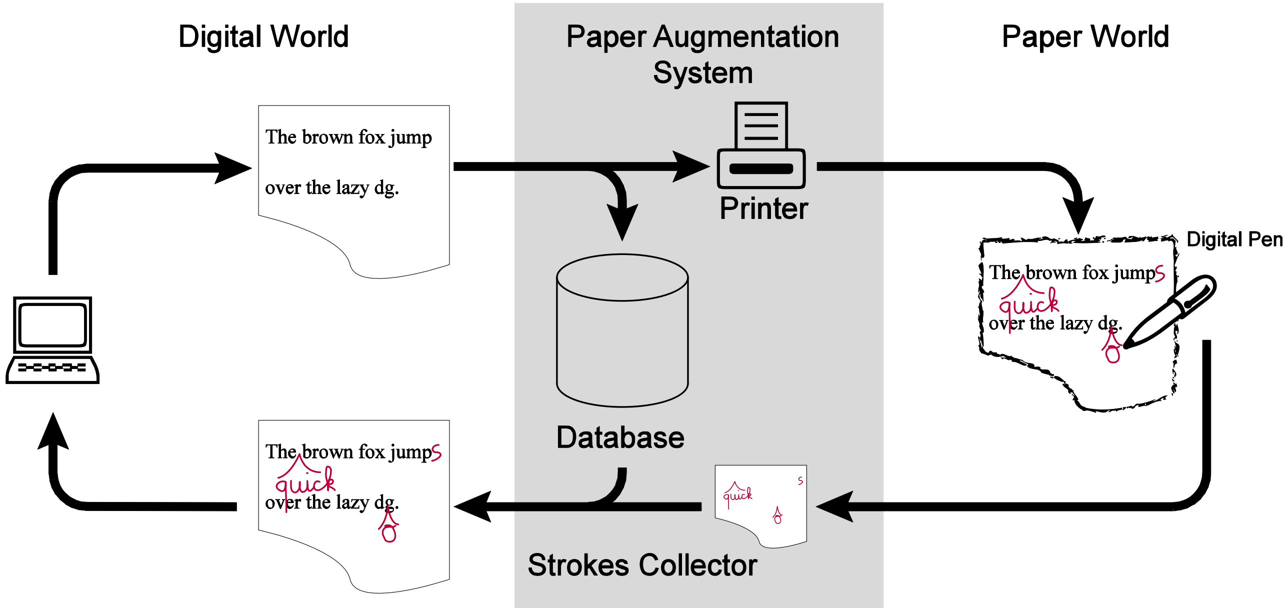

Figure 2: The PADD infrastructure. When paper affordances are needed, a snapshot of the PADD is stored in the database and the PADD is printed. The printer acts like a normal printer but adds a pen-readable pattern to each document. Using a digital pen, the document can now be marked like a normal paper document. The strokes collected by the pen are sent back to the stroke collector which will retrieve the target PADD from the database, and have it process the pen input. The resulting PADD can now be edited, shared, or archived.

building. It is unclear if such a system can be built in the near future, but its functionality can be provided through PADDs. We will come back to this point in the discussion section.

PAPER AUGMENTED DIGITAL DOCUMENTS

PADDs are primarily digital documents: they are stored in a digital format, are edited using computers, and are easy to duplicate, transmit, or archive. They can also be printed on paper when the affordances of paper are needed for any particular task. The different steps in the manipulation of a PADD are shown on Figure 2: PADDs are created as digital documents such as office documents or a CAD drawing. Then, when a paper copy is needed, the document is printed on special paper providing an absolute addressing system. During printing, the system records which page of the document is printed on which sheet of paper, as well as a digital snapshot of the document to be used later when marks will be merged back. The document can now be used as a normal printout, and is easy to navigate, and mark with a digital pen. When users wish to transfer their markings back to the digital version of the PADD, they simply synchronize their pen with a computer. Using information recorded at printing time, the system retrieves the digital document and has it process the markings. At this point, a new cycle could begin: the new documents with the processed marks can be edited, distributed or archived as necessary.

PADD Infrastructure

The PADD infrastructure consists of the four components shown in Figure 2: a documents database, which records the correspondence between digital and physical pages; a printer, which provides printout augmented with a pattern for pen tracking; a digital pen; and a stroke collector which collects strokes from pens, recovers the digital version of the document on which they were drawn, and has it process the marks.

Document database

The role of the document database is to establish and record the correspondence between the digital pages of a PADD and the sheet ID on which they are printed. At printing time, the system first safeguard the *original* copy of the printed document, (not the page description language used to print the document), so that the stroke collectors can retrieve a fully functional PADD. Then, for each page in the document, the system assigns a sheet ID, computes the transformation matrix between document space and paper space (for example, to take into account possible scaling or rotation), and records these pieces of information. It is important to save the transformation matrix on a per document page basis since several document pages can be printed on the same paper page. These pieces of information are saved on a per printout basis, so that the system can handle the common case in which the same document has been printed several times.

The database also provides access control services to manage who can update a PADD and to whom it may be sent upon synchronization. Depending on the application at hand, access control can be set on a document basis, a page per page basis or – since each pen has a unique address – who performed the marks.

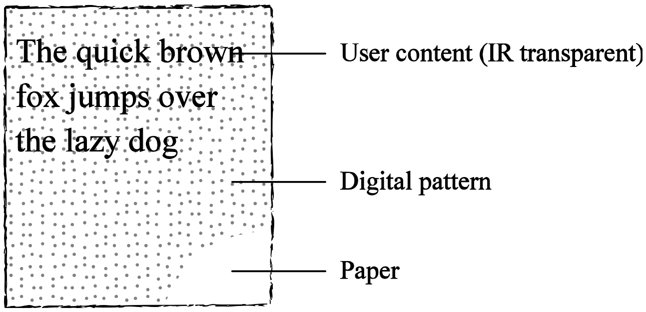

Figure 3: On paper, a PADD is composed by superposing a digital pattern (used by the digital pen to track strokes) and user content. For a PADD using the Anoto pattern (as shown here), user content needs to be printed with an IR transparent ink so that the pen can see the underlying pattern.

Printer

The system is using a printer which includes a page specific pattern in all printing so that marks can be recorded using a digital pen (Figure 3). Our current implementation uses the Anoto pattern [4] for this feature but other system such as DataGlyphs [10], can be used, too.

Depending on the pattern used, the printer can either use sheets that are pre-printed with a given pattern or print the pattern on demand. Furthermore, depending on the digital pattern used, the printer might need to use a special infrared-transparent black ink so printout does not prevent the digital pen from observing the underlying marks.

Digital Pen

The digital pen is used to record strokes made on the printouts produced by the printer. Any system which provides local coordinates, as well as a sheet ID for each stroke, can be used. Depending on the application, the pen could either interact with the PADD infrastructure in real-time by sending the strokes as they are created, or in batch mode, by storing the strokes internally and transmitting them to the infrastructure when a link becomes available.

Strokes Collector

The stroke collector is the intermediary between the digital pen and the PADDs database. Upon being contacted by the pen, it uses the strokes' sheet ID to retrieve the corresponding PADD pages from the database, and has the PADD process the strokes provided by the pen.

Accessing the database is more complicated than it may seem if one wants to address common ways in which people interact with paper. For example, people often regard personal or physical delivery as an important social interaction between the two parties involved [28]. In practice, this means that one could create strokes on a piece of paper which was not printed locally and will need to access a remote database. This problem might be solved by any one of many distributed directory services (such as the internet DNS service [23]). As an example, we show in

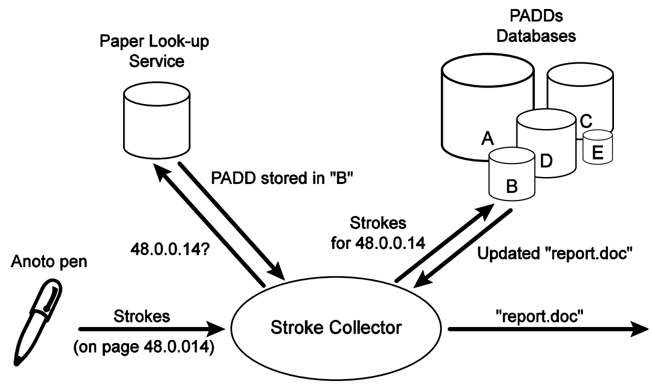

Figure 4: A possible implementation of the stroke collector inspired by Anoto proposed architecture [4]. From left to right: the stroke collector retrieves the strokes from the pen as well as the page ID on which they were created. The stroke collector retrieves the name of the document database managing this page ID using the Paper Look-up Service. It can then contact the database directly to merges the imported strokes with the PADD.

Figure 4 how the system used by Anoto can be extended for our purpose. In this approach, upon receipt of the page address, the stroke collector will contact the Paper Lookup Service and will receive back a Universal Resource Locator for the corresponding documents database. The stroke collector can now collaborate with the document database to process the strokes.

The final step of this process is to have the PADD process the strokes recorded on its paper incarnation. Again, our design was careful to take into account key affordances of paper. For instance, because paper is physical, it is expensive to reproduce and distribute. Providing only a hardcopy to a third party provides a natural, yet limited, protection to the author of the document. Our infrastructure takes this aspect into account by having the PADD, not the stroke collector, process the strokes. This leaves the author of the PADD in control of the distribution of the digital copy of the PADD. Depending of the application and the trust between the author and the editors, the author can either just collect strokes and defer their processing, allow for immediate processing, or provide the author of the strokes a digital version of the PADD with the strokes processed.

Depending on the application, the processing can take a wide variety of forms. Some marks can be just drawn as an overlay layer on top of the original document as in the A-book system [20]; others can be integrated in the text and flow with it; some might be interpreted as commands to process the document in a certain way, using a system such as the one used in the Paper PDA [6] or pidgets in the Anoto system [4]. Some marks might be synchronized with other sources of information, such as in the Recall system [31] and Audio Notebook [29], and some might be processed as they are created to provide immediate

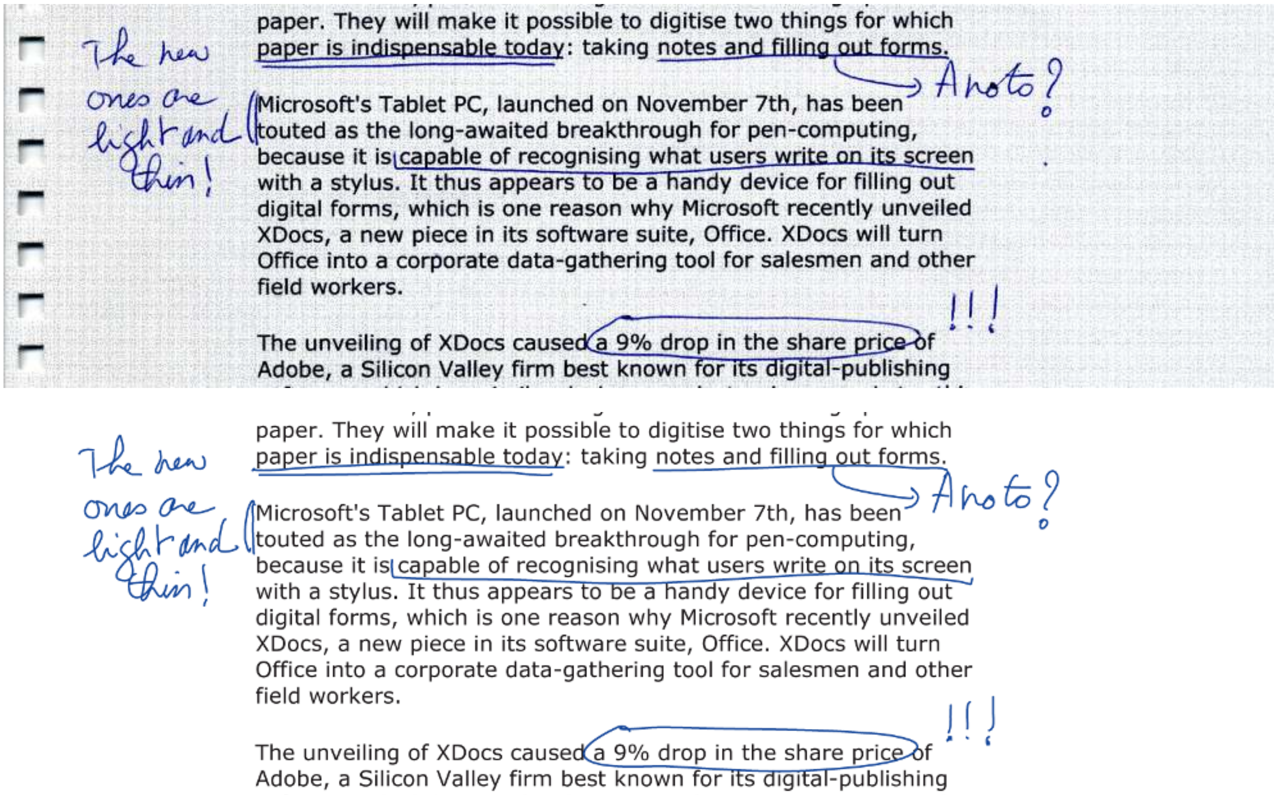

Figure 5: Two representation of the same PADD. Top, we show the scanned version of the paper representation annotated with the io pen. Bottom we show its digital equivalent after synchronization. The PADD was printed on sheet ID 44.0.0.23. Calibration was performed on sheet ID 44.0.0.21. Both sheets are from an A4 notebook provided with the Anoto evaluation kit. Both pictures have been processed as 300 dpi JPEGs. A close-up is shown Figure 6. Original text from [1].

feedback to the users, such as in the Intelligent Paper system [7].

PROTOTYPE

To test the validity of the PADD concept, we implemented a PADD infrastructure prototype. The design goal behind our prototype were to 1) provide a proof of concept of the full life cycle of a PADD; and 2) investigate possible difficulties behind printing and stroke collection since we had concerns about calibration issues. This is not to say that other parts of the design such as the document database or the directory infrastructure used by the stroke collector were not important. However, existing solutions, proposed in different context such as the internet DNS [23] or Anoto's Paper Look-up Service [4], can be tailored for use in a PADD system. We did not focus on this aspect of the solution in this paper.

We implemented our prototype as a plug-in for the Adobe Acrobat system [2]. This choice seemed natural as it allows us to easily compare our system to one of its contender, the Acrobat digital annotation system. Our plug-in plays both the role of the document database and the stroke collector as described below. By itself, the plug-in provides a valid PADD infrastructure with the limitations that it can only be used by a single user, can only deal with PADDs printed locally with one standing printout, and annotations are not interpreted.

Acrobat Paper Augmentation Plugin

Our Acrobat plug-in works in concert with io digital pen and Anoto compatible pen build by Logitech [16]. Users can print Acrobat documents, annotate them using the io Pen, and then merge the annotations into the original document.

Document database

For this prototype, the database is implemented as part of the plug-in. The plug-in keeps track of the current sheet ID available on the printer tray and tracks printing requests so that it can annotate each document page with the sheet ID on which the page was printed. For each page, the plug-in also computes and stores the geometric transformation needed to convert coordinates from paper space to document space. This annotation mechanism makes it easier to manage stroke collection as explained below, and makes the system more reliable since the information for any given page is stored with that page. The annotation protocol provided by the Acrobat format [3] allows the

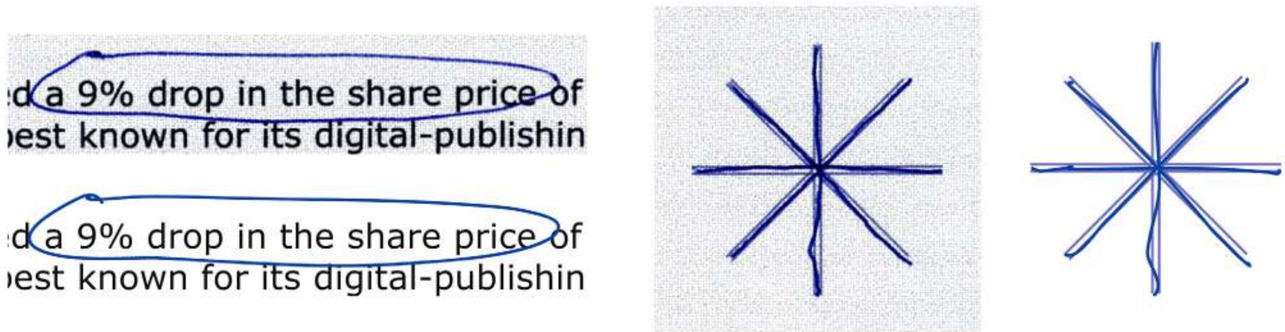

Figure 6: Left, we show a close-up from Figure 5. Note the similar look and feel of the scanned paper version (grey background) and the digital version. The accuracy of the calibration is good enough that a narrow letter like an 'l' could be crossed out with accuracy. Right, we show a calibration test printed on the bottom right corner of sheet ID 44.0.0.23 (and not shown on figure 5). As in Figure 5, both paper and digital versions are shown processed as 300 dpi JPEGs.

guaranteed compatibility of this approach with other plug-ins.

Printer

Potentially the most limiting component in the current prototype is the printer. No commercially available desktop printers can print the Anoto pattern on demand efficiently[1], so our implementation relies on the printer being loaded with pre-printed paper. This causes a reliability problem, since the printer and the database may be out of synchronization. Another hurdle is that common black ink used by printers (both inkjet and laser) is opaque to IR light and masks the underlying page pattern, preventing an accurate pen reading. We avoided this problem by using an inexpensive inkjet printer (HP 5550) and removing the black cartridge. In that configuration, the printer enters a "reserve mode"[2] and uses a mix of Cyan, Magenta and Yellow to produce ink which looks black but is transparent to IR. This solution is preferred to a software solution, such as transforming each color in the document, for it simplicity and robustness. In particular, it is immune to printer firmware modifying ink distribution to save ink.

Stroke Collector

Our plug-in implements a lazy version of the stroke collector by integrating the strokes on user demand. To process the request, the stroke collector first looks up the annotations on each page. Then, it uses this information to retrieve the proper stroke file from the central repository maintained by the io pen system, and transforms the stroke coordinates into document space. Strokes are added to the target page using the standard SDK provided by Adobe [2]. In the current implementation, strokes are not processed and are drawn as one annotation placed as a static overlay on each page.

Results

While limited in scope, our plug-in provided a successful implementation of a PADD system. Using our system, we were able to annotate standard Acrobat documents on paper and then automatically transfer the annotations into the original digital document. We present an example of two versions, digital and paper, of a PADD in Figure 5. On the left is the scanned version of the annotated paper document and on the right is the same PADD shown as an Acrobat document. Figure 6 shows 2 close ups. Not only do the annotations on each document have the same look and feel, but also the calibration errors were small enough to make the system practical in standard proofreading and annotation tasks. While we have not performed formal user testing on our system yet, in part because the difficulty of getting pre-printed paper, the initial reaction of potential users during our presentation at the first Anoto developer conference was very positive.

DISCUSSION

The infrastructure presented in this paper is a unifying framework for several solutions proposed in the past including Xax [13], DigitalDesk [30], Ariel [18], Video Mosaic [17], PaperLink [5], Intelligent paper [7], Paper PDA [11], the EnhancedDesk [14], Anoto [4], and the A-book [20]. This framework identifies the key components involved in building a system bridging the gap between the digital and paper worlds. It also helps us identify several important design trade-off of such systems. We will now present these trade-offs in more details.

Paper as a form filling device

Several systems including Xax [13], Anoto [4], and the Paper PDA [11] have been proposed to collect strokes performed on paper. Among them Anoto is probably the most mature. It is also closely related to the PADD system since it shares the same underlying tracking technology. In its current form, the Anoto infrastructure focuses on form

[1] In our tests, it took a Xerox Phaser 7700DX Color more time to print one page of Anoto pattern than to print this whole paper (or a little more than 1min).

[2] For printer without this capability, one can refill a black ink cartridge with and equal amount of Cyan, Magenta and Yellow. This solution provided a lesser contrast during our tests.

filling applications including transmission of fax (or MMS) through the cell phone infrastructure, transmission of handwritten e-mail (such as the Logitech system), and business applications such as gathering data for healthcare providers. In the Anoto system, each PADD (a "form" in Anoto terminology) is created once and accessed many times by each potential user. The pen can either process strokes as they happen (upon checking the Anoto "Magic Box©" for example) or simply store them (while drawing). The processing of marks does not modify the form, but instead issues the corresponding commands.

Our system extends this "paper as a form filling device" paradigm, to a more general mechanism where paper is used as a general input device for interacting with *user content*. To do so requires a new system architecture, which can automatically manage the correspondence between paper copies and digital copies, provide an easy way to print user content on top of the tracking pattern, and automatically manage the integration of the users' strokes back into the users documents. The infrastructure proposed here as well as our prototype only uses the stroke tracking ability of the io pen and does not rely on the Anoto infrastructure.

Cohabitation versus coupling

Coupling paper interaction and computer interaction is another way to bridge the paper-digital gap. The DigitalDesk [30], Ariel [18], PaperLink [5], Intelligent paper [7], the EnhancedDesk [14] and the A-book [20] are examples of systems exploring coupling. They all require the concurrent use of paper and computers. While the PADD infrastructure as described in this paper could be used to implement these systems, our focus is on cohabitation instead of coupling.

While coupling brings interactivity to paper, a static medium, this improvement comes at the cost of limiting paper affordances. For example, annotating an engineering drawing with Ariel [18] requires the use of a calibrated projection system making it difficult to use the system in the harsh outdoor environment of a building construction site. Because PADDs do not require the use of a nearby computer, they provide all the affordances of paper: they can be used in harsh environments, and are easy to transport and navigate. Our focus is probably best reflected in another of such a system, the Audio Notebook [29] which separates note taking and note processing.

Because our system provides support for cohabitation between paper and digital document, it opens new interaction designs for situations where paper is convenient but poorly integrated with the digital world. For example, PADDs can be used to create a family calendar system which bridges the digital divide between family members, by being accessible not only to tech savvy family members, but also to elderly people who might not be willing to deal with a computer and who prefer paper.

Of course strict cohabitation is not always desirable: for example, by coupling paper interactions and a video edition system, Video Mosaic [17] let video editors rapidly assemble a video sketch mixing video footage and storyboarding elements. While PADDs support strict cohabitation settings, they can also be coupled more tightly with a computer system. An interesting application is their potential use in Air Traffic Control systems. While computer versions of flight strips might be more efficient at managing the traffic, paper versions are often preferred because of their reliability [19]. PADDs can be used in a system where paper strips are used as usual but marks are transmitted to the ACT computer in real-time so that the system can provide advisory information to the controller. By using a somewhat loose coupling between paper and ATC computers, such a system will provide its users with the power and efficiency of a computerized system and the reliability of paper since the system can still be used if the computers go down.

PADD versus tablet computers

With the recent advance in tablet computing, such as the Tablet PC introduced by Microsoft Corporation [22], one may wonder whether the paperless office is finally within reach. This would reduce the need for PADD system. However, we do not foresee an immediate transition.

Ease of annotations

Both advances in hardware, such as better tracking technology, and advances in software, such as ink management SDK [12] and annotation manipulation software such as XLibris [8, 27], are narrowing the quality gap between paper and digital annotation. This is especially true for office workers in a well-defined working environment, who can accommodate the bulk and rigidity of tablet computers. This is not as true for nomadic uses such as editing a paper during a flight, or at a café terrace. In these later cases, we believe that PADDs offer a more flexible ways to interact with a document.

Ease of navigation

While it is true that some research efforts have been made to develop better navigation interfaces (for example the work by Harrison et al. [9]), we believe that paper still has a sharp advantage. Not only does the tangibility of paper make it easier to navigate through a paper document, but also paper is very easy to spread out in one's working environment providing instant navigation. Even putting aside cost issues, such as owning several tablet computers, this kind of interaction would be difficult to emulate using computers. We believe that in this regard, PADDs provide a significant advantage.

Display size

Current printing technology can easily provide large, high-resolution printouts in a variety of formats, from the standard Letter format to A0 format. The resulting display is portable, flexible, resilient, and inexpensive. This contrasts sharply with what can be done with computer screens: it is difficult to provide surface, resolution, and portability in the same package. While new technologies, such as light emitting polymers, will narrow this gap, they are still at the experimental stage for large form factors. We believe that PADDs are the only viable solution to manipulate such documents in the foreseeable future.

Other considerations

There are, of course, many other considerations that make a specific medium useful or not, such as weight, cost, and reliability. While it will be simpler for a lawyer to carry a tablet computer rather than carrying all the documents required by a complex patent litigation case, he might prefer working on paper while riding a crowded subway. An architect might prefer to use paper and a digital pen at the construction site due to the extended battery life but might use a tablet computer while presenting a design to clients.

The main advantage of our proposed infrastructure is that it does not require a radical change in usage patterns. Instead, PADDs let users pick the solution that will work best for the task at hand – which could be paper, a tablet computer or a desktop computer – knowing that they are always interacting with a single, unified PADD.

Infrastructure management

To maintain the correspondence between the digital and the paper version of a PADD, the PADD infrastructure relies on maintaining a large pattern space and a large database of printed documents.

Pattern space

While the Anoto pattern space is very large[3], it may be too small if the PADD system were to be universally applied and each piece of paper printed required a unique ID[4]. This problem is not specific to the Anoto system and will also arise if we had used DataGlyph.

Several solutions can alleviate this problem. First, it is not necessary that all addresses be global. Private address spaces could be repeated as long as the probability of overlap is low. Second, the pen could make more efficient use of the address space. One solution would use a hybrid system combining absolute and relative reading. In such a system, known patterns would be placed further apart on the page as beacons letting the pen filling the gaps using relative tracking such as the MEMO-PEN [25]. This approach would greatly increase the available pattern space at the cost of possibly missing small strokes.

Document database

As described above, the document database has to save a copy of the current version of all printed PADDs. This could be judged as prohibitive: for example, our department sees an average of more than 500MB of data sent to the printer everyday. Several factors alleviate this situation. First, because a PADD may be printed several times, using versioning system, such as the one proposed by Muthitacharoen et al. [24] or leveraging log-based file systems [26] may considerably limit the memory requirement of the document database. Second, the document database stores the original documents, not a printer friendly description of the document. This difference is important because original documents may be more amenable to efficient versioning. For example, changing a single character in a text processor document will be efficiently handled by a versioning system dealing with the original document, but might cause the layout of the document to change leaving little room for versioning between the corresponding printout files.

Calibration

One of the main goals of our prototype was to provide us with insight about calibration errors, the main unknown issue in the implementation of a PADD infrastructure. Calibration errors can come from 3 major sources: the pen might be inaccurate, the printer might place the user content at an inconsistent position on the page, or the relationship between the Anoto pattern and the sheet of paper might vary.

Pen reading

The stroke positions provided by the Logitech io pen were very accurate and provided at a high sampling rate. Our tests showed that data provided by the pen was very reliable even when annotations were performed on top of blocks of text, confirming that simply mixing the three process colors was a simple and efficient way to get a black that is compatible with the pen reading.

Printer repeatability

We found that even entry-level inkjet printers, such as the HP 5550 in our prototype, deliver a very repeatable image placement. The only problem we faced was that the relative location of the printed document with respect to the paper page boundary varied between odd and even pages. This was easily taken care off, since our infrastructure records the paper space to digital space transformation on a page-by-page basis. This confirmed that PADDs do not required the more complex calibration system required by augmented reality systems such as the Digital Desk [30], or Ariel [18].

Our experience confirmed that a printer for this system could either be implemented by having the printer produce the pattern on demand or by using pre-printed paper. The former solution will provide the best results because it guarantees the synchronicity of the printer and the document database, and a perfect alignment between the printed document and the pen pattern since the two will be printed in the same referential. The latter solution will probably be less expensive and its reliability could be

[3] In excess of 2^{48} pages [4].

[4] In 2000, the United States consumed 15.7 M tons of uncoated free-sheet paper [28], or around 2^{42} sheets.

greatly improved if the printer were able to read the pattern on the page to be printed, guaranteeing that the printer and the document database stay synchronous.

Pre-printed paper repeatability

Small variations in the relative position of the pre-printed Anoto pattern, with respect to the sheet of paper, can limit the accuracy of the calibration. Unfortunately, few pre-printed Anoto paper provider are currently available, so we can only present a limited report. In out tests, we found that the paper provided by Anoto was very repeatable from page to page (see Figure 6) but probably quite expensive to produce. We also ran tests with the less expensive paper provided with the Logitech io pen and found that it exhibited small variations (up to ±2pt) making some marks look out of place. From this experience, we believe that having the printer read the pattern at printing time to align the output accordingly will improve the reliability of the system for low cost paper.

FUTURE WORK

The system described above corresponds to an early stage of the implementation of a full PADD infrastructure focusing on personal use. We would like to extend this prototype implementation in two directions. First, we expect to develop a full PADD infrastructure, beginning with our work group and then for a wider audience. Second, we are working on application specific ways to merge marks into the PADD documents. For example, we are working on adapting Golovchinsky et al. work on moving markup [8] so that markups gathered on a paper copy of a word processor document could be merged in the original document as comments flowing with the text. We would also like to explore how PADDs could be used in conjunction with Recall [31], a system designed to transparently capture, share and re-use sketch based design knowledge. Recall is currently a tablet PC based solution but its scope could be increased if paper was used as a primary capture medium.

Another avenue of research will be to use our system as an analytical tool to explore how the use of digital documents and their paper representations differ. While several authors, including Sellen et al. [28] and Marshall [21] have provided us with valuable ethnographic data, it has often been difficult to gather quantitative data to pinpoint what sets digital and paper documents apart. We believe that the PADD infrastructure can be used as a valuable analytical tool by providing a large accurate corpus of data for analysis.

CONCLUSION

In this paper, we presented Paper Augmented Digital Documents (PADDs), digital documents which can be edited both in digital and paper form. We described the PADD infrastructure necessary to interact with PADDs, and presented our early experiences in implementing a simple single-user but fully functional PADD system based on an Acrobat plug-in.

We see PADDs as basic tool toward an office environment where paper is only used as a transient medium chosen when its affordances fit best the task at hand, and then can be discarded when the information collected on paper is merged back with the digital counterpart.

From previous work, practice studies, and discussions with potential users, we believe that our system addresses the need to seamlessly bridge the gap between the digital and the paper world in many tasks such as proofreading, editing, and annotation of large format document like blueprints

ACKNOWLEDGEMENTS

The author wish to thanks Ben Bederson, Samrat Bhattacharjee and Corinna Löckenhoff for their comments on early drafts of this paper. Nino Tarantino and Ewa Reppling, from Anoto, provided technical assistance. David Henry and Vanessa Torres, from Logitech, provided an io pen. Finally the author wishes to thanks reviewers for they useful suggestions.

REFERENCES

1. Saving Trees. The Economist, November 9th 2002, (http://www.economist.com/displaystory.cfm?story_id=1431718).
2. Adobe System Incorporated, *Acrobat Core API Reference*. (http://partners.adobe.com).
3. Adobe System Incorporated, *PDF reference: Adobe portable document format version 1.4*. 3 ed. 2001: Addison-Wesley.
4. Anoto, *Development Guide for Service Enabled by Anoto Functionality*. 2002, Anoto (http://www.anoto.com).
5. Arai, T., D. Aust, and S.E. Hudson. PaperLink: a technique for hyperlinking from real paper to electronic content. *Proceedings of CHI'97*, pp. 327 - 334.
6. Avrahami, D., S.E. Hudson, T.P. Moran, and B.D. Williams. Guided gesture support in the paper PDA. *Proceedings of UIST'01*, pp. 197 - 198.
7. Dymet, M. and M. Copperman. Intelligent Paper. *Proceedings of EP'98*, pp. 392 - 406.
8. Golovchinsky, G. and L. Denoue. Moving markup: repositioning freeform annotations. *Proceedings of UIST'02*, pp. 21 - 30.
9. Harrison, B.L., K.P. Fishkin, A. Gujar, C. Mochon, and R. Want. Squeeze me, hold me, tilt me! An exploration of manipulative user interfaces. *Proceedings of CHI'98*, pp. 17 - 24.
10. Hecht, D.L. Embedded Data Glyph Technology for Hardcopy Digital Documents. *Proceedings of SPIE Color Hard Copy and Graphic Arts III*, pp. 341 - 352.
11. Heiner, J.M., S.E. Hudson, and K. Tanaka. Linking and messaging from real paper in the Paper PDA. *Proceedings of UIST'99*, pp. 179 - 186.

12. Jarett, R. and P. Su, *Building Tablet PC Applications*. 2002: Microsoft Press.
13. Johnson, W., H. Jellinek, J. Leigh Klotz, R. Rao, and S.K. Card. Bridging the paper and electronic worlds: the paper user interface. *Proceedings of CHI'93*, pp. 507 - 512.
14. Koike, H., Y. Sato, Y. Kobayashi, H. Tobita, and M. Kobayashi. Interactive textbook and interactive Venn diagram: natural and intuitive interfaces on augmented desk system. *Proceedings of CHI'00*, pp. 121 - 128.
15. Levine, S.R. and S.F. Ehrlich, The Freestyle System: A Design Perspective, in *Human-Machine Interactive Systems*, A. Klinger, Editor. 1991. pp. 3-21.
16. Logitech, *IO digital pen. http://www.logitech.com*,
17. Mackay, W. and D. Pagani. Video mosaic: laying out time in a physical space. *Proceedings of MM'94*, pp. 165 - 172.
18. Mackay, W.E., D.S. Pagani, L. Faber, B. Inwood, P. Launiainen, L. Brenta, and V. Pouzol. Ariel: augmenting paper engineering drawings. *Proceedings of CHI'95*, pp. 421 - 422.
19. Mackay, W.E., A.-L. Fayard, L. Frobert, and L. Médini. Reinventing the familiar: exploring an augmented reality design space for air traffic control. *Proceedings of CHI'98*, pp. 558 - 565.
20. Mackay, W.E., G. Pothier, C. Letondal, K. Bøegh, and H.E. Sørensen. The missing link: augmenting biology laboratory notebooks. *Proceedings of UIST'02*, pp. 41 - 50.
21. Marshall, C.C. Annotation: from paper books to the digital library. *Proceedings of DL'97*, pp. 131-140.
22. Microsoft, *TabletPC*. (http://www.tabletpc.com).
23. Mockapetris, P.V. and K.J. Dunlap, Development of the Domain Name System. ACM SIGCOMM Computer Communication Review, 1995. **25**(1): pp. 112 - 122.
24. Muthitacharoen, A., B. Chen, and D. Mazières. A low-bandwidth network file system. *Proceedings of SOSP'01*, pp. 174 - 187.
25. Nabeshima, S., S. Yamamoto, K. Agusa, and T. Taguchi. MEMO-PEN: a new input device. *Proceedings of CHI'95*, pp. 256 - 257.
26. Rosenblum, M. and J.K. Ousterhout, The design and implementation of a log-structured file system. ACM Transactions on Computer Systems, 1992. **10**(1): pp. 26 - 52.
27. Schilit, B.N., G. Golovchinsky, and M.N. Price. Beyond paper: supporting active reading with free form digital ink annotations. *Proceedings of CHI'98*, pp. 249 - 256.
28. Sellen, A.J. and R.H.R. Harper, *The Myth of the Paperless Office*. 1st ed. 2001: MIT press.
29. Stifelman, L., B. Arons, and C. Schmandt. The audio notebook: paper and pen interaction with structured speech. *Proceedings of CHI'01*, pp. 182 - 189.
30. Wellner, P., Interacting with paper on the DigitalDesk. Communications of the ACM, 1993. **36**(7): pp. 87 - 96.
31. Yen, S.J., Capturing multimodal design activities in supportof information retrieval and process analysis, PhD thesis, Stanford University. 2001