Paper Augmented Digital Documents

François Guimbretière

Department of Computer Science
Human-Computer Interaction Lab
University of Maryland,
College Park, MD, 20742
francois@cs.umd.edu

ABSTRACT

Paper Augmented Digital Documents (PADD), are digital documents that can be manipulated either on a computer screen or on paper. PADD, and the infrastructure supporting them, can be seen as a bridge between the digital and the paper worlds. As digital documents, PADD are easy to edit, distribute and archive; as paper documents, PADD are easy to navigate, annotate and well accepted in social settings. The chimeric nature of PADD makes 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 PADs using today's technologies and reports on the lessons we learned while implementing the first PADD system.
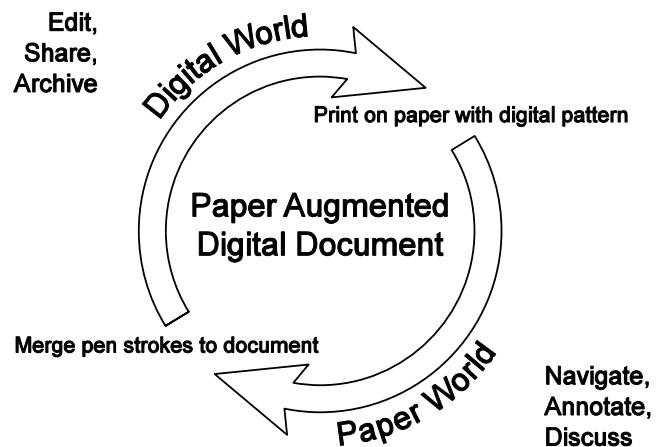
Keywords: Paper Augmented Digital Document, Paper based user interface, digital pen.

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" [25], 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 [27], Ariel [16] and the A-book [18], proposed bringing digital resources to paper. Other, such as Xax [11], Intelligent Paper [5], Audio Notebook [26] and Anoto [2], used paper as the input medium. Others still, such as the Freestyle system [13] or XLibris [24], explored how paper affordances could be provided on tablet computers, such as the recent Tablet PC.

The work presented here explores a fourth track, which has been given little attention in the past: *cohabitation*. We believe that the best way to bridge the gap between paper

Figure 1: Life cycle of a Paper Augmented Digital Document. PADD 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

and digital documents is to let users access any given document either through a digital interface or through a paper based interface as they see fit.

Toward that goal, we introduce the concept of Paper Augmented Digital Documents (PADD). PADDs are digital documents with all the digital affordances, but which also let users interact with them as printed documents. When printed, the documents offer all the affordances of paper and can record marks performed on them using a digital pen. At any time, the input gathered during interactions on paper can be merged with the original document to be processed as a new editing cycle starts (Figure 1). This cycle of use reflects the transient role of paper in the few successful paperless companies, such as DanTech in [25], and is well adapted to many activities which currently rely heavily on paper, such as proofreading, editing of drafts, and annotations of large format documents, like blueprints.

In this paper, we present the concepts behind PADD, how these concepts can be implemented using current technologies such as the Logitech io pen [14], and report on

our implementation experience of a prototype system based on an extension of the Adobe Acrobat application [1].

MOTIVATIONS

This work was inspired by a salient breakdown in the everyday life of a majority of knowledge workers [25]: 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 in to 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 [25] we believe that the hegemony of paper as a prime medium for review is rooted in its three fundamental affordances:

- Paper is easy to annotate. First because it is porous and holds ink, but also because it is light in small quantities, and flexible. These characteristics make it the perfect candidate for use in difficult environments, such as a construction site or public transportation where tight space, awkward environment 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 form factor sheets, such as the ones typical in engineering and architecture, or by creating a dynamic display of smaller pieces, as when combining several document pages.

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, 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 of 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 documents.

Typical uses

While PADD 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. She can next distribute her comments, or start acting upon them. 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 PADD 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 [26], or video stream, such as in the RECALL system [28], 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 [24] on the software side, and the Tablet PC [19] 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 reliability to accept the environment of an unfinished building. It is unclear if such a system can be built in the near future, but its functionality can be provided through PADD. We will come back to this point in the discussion section.

PADD

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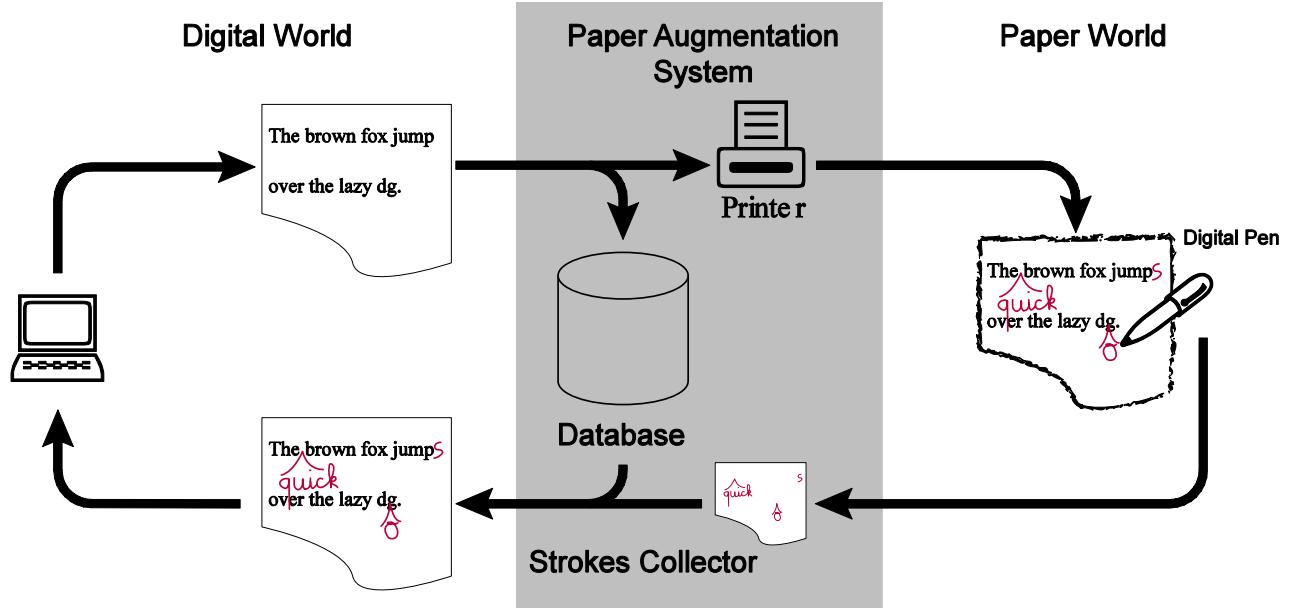
Figure 2: The PADD infrastructure. When paper affordances are needed, a snapshot of the PADD is stored in the PADB and the PADD is printed on the PA printer. The PA printer acts like a normal printer but adds a pen-readable pattern to each document. Using a digital pen, the document can now be marked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

paper when the affordances of paper are needed for any particular task. Figure 2 shows a typical PADD life cycle: PADD are created as digital documents such as a word processor document 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 for marks to be merged back. The document can now be used as a normal printout, easy to navigate, mark with a digital pen, and spread out. When users wish to transfer their markings back to the digital version of the PADD, they merely 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: the Paper Augmentation Database, which records the correspondence between digital and physical page; the Paper Augmentation Printer, which provides printout augmented with a pattern for pen tracking; a digital pen; and the Paper Augmentation 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. We will now describe each component in more detail.

Paper Augmentation Database (PADB)

The role of the PADB 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 PADB assigns a sheet ID, computes the transformation matrix between document space and paper space, and records these pieces of information, for each page in the document. The PADB also creates a copy of the document. The PADB needs to safeguard the *original* copy of the document, not the page description language used to print the document, so that the Stroke Collectors can retrieve a fully functional PADD.

The PADB also provides access control services to manage who can update the 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 as a unique address – who performed the marks.

Paper Augmentation Printer (PAP)

The Paper Augmentation Printer acts as a normal printer from the user perspective, but 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 [2] pattern for this feature but other system such as DataGlyphs [8], can be used, too.
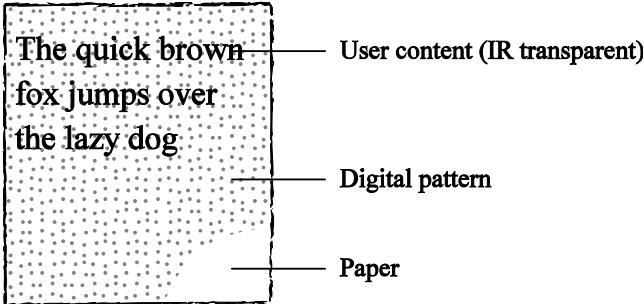
Figure 3: On paper, a PADD is composed by superposing a digital pattern (used by the digital pen to track movement) and user content. For PADD using the Anoto pattern (as shown here), user content needs to be printed with an IR transparent ink so that the pen can see the underlying pattern.

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

Paper Augmentation Strokes Collector (PASC)
Paper Augmentation Stroke Collector is the intermediary between the digital pen and the PADB. Upon being contacted by the pen, the PASC retrieves the corresponding PADD pages from the PADB, and has the PADD process the marks provided by the pen.

Accessing the PADB is more complicated that 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 [25]. In practice, this means that one could create marks on a piece of paper which was not printed locally and will need to access a remote PADB. This problem might be solved by any of many distributed directory services (such as the internet DNS service [20]). As an example, we show in Figure 4 how the system used by Anoto can be extended for our purpose. In this approach, upon reception of the page address, the PASC will contact the Paper Lookup Service and will receive back a Universal Resource Locator for the corresponding PADB. The PASC can now collaborate with the document PADB to process the markings.
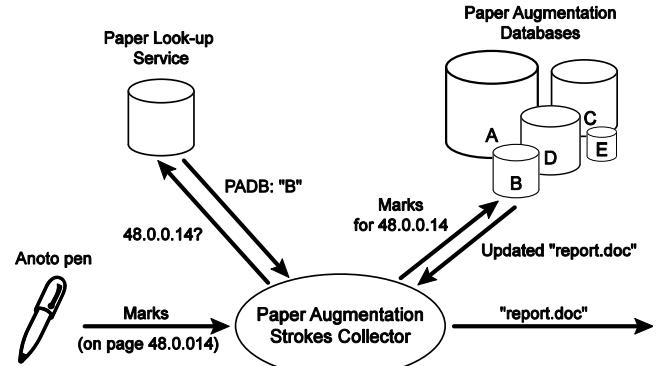

Figure 4: A possible implementation of the PASC following Anoto proposed architecture [2]. From left to right: the PASC retrieves the strokes from the pen as well as the page address on which they were created. The PASC then requests the PADB, which manages this page from the Paper Look-up Service. The PASC can then contact the PADB directly and merges the imported strokes with the PADD.

The final step of this process is to have the PADD process the marks 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 PASC, process the marks. 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 marks and defer their processing, allow for immediate processing, or provide a digital version of the PADD with the marks processed to the author of the marks.

Depending of the application, the processing can take a wide variety of forms. Some marks can be just drawn as an overlay; 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 system used in the Paper PDA [4] or pidgets in the Anoto system. Some might be synchronized with other sources of information, such as in the Recall system [28] and Audio Notebook [26], and some might be processed as they are created to provide immediate feedback to the users, such as in the Intelligent Paper system [5].

PROTOTYPE
To test the validity of the PADD concept, we implemented a PADD infrastructure prototype. The design goal behind our prototype was to 1) provide a proof of concept of the full life cycle of a PADD document; 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 PADB or
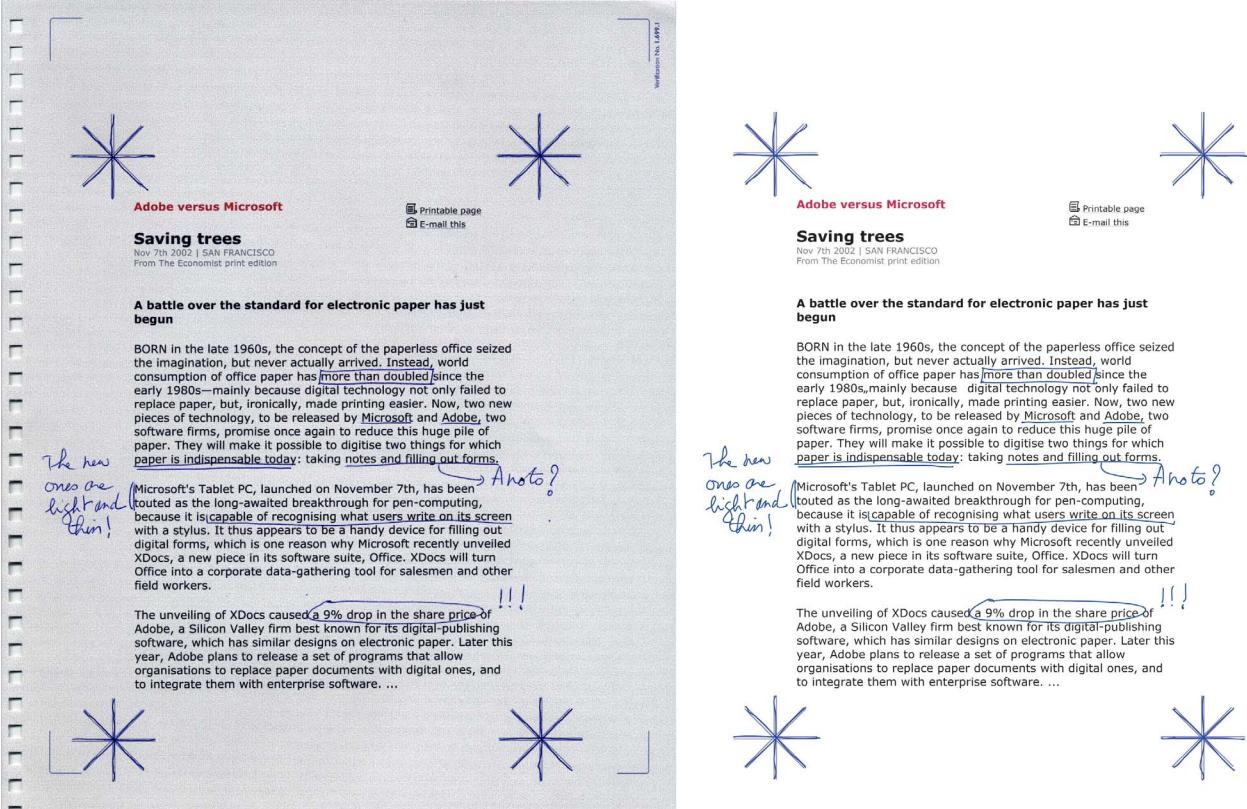
Figure 5: Two representation of the same PADD. On the left we show the scanned version of the paper representation annotated with the io pen. On the right its digital equivalent after synchronization. The PADD was printed on sheet address 44.0.0.23. Calibration was performed on sheet address 44.0.0.21. Both sheet are from a A4 notebook provided with the Anoto evaluation kit. Both pictures have been processed as JPEG (300 dpi). Close-ups are shown Figure 6.

the directory infrastructure used by the PASC were not important. However, existing solutions, proposed in different context such as the internet DNS [20] or Anoto's Paper Look-up Service [2], 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. This choice seemed natural as it allows us to easily compare our system to its contender the Acrobat digital annotation system. Our plug-in plays both the role of the PADB and the PASC as described below. By itself, the plug-in provides a valid PADD infrastructure with the limitation that it can only be used by a single user and only deal with PADD printed locally.

Acrobat Paper Augmentation Plugin
Our Acrobat plug-in works in concert with a Logitech io digital pen [14]. Users can print Acrobat documents, annotate them using the io Pen, and then merge the annotations into the original document.

Paper Augmentation database
For this prototype, the PADB 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 this page. The annotation protocol provided by the Acrobat format [1] allows the guarantee compatibility of this approach with other plug-ins.

Paper Augmentation Printer
Potentially the most limiting component in the current prototype is the PAP. No commercially available printers can print the Anoto pattern on demand efficiently, so our implementation relies on the printer being loaded with pre-printed paper. This causes a reliability problem, since the printer and the PADB may become desynchronized. 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.
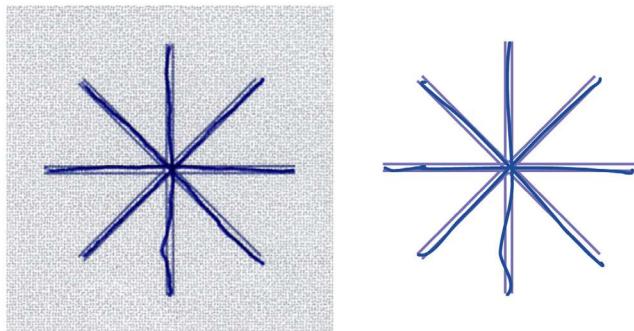
Figure 6: Close-up from Figure 5: Left we show the bottom right calibration mark. Right the text circled in the last paragraph. Note the similar look and feel of the scanned paper version (grey background) and the digital version. The accuracy of the calibration is good enough that a narrow letter like an 'i' could be crossed with accuracy. As in Figure 5 both versions are shown processed as JPEG (300 dpi).

In that configuration, the printer enters a "reserve mode"[1] and uses a mix of Cyan, Magenta and Yellow to produce an 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.

Paper Augmentation Stroke Collector
Our plug-in implements a lazy version of the PASC by integrating the strokes on user demand. To process the request, the PASC looks up the annotations created by the PADB on each page. It then 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 [1]. In the current implementation, all strokes in a page are part of the same annotation placed as an 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 document 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 to get pre-printed paper, the initial
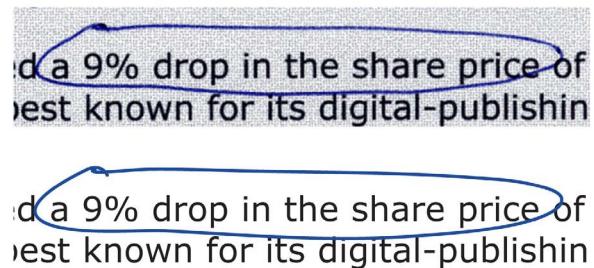
reaction of potential users during our presentation at the first Anoto developer conference was very positive.

DISCUSSION
The infrastructure presented in this paper is a unifying framework for several solutions proposed in the past, including Xax [11], Intelligent paper [5], Anoto [2], and the Paper PDA [9]. For example, in the Anoto system, each PADD (a "form" in Anoto terminology) is created once and accessed many times by each potential users. The pen can either process strokes as they happen (upon checking the Anoto "Magic Box©" for example) or in batch mode (while drawing). Mark processing does not modify the form, but instead issues the corresponding command. Unlike these systems, PADD provide a solution for a full digital-paper-digital cycle.

Because our system provides support for cohabitation between paper and digital document, it also opens new interaction designs for situations where paper is convenient but poorly integrated with the digital world. We have already discussed several possible applications, such as editing text or annotating large drawings used in engineering, but other examples include:

- Bridging the digital divide that might exist between members of a collaborating group. For example, our lab is studying how these PADD can be adapted as a family calendar system that is accessible not only to tech savvy family members but also to elderly who might not be willing to deal with a computer and who prefer paper.

- Bridging the gap between computational resources and paper reliability. For example, while computer versions of flight strips might be more efficient at managing the traffic, paper versions are still preferred because of their reliability [17]. PADD 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

---

[1] 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.

that the system can provide advisory information to the controller.

A main advantage of our proposed infrastructure is that it does not require a radical change in usage patterns. Instead, PADD 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.

PADD versus tablet computers

With the recent advance in tablet computing, such as the Tablet PC introduced by Microsoft Corporation [19], 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 [10] and annotation manipulation software such as XLibris [6, 24], 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 like editing a paper during a flight, or at a café terrace. In these later cases, we believe that PADD offers a more flexible way to interact with a document.

Ease of navigation

While it is true that some research efforts have been made to develop better navigation interfaces [7], 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 paper makes it very easy to spread documents 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 PADD 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. For all these cases, PADDs provide the flexibility to let each user choose the most efficient solution.

Infrastructure management

To maintain the correspondence between the digital and the paper version of a PADD, the PADD infrastructure relies on maintaining a large database of printed documents and a large pattern space. In this section, we explore how these requirements can be met.

PADB

As described above, the PADB 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, PADB 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 word processor document will be efficiently handled by a versioning system dealing with the original document, but might cause the layout of the full document to change leaving little room for versioning between the corresponding printout files.

Second, because the a PADD may be printed several times, using versioning, such as the one proposed by Muthitacharoen et al. [21], or leveraging file system features, such as log-based file systems [23] may considerably limit the memory requirement of the PADB.

Pattern space

While the Anoto pattern space is very large[2], it may be too small if the PADD system were to be universally applied and each piece of paper printed required a unique ID[3]. This problem is not specific to the Anoto system but 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. This approach would greatly increase the

---

[2] In excess of 2^{48} pages [2].
[3] In 2000, the United States consumed 15.7 M tons of uncoated free-sheet paper [25], or around 2^{42} sheets.

available pattern space at the cost of possibly missing small strokes.

Prototype implementation issues

The implementation of our first PADD infrastructure provided us with valuable insights for the development of a full scale PADD system. We focus here on calibration since it was 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 pen were very accurate and provided at a high sampling rate. Our tests showed that data provided by the Logitech io 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 also found that even entry-level inkjet printers, such as 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 since our infrastructure records the paper space to digital space transformation on a page-by-page basis.

Our experience confirmed that a PAP printer 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 PADB, 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 greatly improved if the printer were able to read the pattern on the page to be printed, guaranteeing that the printer and the PADB 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. In out tests, we found that the paper provided by Anoto was very repeatable but probably quite expensive, while the paper provided with the Logitech pen exhibited small variations but is very cheap. 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.

RELATED WORK

Many systems have been proposed to narrow the gap between the paper and the digital world. One research direction proposed to use paper as an input medium to gather information to be processed by a computer at a later time. These systems include XaX [11], the Paper PDA [9], Anoto [2]. They all share the characteristic that the interaction on paper does not require the use of a nearby computer. Anoto is probably the most developed system and is currently used in form filling applications and to transmit fax (or MMS) through the cell phone infrastructure. Our proposed PADD system extends these systems by providing an easy way to close the digital-paper-digital cycle.

Another research direction was to augment normal paper interactions by coupling them with the digital world. Such systems required the concurrent use of paper and computers. These systems include the DigitalDesk [27], Ariel [16], VideoMosaic [15], PaperLink [3], Inteligent paper [5], the EnhancedDesk [12], and the A-book [18]. While the PADD infrastructure as described in this paper could be used to implement these systems, our focus is different: we believe that cohabitation rather than coupling will serve user's needs better. Our focus is probably best reflected in another of such a system, the Audio Notebook [26] which separates note taking and note processing.

Yet another research direction is to provide paper affordances on digital computers. This direction of research produced systems like the Freestyle system [13], the tangible manipulation interface proposed by Harrison et al. [7], and the XLibris system [6, 24]. New technology like the Tablet PC [10] is making this approach more and more appealing, yet current tablet computers are still too bulky and expensive for casual use like annotating a printed document in a café.

Our system relies on paper based tracking technology. While our prototype is using the Anoto system [2], other technologies have been proposed including DataGlyphs, [8], and MEMO-PEN [22]. Both Anoto and DataGlyphs are well adapted for our system since they offer a large absolute address space.

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 extending the work presented by Golovchinsky et al. in [6] so that strokes can be included within an active word processing document and flow with the text. We are also collaborating with the CIFE group at

Stanford to extend their Recall system [28] so that it can be used with paper, as well as with a tablet computer.

We also intend to use our system as an analytical tool to explore how the use of digital documents and their paper representations differ. While qualitative studies have been performed in that direction [25], 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 in that direction by providing a large accurate corpus of data for analysis.

CONCLUSION

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

We see Paper Augmented Digital Documents (PADD) 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 yet can be discarded when the information collected on paper is merged back with the digital counterpart and the affordances provided by digital media are more appropriate.

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.

REFERENCES
1. Adobe, Acrobat. http://www.adobe.com.
2. Anoto, *Development Guide for Service Enabled by Anoto Functionality*. 2002: Anoto.
3. 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.
4. 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.
5. Dymet, M. and M. Copperman. Intelligent Paper. *Proceedings of EP'98*, pp. 392 - 406.
6. Golovchinsky, G. and L. Denoue. Moving markup: repositioning freeform annotations. *Proceedings of UIST'02*, pp. 21 - 30.
7. 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.
8. Hecht, D.L. Embedded Data Glyph Technology for Hardcopy Digital Documents. *Proceedings of SPIE Color Hard Copy and Graphic Arts III*, pp. 341 - 352.
9. 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.
10. Jarett, R. and P. Su, *Building Tablet PC Applications*. 2002: Microsoft Press.
11. 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.
12. 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.
13. 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.
14. Logitech, IO digital pen. http://www.logitech.com.
15. Mackay, W. and D. Pagani. Video mosaic: laying out time in a physical space. *Proceedings of International Multimedia Conference*, pp. 165 - 172.
16. 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.
17. 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*, pp. 558 - 565.
18. 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.
19. Microsoft, TabletPC. http://www.tabletpc.com,
20. Mockapetris, P.V. and K.J. Dunlap, Development of the Domain Name System. ACM SIGCOMM Computer Communication Review, 1995. **25**(1): pp. 112 - 122.
21. Muthitacharoen, A., B. Chen, and D. Mazières. A low-bandwidth network file system. *Proceedings of the eighteenth ACM symposium on Operating systems principles*, pp. 174 - 187.
22. Nabeshima, S., S. Yamamoto, K. Agusa, and T. Taguchi. MEMO-PEN: a new input device. *Proceedings of CHI'95*, pp. 256 - 257.
23. 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.

24. 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.

25. Sellen, A.J. and R.H.R. Harper*, The Myth of the Paperless Office*. 1ˢᵗ ed. 2001: MIT press.

26. Stifelman, L., B. Arons, and C. Schmandt. The audio notebook: paper and pen interaction with structured speech. *Proceedings of CHI'01*, pp. 182 - 189.

27. Wellner, P., Interacting with paper on the DigitalDesk, in *Communications of the ACM*. 1993. pp. 87 - 96.

28. Yen, S.J., Capturing multimodal design activities in supportof information retrieval and process analysis, PhD thesis, Stanford University. 2001