

A Systems View of Annotations

Steve B. Cousins, Michelle Baldonado

Information Sciences and Technology Lab

Xerox Palo Alto Research Center

Email: {cousins, baldonado}@parc.xerox.com

Andreas Paepcke

Computer Science Department

Stanford University

Email: paepcke@db.stanford.edu

ABSTRACT

Motivated by our design of an annotation system (Notable), we have developed a framework for annotation systems, based on dimensions of the platform used to create the annotations, the platform used to read them, the annotations themselves, the targets of the annotations, and the correspondence between annotation and target. We demonstrate the framework by using it to compare 10 different annotation systems, and to find unexplored possibilities in the design space.

INTRODUCTION

The notion of ‘annotation’ refers to many different practices. It refers to highlighting words in a text, placing marks or words in a margin, attaching Post-It™ notes to objects, publishing an annotated edition of an important work, etc. The term ‘annotation’ has also been used to describe more formal activities, such as assigning metadata to a literary work. As Marshall points out, even a photograph of a sky tucked into a meteorology textbook might be called an annotation [3]. With the growth of digital media, the document tradition is moving forward while the annotation tradition is being left behind. This paper sheds light on the space of annotation systems, in order to help designers of annotation systems see their way forward.

The move towards electronic documents only makes it more difficult to say precisely what is meant by ‘annotation,’ and therefore to understand what is expected of a system that provides an annotation facility. Some writers argue that appearance *in situ* is a necessary characteristic of any successful annotation system [6]. Others argue that the location of an annotation merely affects the style of the annotation. For example, studies of voice annotations have shown that independent annotations such as comments on a paper left on voice mail tend to be higher-level comments than those written directly on a paper [1]. There is room for a range of digital annotation systems that support both *in situ* and independent annotations.

At the former federal penitentiary Alcatraz, visitors are given devices that play audio annotations specific to various landmarks within the facility. While this is far from a traditional annotation platform, it gives a hint about how much technology changes our view of concepts like annotations. This particular system highlights a subtle point: the system used to author the annotation does not need to be the same as the one used to render the annotation (the technology used at Alcatraz, tape players, does not even support recording).

We have developed a broad definition of annotation that encompasses these examples and many more:

Definition of annotation: An annotation is a commentary on an object that:

- the annotator intends to be separable from the object itself
- the reader interprets to be separable from the object itself

With this definition objects, as well as documents, can be annotated (imagine a sticker on a computer monitor that says “broken”). There are three roles involved in the practice of annotation: the *author* created the object that is being annotated; the *annotator* is the one commenting on the object; and the *reader* is the agent making sense of the commentary. The same person often fills two or more of these roles, and sometimes a role is filled by a group of people. When a student highlights a textbook for later study, she is both annotator and (while studying) reader. When a group of people pass around marked-up drafts of a paper they are jointly writing, they fulfill all three roles.

This paper takes a broad, systems view of annotations. The primary contribution is a framework for comparing annotation systems, both paper-based and electronic. The analysis is motivated by an annotation system we have developed (Notable) that separates the annotation “platform” from the object being annotated, and uses search to bridge the gap between the annotation and the annotation target.



Figure 1: Annotating a document with Notable: (a) initiating a search for the document to annotate, (b) selecting it from the search results, and (c) creating an annotation.

We designed Notable to fulfill a desired usage scenario: to be able to create, on a handheld device, annotations that can be linked to paper documents. The motivation for this scenario arose from our realization that paper and networked handheld devices are complementary for annotation tasks. Paper works well for reading because it offers high resolution, portability, and tangibility. Networked handheld devices work well for authoring annotations because they offer connectivity, portability, and a Post-It™ - like display size.

The canonical Notable example involves a conference reviewer who receives a paper in the mail. She scans the paper document into a searchable repository to create an electronic surrogate version of the document for reference. (If the paper came in an electronic form, she could add it to the repository directly). She can then later (perhaps at home in the evening) search for the document's electronic handle using the search capability of Notable (Figure 1-a), select it if the search produces multiple results (Figure 1-b), and enter annotations on the device as she reads from the paper. When she is finished, the annotations she has made are transmitted to the repository, and linked to the document. The annotations can later be viewed and searched on a variety of platforms, not just the handheld device.

Three features particularly distinguish Notable from other annotation systems. First, Notable allows the user to create annotations and read the original work

in two different media (handheld device and paper). Second, it enables the user to link annotations and texts via search. Third, it supports users in mobile environments by letting annotation occur solely on very small devices (assuming that scanning takes place prior to annotating).

In presenting Notable to potential users, we learned that they identified with the scenarios, but wanted to be able to annotate documents at a finer level of granularity. This desire presented us with a design question, as we had learned that audio annotations can lead to higher-level comments. We decided to do a detailed analysis, in order to learn from other systems; this analysis led to a framework for annotations.

In the next section we describe nine different annotation systems, and then proceed to describe the framework we have developed for characterizing annotation systems. The framework allows us to quickly compare systems, and to notice gaps in the space where new systems might be developed, or new features might be added to existing systems.

ANNOTATION SYSTEMS EXAMPLES

Annotation systems range from the traditional mark-up of published works to the non-traditional, such as the Alcatraz example in the introduction. In this section, we survey a few different systems. We will use these examples to motivate and describe our framework in the next section.

Direct paper annotation

The paper medium affords a plethora of annotation practices, ranging from scribbles on paper to the creation of annotated editions or published movie reviews. We use the phrase “direct paper annotation” to refer to the traditional example of annotating a paper document by marking it up with pen or highlighter.

Paper with Post-It™ notes

The ubiquity of Post-It™ notes in the office setting suggests that direct paper annotation is not sufficient for all tasks and purposes. Advantages of Post-It™ notes include their visibility, their removability, and their convenient size. Visibility is important to creating annotations that are likely to capture the reader’s attention. Removability means that comments can disappear when they are no longer appropriate. It also frees an annotator to add comments to a document that cannot be altered, such as a library book. Finally, the Post-It™ size is both large enough to contain a note that could not be easily squeezed into the whitespace available on a page, and small enough to avoid occluding too much of the paper to which the Post-It™ is attached.

Annotated edition

An annotated edition of a work (e.g., an annotated edition of the Talmud or an annotated edition of a Shakespeare play) collocates the text of the original work with additional paratext by the editor. Often, the original text and the annotations are placed side by side on the page or on opposing pages of the book. Authoring an annotated edition differs from creating direct paper annotations in that the author can choose a spatial layout that grants more room to the annotations and that more clearly links the annotations to the corresponding pieces of text. For example, extra-wide margins are common in annotated editions.

X Libris

The XLibris system [8], a combination of a high-resolution pen tablet and a paper-like user interface, provides users with the ability to create free form digital ink annotations. In fact, the practice of annotating in XLibris is fairly similar to annotating on paper, modulo the differences in tactile feedback and display resolution. XLibris enriches these paper-like annotations by adding the possibility of active annotations. For example, the act of circling causes a query to be computed based on the text within the encircled region.

Multivalent annotations

Multivalent Documents (MVD) [6] are composed of content layers and dynamic behaviors. In this context, multivalent annotations correspond to specialized behaviors on a layer. Examples of annotation-based behaviors include annotations that highlight the underlying text, annotations that are hyperlinks, annotations that are interpreted as copy editor markup (when active, these annotations can in fact be executed), and ink annotations similar to those supported by XLibris.

Microsoft Word annotations

The MS Word desktop word processing software also has an annotation feature. Composition of a textual annotation takes place in a subwindow within the main editing window. In the text, comments are indicated by inlining the annotator’s name inside of double brackets. Readers of the document may filter comments by author and may also opt to hide all comments.

ComMentor

ComMentor [7] is an annotation system for the World Wide Web. Both annotators and readers utilize an augmented version of the Mosaic Web browser. Composition of annotations takes place in a separate window, but annotations are “pinned” to the underlying document and rendered as character-sized inlined images (often, a picture of the annotator). At an annotation point, a reader can view the annotation in a Post-It™ size display or in a separate Web browser window. All annotations can be built upon with further annotations, as ComMentor offers the annotator a set of annotation types to choose from, including follow-ups and trails.

ComMentor differs significantly from the Word model in that the annotator does not have the ability to modify the original document. ComMentor uses several strategies to make the links robust in the face of changes to the original document.

Additional important features of ComMentor include its notions of annotation sets (analogous to annotation topics) and groups, which allow for flexible handling of a spectrum of private to public annotations.

NotePals

NotePals [2] is a group-based note-taking system designed for small handheld devices (e.g., Palm Pilots). In our context, we understand NotePals as a group-based annotation system for meetings. Users of

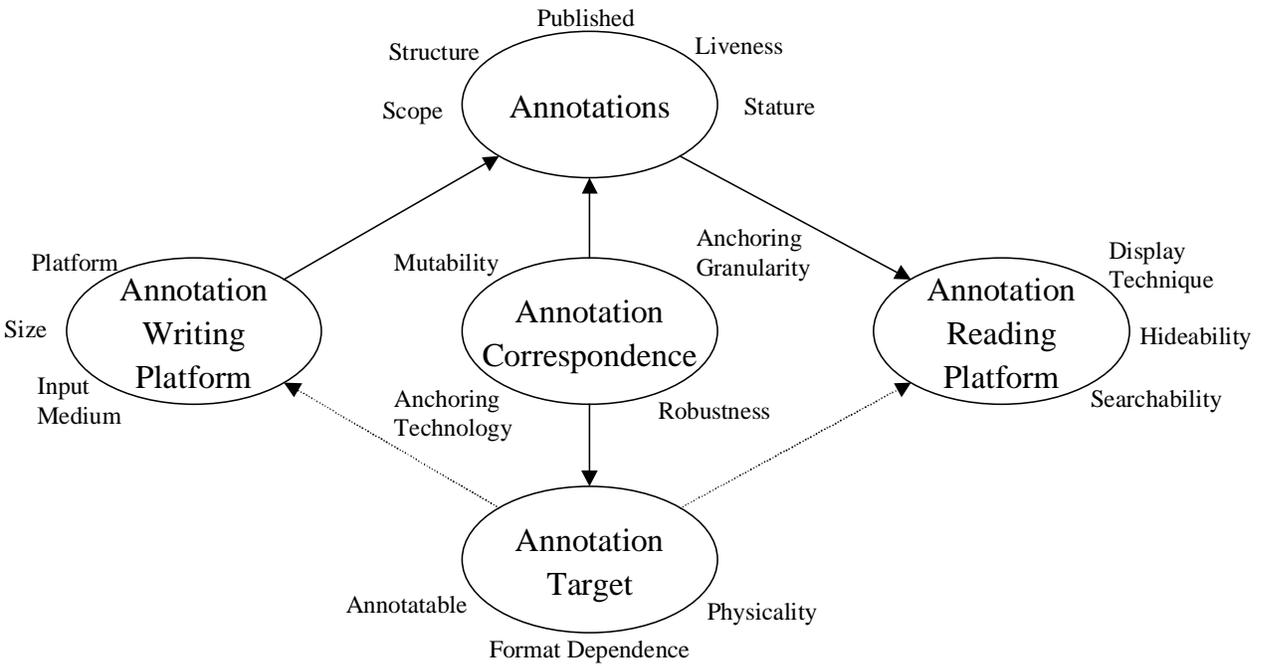


Figure 2. Five annotation system components and their design dimensions

NotePals assign their notes a particular category and time span, then author ink-based notes on the meeting. At synchronization time, all notes from the group that have the same category and time span are merged together and made accessible to the group. Accordingly, members can take turns authoring notes and can easily share ink-based notes. Subsequent work on NotePals has refined the level of granularity for which notes can be taken by using timestamps to link individual notes with audio, agendas, and presentation slides.

Tapestry

Tapestry [4] is both an e-mail system and an annotation system. Subscribers to e-mail lists are able to provide annotations on messages sent to those lists. These annotations can then be used for collaborative filtering, allowing other subscribers to limit their e-mail messages to those that satisfy certain annotation constraints (e.g., are labeled “interesting” by three of your colleagues).

FRAMEWORK

Annotation systems, from a computer science perspective, can be analyzed in terms of characteristics of the *annotation writing platform*, the *annotation reading platform*, the *annotation targets*, the *annotations* themselves, and the *correspondence between annotations and their targets*. The annotation writing platform is the technology used by

the annotator, and the reading platform is used by the reader. In the Alcatraz example from the introduction, the reading platform is the listening device that patrons carry, and the writing platform was a tape recording and editing system. The annotation target is the object being annotated, like Al Capone’s cell, or a document. The correspondence can be as direct as ink on paper or as indirect as Notable. We will enumerate dimensions within each of these areas in the subsections below.

For each of these five major dimensions, we articulate a small set of dimensions and list possible values for each. Figure 2 summarizes the space we have defined. In the rest of this section, we describe our current understanding of each dimension and sub-dimension, giving examples from the systems above. In the next section we describe each system completely in terms of these characteristics.

Annotation writing platform

The annotation writing platform is the system used to create the annotations. We compare the writing platforms of various annotation systems on the basis of the input medium they use, whether the platform is collocated or separate from the platform used to view the annotation target, and the size of the annotations supported.

Input medium: Ink vs. keystrokes vs. multimedia

The fundamental distinction for the systems we have analyzed is whether the annotations are made with ink (real or digital) or with keystrokes. Ink-based annotations are not generally interpretable by the system, whereas keystroke-based annotations generate a character representation.

The technology used to create the annotation influences but does not determine the input medium. Pen-based systems like Notable can be used to create keystroke-based annotations, using virtual keyboards or special writing systems like Graffiti™. A mouse can be used to make ink-like annotations, although drawing with a mouse is generally more awkward than drawing with a stylus. The distinguishing feature is whether the annotation system can interpret the annotation marks.

Highlighting is a special case of ink-based annotation. It may be done with a regular pen by underlining or bracketing text, or with a special highlighter pen, which changes the background color of the text.

Finally, although none of the systems we formally present is a voice or video annotation system, we include multimedia as an explicit choice along this dimension. We expect that multimedia annotations will be closer to ink than keystrokes, since they can not be easily interpreted by the system.

Platform: Separate vs. collocated

Collocated platforms render the annotation target within the annotation authoring environment. All of the systems in our list except Notable and NotePals are collocated. In paper-based systems, the annotator interacts directly with a rendering of the document. MVD and X Libris attempt to simulate paper in their approach. ComMentor falls in between, showing target and annotation authoring surfaces side by side.

Another writing platform that is not collocated is voice mail. For example, an advisor might leave comments on a graduate student's thesis by leaving a detailed message.

Display and storage size: Margins vs. fixed vs. unlimited

This dimension describes the capacity of the comment area, or the amount of commentary supported. Are the annotations written in margins, written on 2-inch square PostIts™ (fixed size), typed alongside the text (possibly causing repagination), typed in a separate window, dictated onto a tape with limited duration, etc?

Annotation reading platform

The annotation reading platform is the system that a reader interacts with to view the annotation target and the annotations. With paper-based annotations, it is often the same as the annotation writing platform, but in general it can be a completely different system. An annotated edition of a book provides a clear example: the annotator (a.k.a. editor) does not write in the margins, but uses a word processor or other technology to produce the book; the reader does not have access to the annotation writing platform.

Display technique: Inline vs. fluid vs. overlaid vs. independent

With electronic documents, there are even more ways to display annotations than with paper documents. By our definition of annotations, the base requirement of an annotation display technique is that both the annotator and reader can distinguish the commentary from the source text.

With direct paper annotation and digital systems that simulate it, annotations are displayed with marks overlaid on the text and margins. Post-It™ techniques (real and simulated) also involve overlaying.

Annotated editions typically use an inline display technique. Text size, font, color, and position are all used to distinguish the annotation from the source text, but the annotations are still integral to the document, and affect the size of its pagination and other layout properties. MS Word annotations are even more obviously inline: they can even be sentences within a paragraph of the target document.

Fluid display techniques use animation to present supporting information [10]. Fluid techniques can be inline, as when text lines move apart to make room for an annotation, or overlaid, as when the source text fades and the annotation is rendered on top of it.

Finally, it is possible to render some annotations independently of their source text. Metadata is often rendered in this way, especially metadata about ratings. Voice annotations, or comments sent via email also fall into this category. Generally, annotations with document granularity are displayed independently from the annotation target.

Hideability

Can the annotation target be viewed without the annotations? This attribute singles out pen and paper annotations, but could also apply to some digital systems, especially those designed with security requirements. For example, copyright holders may insist that their copyright appear on every copy, or parents may insist that the television rating always be present with any video object rendered in the house.

Searchability: Searchable vs. not searchable vs. limited search

Limited search refers to the ability to search for marks, but not necessarily to distinguish them. The X Libris system lets users skip to the next annotation, but does not interpret the marks directly (so a user cannot search for the word “huh?” in the margin). In general, digital systems have some kind of search potential. Systems with keystroke input or stroke interpretation would be fully searchable.

Annotation targets

The annotation target is the object that is being annotated. The two dimensions here are “direct vs. indirect” and “digital vs. physical.”

Annotatable: Direct vs. indirect

By directly annotatable, we mean that the target is writable by the annotator. For example, to use the annotation facilities in MS Word, the annotator has to be able to change the Word document, because the annotation becomes part of the document representation. A paper that I control is directly annotatable. The Declaration of Independence is only indirectly annotatable (I can write about it, but I can't write on it without getting in a lot of trouble). Systems that support indirect annotation must provide a means for locating the annotations and rendering them with (or otherwise attaching them to) the annotation target.

Physicality: Digital vs. physical vs. either

Does the annotation target have to be a physical object, a digital object, or can it be either type? Notable supports annotation of both digital and physical annotation targets. In contrast, most electronic systems support only digital annotation targets, while most paper-based systems support only physical annotation targets.

Format dependence: strict vs. medium vs. none

Does the system allow arbitrary objects to be annotated, or does it only support objects in a particular format? Pens can be used to annotate most paper. Post-Its™ can annotate paper or many objects. By contrast, annotated editions are rather strict: as a technology it only applies to published works like books. In the digital world, systems range from very strict (MS Word annotations only work on MS Word format) to arbitrary (Notable can be used to annotate anything that can be named) as well.

Annotations

We include five characteristics of annotations: whether or not the annotation is published, for whom

it is published, whether it is formal or informal, whether or not it is active, and whether annotations themselves are documents.

Published vs. unpublished

For whom is this annotation intended and available? In the paper world my annotations in my copy of a book are only as sharable as the book- though this may be a wider audience than the annotator intends, as Marshall points out with her used textbook examples [3]. The other extreme in the paper world is annotated editions of works.

Scope: Global vs. institutional vs. workgroup vs. personal

This dimension addresses the intended audience for annotations. Note that an annotator may want to make annotations on the same document with different scopes. An annotator of a document about Java might provide different comments for the marketing department than for senior management. Both types are published, but each has a different scope. Similarly, I might distinguish annotations in a Word document with different colors, where each color is for a different recipient.

Structure: Formal vs. informal

Can the meaning of the annotation be described according to a formal system? The most common example of formal in this dimension is metadata, where the annotation value is associated with a particular slot. For example, a document may have a metadata slot for publication date, where the value can be characterized by some rules describing how dates are represented. On the other end of the spectrum are scribbles on paper.

When annotations are formal, they may have complex structure. For example, in Notable, indirect annotations are complex, in that they have a property that says which document they are annotating.

Liveness: Active vs. inactive

Traditional paper-based annotations have always been inactive, but in the digital world, the possibility of active annotations arises. For example, a piece of metadata like “last-edited” or “word-count” could be an active property that is constantly updated.

Stature: Document vs. fragment

Are annotations themselves documents? The Talmud is an example of this from the paper world, in which there is a long history of published annotations that are themselves annotated. Highlighter marks on paper are at the other end of this spectrum, which we refer to as fragments. MVD's layers of annotations fall somewhere in the middle: each layer has its own URL, but may not be independently displayable.

Annotation correspondences

The final dimension of our analysis focuses on how the correspondence between annotation and annotation target is specified and maintained. Marshall describes various mechanisms such as brackets and arrows that annotators use to mark up textbooks [3]. Our analysis is at a somewhat higher level. We lump all of the techniques Marshall describes together under “pixel granularity using physical juxtaposition” (two of our dimensions below). On the other hand, we also add two characteristics for which paper implies a particular value: whether the annotation target is mutable, and how robust the correspondence is.

Anchoring granularity: Pixel vs. character vs. page vs. document

Some tasks like copy editing require people to draw arrows and be able to point to particular locations on the page, while other annotation tasks simply require comments on a document as a whole. Between these extremes, which we call editing at pixel and document granularity respectively, are character and page levels of granularity. At character granularity, annotators can select a span of words and comment on them, but cannot easily link two spans together (as they could in pixel granularity by drawing an arrow). In general, if the annotation system only supports a coarse granularity, annotators are forced to use linguistic mechanisms to make finer-granularity comments.

Mutability: Mutable vs. immutable

Are annotation correspondences maintained in face of changes in the annotation target? If they are not, or if this question does not apply because the annotation target cannot change (as in an instance of a paper document), we term the target immutable. ComMentor includes a number of mechanisms to keep an annotation attached to the right location, in the face of changes to the annotation target. Notable has document granularity, so keeping the annotation attached is not difficult, but it may become irrelevant as the underlying annotation target changes.

Anchoring Technology: Physical juxtaposition vs. URL

How does the annotator refer to the annotation target? At a high level, the choice is by marking on it directly, or by referring to it (what we call URL correspondence). There are refinements of URL correspondence, of course. ComMentor used URL plus additional data to get character granularity. Tapestry pre-dated the web, so instead of URLs it used message IDs. In work at Berkeley, Landay’s group uses time to determine the correspondence of annotations with parts of presentations [2].

Robustness: Permanent vs. removable vs. “URL-stable”

Can correspondence be broken without removing annotation? URL stability refers to the correspondence being broken by normal web “bit rot.” When the correspondence is not permanent, a related issue is who maintains the correspondence.

USING THE FRAMEWORK

The framework we have created can be used for both analysis and design. We demonstrate its use in analysis by creating a table that allows the systems described above to be directly compared. Given the analysis, we can look for gaps in the design space.

Framework application 1: Analyzing existing systems

Table 1 compares the systems we have described above. Occasionally a system is ambiguous with respect to a property in one of the dimensions. In these cases, we mark the keyword with either a “+” (does more) or “-” (is more limited than), or replace it with “varies” when a system can implement a range of keywords. Entries with caveats are marked with a “(*)”; we do not have space to discuss those caveats in this paper.

The table can be used for pairwise comparisons, and we leave most of the 45 possible comparisons for the reader to explore. Here we describe a few that we have noted.

X Libris was designed to work like a direct paper annotation system, and clearly its designers did a good job. It only differs from paper on four of the 18 characteristics: hideability, searchability, physicality, and liveness. Hideable, searchable (limited vs. none), and active annotations (X Libris can find documents related to highlighted sections) are all improvements to direct paper. Physicality is the only remaining difference, and is quite understandable. Paper does have some technological advantages that do not show up in the table. For example, it has higher resolution and contrast than the hardware used in the X Libris prototypes.

MVD is very similar to X Libris. They differ in that MVD’s targets are indirect and have only medium format dependence, and its annotations are published and global, with somewhat different values in the annotation correspondence column. This all makes perfect sense given the different intended users: X Libris is designed for a reader to make notes to herself, while MVD supports making notes for others to see.

	Writing Platform <i>1. input medium</i> <i>2. platform</i> <i>3. size</i>	Reading Platform <i>1. display technique</i> <i>2. hideability</i> <i>3. searchability</i>	Annotation Targets <i>1. annotatable</i> <i>2. physicality</i> <i>3. format dependence</i>	Annotations <i>1. published</i> <i>2. scope</i> <i>3. structure</i> <i>4. liveness</i> <i>5. stature</i>	Annotation Correspondence <i>1. granularity</i> <i>2. mutability</i> <i>3. technology</i> <i>4. robustness</i>
Direct paper	1. ink 2. collocated 3. margins	1. overlaid 2. not hideable 3. not searchable	1. direct 2. physical 3. medium	1. unpublished 2. personal 3. informal 4. inactive 5. fragment	1. pixel 2. immutable 3. juxtaposition 4. permanent
Post-Its™	1. ink 2. separate 3. fixed overlay	1. overlaid 2. hideable 3. not searchable	1. direct 2. physical 3. none	1. unpublished 2. personal 3. informal 4. inactive 5. document	1. page+ (*) 2. mutable 3. juxtaposition 4. removable
Annotated Edition	1. keystrokes 2. depends (*) 3. unlimited	1. inline 2. not hideable 3. limited search (*)	1. direct 2. physical 3. strict	1. published 2. global 3. informal 4. inactive 5. document	1. page 2. immutable 3. juxtaposition 4. permanent
X Libris	1. ink 2. collocated 3. margins	1. overlaid 2. hideable 3. limited search	1. direct 2. digital 3. strict	1. unpublished 2. personal 3. informal 4. active 5. fragment	1. pixel 2. immutable 3. juxtaposition 4. permanent
MVD	1. ink- (*) 2. collocated 3. margins	1. overlaid 2. hideable 3. limited search	1. indirect 2. digital 3. none	1. published 2. global 3. informal 4. active 5. document- (*)	1. pixel 2. mutable- (*) 3. juxtaposition 4. URL-stable
MS Word	1. keystrokes 2. collocated 3. unlimited	1. inline 2. hideable 3. searchable	1. direct 2. digital 3. strict	1. unpublished 2. personal 3. informal+ (*) 4. active 5. fragment	1. character 2. mutable 3. juxtaposition 4. removable
ComMentor	1. keystrokes 2. collocated 3. fixed overlay+ (*)	1. independent 2. hideable 3. searchable	1. indirect 2. digital 3. medium	1. published 2. varies (*) 3. informal 4. inactive 5. document	1. character 2. mutable 3. URL 4. URL-stable
NotePals	1. ink 2. separate 3. unlimited	1. independent 2. hideable 3. limited search	1. indirect 2. varies (*) 3. none	1. published 2. workgroup 3. informal 4. inactive 5. document	1. document+ (*) 2. immutable 3. URL+ (*) 4. removable
Tapestry	1. keystrokes 2. collocated 3. unlimited	1. independent 2. hideable 3. searchable	1. indirect 2. digital 3. strict	1. published 2. workgroup+ (*) 3. formal 4. active 5. fragment	1. document 2. immutable 3. URL- (*) 4. permanent- (*)
Notable	1. keystrokes- (*) 2. separate 3. unlimited	1. independent 2. hideable 3. searchable	1. varies (*) 2. varies (*) 3. none	1. published 2. varies (*) 3. varies (*) 4. inactive 5. document	1. document 2. mutable- (*) 3. URL 4. URL-stable

Table 1: Comparison chart for existing systems using the framework

Notable and NotePals are both Palm Pilot-based annotation systems. The table shows that they are similar in some respects, although they were designed for very different tasks. By this analysis, if NotePals used URLs to refer to events and could render keystroke-based annotations in its annotation reader, the Notable annotation writing platform could be used as input to the NotePals reading platform.

Notable could also be used to construct Tapestry annotations, if it were modified to work with Tapestry's document IDs, and written in Tapestry's format. The two systems are very similar along many dimensions.

By contrast, Tapestry is different from Post-Its™ on almost every characteristic. They share "hideable," but otherwise differ in every way. Contrary to our intuition going in, Notable is not a complete outlier in annotation systems: it falls between Tapestry and Post-Its™, and has similarities to a number of systems.

Framework application 2: Design

The analysis above can be used for design in a number of ways. The easiest way is to consider, for a particular characteristic of some system, what would happen if its value changed. For example, Notable has inactive annotations; what would it mean for Notable to have active annotations? Similarly, how could Notable support finer granularity than document level? Such questions are starting points for design.

Notable's granularity was one of the motivations for the analysis, and from the analysis we came up with a number of ideas on how to support a tighter annotation correspondence. An obvious idea was to render thumbnails of pages on the Pilot, but thinking in terms of levels of granularity suggested storing an outline or even the whole text of the document. Pixel-level granularity would be difficult to achieve on a 160x160 display, but character or paragraph-level granularity would be possible. Tapestry also had document-level granularity, but users there could get finer granularity using linguistic techniques, such as cutting and pasting text to be commented upon. Notable could support linguistic mechanisms as well, taking advantage of whatever representation of the document was available.

Another concrete design idea that came out of this exercise for us was the idea of adding a camera to the Palm Pilot in Notable, and letting the resulting picture be annotated on the Pilot. This change leverages Notable's unique ability to handle a wide range of annotation targets, including physical objects like

sculptures. This idea falls right out of trying to toggle an attribute value in Table 1: in this case, asking what it would take to make Notable's writing platform collocated.

MS Word is the only digital system that supports inline annotations. Would it be desirable for other systems to support them as well? How could that be done, for each system? Which systems could easily support voice annotations? Which could benefit from fluid UI techniques? These are the types of questions that this analysis helps to raise.

RELATED ANALYSIS

Various authors have outlined broad characterizations of annotations for various purposes. Marshall has studied annotation from an anthropological point of view [3], and proposed a number of annotation dimensions. Phelps and Wilensky designed a digital annotation system as part of their MVD work, and propose a number of desiderata for digital annotation systems [6]. O'Hara and Sellen [5] and Wojahn et al. [9] have studied people engaged in annotation tasks, and come to conclusions that inform the design of annotation systems.

CONCLUSIONS

Our perspective differs from those previously published because of its scope and focus on systems-oriented characteristics. Our framework can meaningfully distinguish a wide range of annotation systems, including the ten presented in this paper. In performing this analysis, we have incorporated Marshall's dimensions and Phelps' desiderata, but the taxonomy we present goes beyond those perspectives. The resulting analysis helps to compare systems, and to design new ones. It also provides a framework for comparing the results of annotation studies previously published.

ACKNOWLEDGEMENTS

We appreciate the support given to this project by Per-Kristian Halvorsen, Dave Marimont, Anthony LaMarca, and Eric Bier, and the hard work of Brian Lee, who wrote the first version of the Palm Pilot prototype.

We are also thankful to the Stanford Digital Library Project for sharing with us the Palm C++ libraries developed for that project.

REFERENCES

1. Chalfonte, B. L., Fish, R. S. and Kraut, R. E. Expressive richness: a comparison of speech and text as media for revision, in *Proceedings of CHI '91* (New Orleans LA, April-May 1991), 21-26.

2. Davis, R. C., Brotherton, J. A., Landay, J. A., Price, M. N. and Schilit, B. N. NotePals: lightweight note taking by the group, for the group. UC Berkeley Computer Science Division Technical Report UCB//CSD-98-997.
3. Marshall, C. Toward an ecology of hypertext annotation, in *Proceedings of ACM Hypertext '98* (Pittsburgh PA, June 1998), 40-49.
4. Nichols, D., Oki, B. M., Goldberg, D. and Terry, D. Using collaborative filtering to weave an information tapestry, *Communications of the ACM*, 35, 12 (1992), 61-70.
5. O'Hara, K. and Sellen, A. A comparison of reading paper and on-line documents, in *Proceedings of CHI '97* (Atlanta GA, March 1997), 335-342.
6. Phelps, T. A. and Wilensky, R. Multivalent annotations, in *Proceedings of the First European Conference on Research and Advanced Technology for Digital Libraries* (Pisa Italy, September 1997).
7. Röscheisen, M., Mogensen, C. and Winograd, T. Interaction design for shared World-Wide Web annotations, in *Proceedings of CHI '95* (Denver CO, May 1995), v. 2, 328-329.
8. Schilit, B. N., Golovchinsky, G. and Price, M. Beyond paper: supporting active reading with free form digital ink annotations, in *Proceedings of CHI '98* (Los Angeles CA, April 1998), 249-256.
9. Wojahn, P. G., Neuwirth, C. M. and Bullock, B. Effects of interfaces for annotation on communication in a collaborative task, in *Proceedings of CHI '98* (Los Angeles CA, April 1998), 456-463.
10. Zellweger, P., Chang, B., and MacKinlay, J. Fluid links for informed and incremental link transitions, in *Proceedings of Hypertext '98* (Pittsburgh PA, June 1998).