

PenLight: Combining a Mobile Projector and a Digital Pen for Dynamic Visual Overlay

Hyunyoung Song^{1,2}, Tovi Grossman¹, George Fitzmaurice¹, François Guimbretière^{2,3},
Azam Khan¹, Ramtin Attar¹, Gordon Kurtenbach¹

¹Autodesk Research
210 King St. East,
Toronto, ON, M5A 1J7 Canada
{firstname.lastname}@autodesk.com

²University of Maryland
Department of Computer Science
Human-Computer Interaction Lab
College Park, MD 20742 USA
{hsong, francois}@cs.umd.edu

³Cornell University
Computing and Information Science
301 College Ave,
Ithaca NY 14850 USA
francois@cs.cornell.edu

ABSTRACT

Digital pen systems, originally designed to digitize annotations made on physical paper, are evolving to permit a wider variety of applications. Although the type and quality of pen feedback (e.g., haptic, audio, and visual) have a huge impact on advancing the digital pen technology, dynamic visual feedback has yet to be fully investigated. In parallel, miniature projectors are an emerging technology with the potential to enhance visual feedback for small mobile computing devices. In this paper we present the PenLight system, which is a testbed to explore the interaction design space and its accompanying interaction techniques in a digital pen embedded with a spatially-aware miniature projector. Using our prototype, that simulates a miniature projection (via a standard video projector), we visually augment paper documents, giving the user immediate access to additional information and computational tools. We also show how virtual ink can be managed in single and multi-user environments to aid collaboration and data management. User evaluation with professional architects indicated promise of our proposed techniques and their potential utility in the paper-intensive domain of architecture.

ACM Classification Keywords

H5.2. [User Interfaces]: Input devices and strategies

General Terms:

Design, Human Factors

Keywords

Digital pen input, spatially-aware display, mobile projector, multi-layer interaction.

INTRODUCTION

In recent years, digital pens, that capture the ink strokes made on physical paper, have become widely available. These devices have revived the HCI community's interest

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2009, April 4–9, 2009, Boston, Massachusetts, USA.

Copyright 2009 ACM 978-1-60558-246-7/09/04...\$5.00.

in paper-based interfaces [3, 16, 18, 28, 36], as they combine the versatility and simplicity of paper [24, 32], but with digital enhancements. Such enhancements are not limited to the capture and recording of annotations, but can also be extended to support paper-based command systems. For example, Anoto [1] based applications allow users to interact with images of icons printed on the paper to provide computational results. Alternatively, the PapierCraft system [24] supports stroke-based commands to allow for active reading.

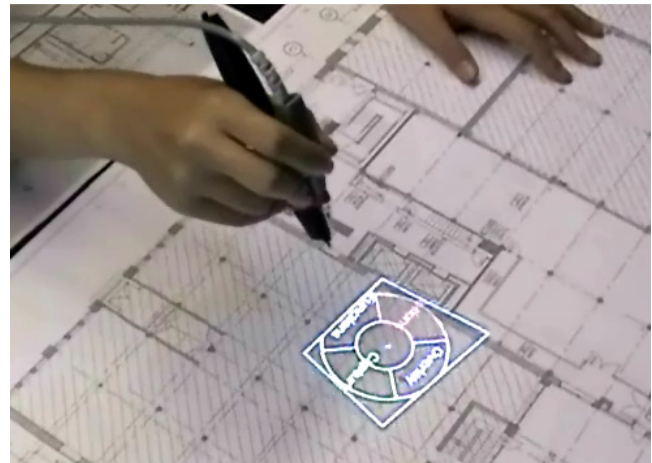


Figure 1. Our vision of the PenLight system.

A challenge with such systems is that while the pen provides the user with rich and dynamic input capabilities through the creation of ink and command strokes, current digital pen devices have very limited output capabilities. In its basic form, the user would receive no feedback at all. To address this issue, most digital pens have been enhanced with various forms of feedback, such as audio [13], haptic and visual feedback [25]. However, the visual feedback explored so far is limited to what can be displayed on the pen barrel itself, such as colored LEDs [23], or small OLED displays [25]. While such displays may be suitable for the most basic digital pen applications (e.g. querying simple text), it might be difficult to pursue more complex and intensive applications (e.g. searching for an word within a text) with such limited forms of visual feedback.

One potential solution to supporting rich applications on a digital pen is to mount a projector on the digital pen. To project a visual overlay in the context of a paper document [27, 36], the projector needs to be aware of its spatial location relative to the paper. Capturing the 3D location of the pen tip on or above the paper surface would allow the system to display virtual information which is relevant to the existing physical content on the paper, which may have been either printed or hand-written. A projection pen would thus increase the user's ability to work with functionality that requires visual feedback, such as viewing the results of computations, and overlaying contextual information.

A spatially-aware digital pen projection has yet been explored and will introduce new types of interactions and challenges. Currently, state-of-the-art miniature projectors are getting smaller, and soon will become small enough to be unobtrusively mounted on a pen. This also applies to the 3D optical tracking technology. To explore our vision of the projection pen, before these technologies are readily available, we implement a prototype of this configuration.

In this paper, we present PenLight (Figure 1), our proof-of-concept system for a digital pen embedded with a spatially-aware miniature projector. We present the interaction design space that the PenLight configuration introduces followed by a description of our high fidelity prototype implementation. Our system is implemented within an architectural application domain, chosen due to the significant use of paper throughout the current practices of the design, construction, and review phases. PenLight was used to conduct an informal user study with professional architects. Among several interaction techniques, overlaying building information on top of the blueprint and sharing annotations between remote users was most appreciated. Lastly, we present an analysis of possible challenges in building the actual setup of the projector pen.

RELATED WORK

Relevant areas of research associated with the PenLight include interactive paper, spatially aware displays, handheld projectors, and multi-layer interaction.

Interactive Paper with Digital Pens

The main goal of the PenLight system is similar to the goal of some previous systems [20, 27, 36]: to visually augment physical paper to enable virtual functionality.

The DigitalDesk [36] extends the computer workstation to include the affordances of a real desk such as tactile manipulation. PenLight takes the opposite approach to extend a physical pen to include the affordances of a workstation to be suitable for lightweight mobile system.

Paper-based systems, systems such as PapierCraft [24], ButterflyNet [38], PaperPoint [32] explored the use of a digital pen to directly interact with physical printouts. Digital operations presented in these systems capture and manage the annotations made on the document. Using the digital pen, users can query a limited amount of relevant

information using audio [13] or a nearby display [24, 32]. PenLight differs from previous systems in that query results are flexible in size and are projected in context of the paper.

Another group of paper-based research focuses on managing the link between paper and electronic content [3, 18, 28]. However, these systems explore an indirect link where input and output exist on separate devices. In contrast, PenLight examines a direct link between input and output.

The type of feedback provided by the digital pen [13, 25] plays a major role in diversifying the possible applications. Haptic vibration [22] and audio feedback [13] was provided by the first generation of digital pens. Liao et al. [23] presented a guideline describing how to combine color LEDs, tactile feedback, or audio feedback into coherent pen-top interfaces to improve the accuracy and the error rate. Recently, an 8 by 20 millimeter OLED display was embedded into the barrel of a digital pen [25] to display the result of the pen function. This has enabled commercial applications such as the display of a translated word. However, richer forms of visual feedback, which PenLight provides, have not been previously explored.

Spatially Aware Displays

PenLight utilizes a peephole display metaphor that has been used in earlier systems [7, 11, 28, 34, 37]. While some of these previous systems support navigation of virtual content by a fixed size viewing window, PenLight's viewing window dynamically changes, based on the location of the pen relative to the paper.

Yee [37] and Cao's work [7] explore pen input combined with a display in a bimanual setting to define input and output areas in the environment or to support a travelling input area. The PenLight system demonstrates different interaction techniques when the pen input and the display are integrated and used simultaneously.

Handheld Projectors

With the recent advancement in mini-projector technology, projectors are being embedded into a variety of handheld mobile devices such as cell phones and PDAs [26, 29]. To our knowledge, no one has previously explored the potential of augmenting digital pen applications with an on-board projector display.

Cao's work [7] looks into specifying projectable areas in the environment to create interactive information spaces. Similarly, PenLight explores different implicit information spaces, defined by the contents of the paper, both on and above its surface. Cao's multi-user scenario [8] also investigates how a hand-held projector can be used in a collocated multi-user scenario. PenLight manages pen input between remote users that share the same printout.

In handheld projector systems, the size and the resolution of the display also changes based on the proximity of the projector to the surface. Cao explores the granularity of visual content [7] at different distances. Zoom-and-Pick

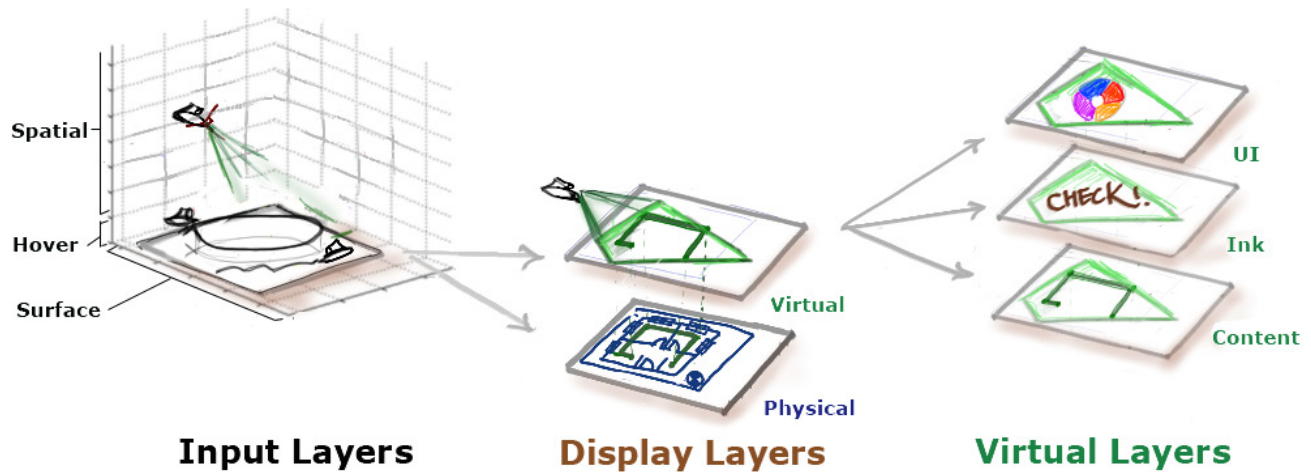


Figure 2. Layers in the PenLight system.

[14] uses proximity to improve selection accuracy. PenLight uses proximity information to control multi-scale widgets and to navigate virtual layers of information.

Multi-layer Interaction

The PenLight system interacts with multiple input layers above the paper surface. Multi-layer input interaction has previously been explored in devices such as tabletops [33], tablets [15] or pure virtual environments [6].

PenLight also explores the concept of *multivalent documents* [31] that consists of multiple abstract layers of distinct but closely coupled content. This concept is especially prevalent in the application domain that we are exploring. In architecture, building information modeling [9] comprises managing multiple data sets (different floor plans and section views with additional metadata to describe materials and processes) all intimately related to each other as part of a single virtual 3D model.

INTERACTION DESIGN SPACE

The idea behind PenLight is to provide richer visual feedback while interacting with paper. Our vision of the PenLight system consists of two components that will be available in the immediate future: (1) a pen sized projector and (2) a digital pen with 3D optical tracking.

This configuration opens up a unique interaction design space. We have partitioned the interaction design space into *input layers* and *display layers* (see Figure 2).

Input Layers

PenLight supports a spatial input layer in free space, a hover input layer just above the surface, and a surface input layer on the physical surface (Figure 2, left).

Spatial Input Layer

The spatial awareness of PenLight enables above-the-surface interaction [19]. As the PenLight is tracked in 3D space, we can consider a large spatial area above the paper as an input layer. The main use of this layer would be for command use and to position or rescale the projection.

Hover Input Layer

The hover layer is the layer above the surface, where the height information is less important for input. The primary use of this layer is for command input and manipulating the virtual cursor inside the projection area, as suggested previously [15].

Surface Input Layer

The surface layer is where the pen tip is in physical contact with the surface (typically the paper). We highlight two properties of the content of this layer: visibility and context.

Visibility: The visibility of the surface input layer indicates whether or not input within the layer will produce a visible trail of ink. With a standard physical pen, this input is visible. However, it may be desirable to provide input on the surface layer, without leaving a trail of ink. For example, when providing command input, an ink trail which was used for selection is of no use after the menu item is selected [24]. Also, it may be useful to support invisible ink annotations created on top of the original of a physical image, to avoid undesirable clutter, and to preserve the original.

Context: An important property of digital pens is that they are aware of the content that has been created, and of the pen location on the physical surface [24]. Thus, input created on the surface layer can be either high level global system commands, or contextual, acting on the data which is in proximity to the input.

Display Layers

Display layers consist of the physical surface layer and the virtual display layer (Figure 2, middle). The virtual layers can be overlaid user interfaces, ink, or data (Figure 2, right).

Physical display layer

The physical display layer is the layer which physically exists on the paper. This can consist of a number of different elements. There may be (1) *printed content*, such as a diagram or a 2D building layout, (2) *ink*, created by the user, and (3) *user interface elements*, such as menus and icons, preprinted onto the paper [13, 25].

Virtual Display Layer

Above the physical layer are virtual display layers that can be conveniently described in terms of *display elements*, and *display metaphor*.

Display Elements: We consider three categories of display elements which can be projected onto the virtual layer. Two traditional forms of display elements are the *user interface elements*, and the user generated data, or in the case of PenLight, *ink*. A third form of display element is *auxiliary data* relevant to the printout stored in other databases, which is not explicitly created by the user with a pen. Often, only a subset of associated virtual content is transferred to the physical printout during the printing process. This form of content could be useful for displaying aspects of the data which are not already shown on the physical display layer.

Display Metaphor: There are two metaphors we use for displaying virtual data: *content locked on-surface* (*peephole, overlaid*) and *content locked in-hand* (*default, displaced*). In the *content locked on-surface* metaphor, the *peephole* reveals a virtual world that is stationary relative to the physical world [11, 34, 37]. As PenLight is aware of the contents and location of the physical display layer, virtual data is directly *overlaid* in the context of the physical printout, *locked on-surface*. For example, ink annotations made by a remote collaborator could be positioned on top of the content which they are referring to, or virtual content which augments the physical content can be registered with the printed content and be overlaid.

In the *content locked in-hand* metaphor, imagery is projected without any calibration or transformation. As this metaphor uses the default projection style, it does not rely on the projector to be spatially aware. It can be used as an alternative display metaphor when tracking is likely to fail. Additionally, this metaphor is useful if the user wants to change the position or scale of the content as it moves. Such content could be *displaced*, or displayed indirectly.

PENLIGHT

Hardware Implementation

Currently, miniature projectors are small enough to be embedded into small gadgets such as cell phones [26, 35] and off-the-shelf digital pens [1] have internal cameras for 2D tracking. With today's camera technology, it would actually be possible to acquire the 3D location of the pen using the integrated pen camera, by analyzing the Anoto pattern on the paper, even when it is above the surface. However, combining these technologies in their current state would not produce a prototype suitable for use. Our current implementation of PenLight (Figure 3) simulates the unique configuration of a pen sized projector mounted on a digital pen before such hardware and technology is available.

Pen Input

For pen input, we use a Destiny IO2 Bluetooth digital pen (Figure 3 left). The digital pen allows the creation of physical ink and high resolution 2D tracking so that the system can store the created pen strokes. The 2D tracking is accomplished with a camera inside the pen that recognizes its location on the page and the page number, by reading a small high-resolution Anoto pattern which is physically printed on the page. A pressure-sensitive tip switch on the pen senses when the pen is in contact with the paper. A wireless Bluetooth connection links the pen with the CPU, so that the pen strokes can be stored virtually, and if desired, displayed, in real time.

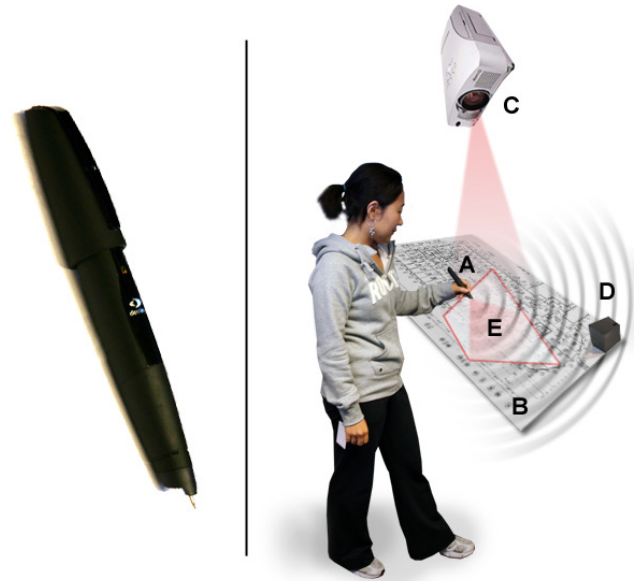


Figure 3. (left) Digital Pen. (right) System Setup: Our system was implemented using three components: (a) digital pen, (b) digital paper, (c) overhead projector, (d) 3D tracker, (e) projection image.

We chose to simulate the virtual surface input layer using a physical transparency, positioned under the pen with the non-dominant hand. The user can directly stroke on this surface without leaving an ink trail on the physical paper. For any of the interaction techniques which involve input on the actual surface, the user can choose to input them on this virtual surface input layer.

3D Tracking

To project imagery which overlays and augments the physical paper, the 3D location of the pen, relative to the paper, must be known. We acquire this 3D information by fixing a Polhemus FastTrak 3D magnetic tracker (Figure 3 right) onto the digital pen. The tracker is in the shape of a small pen, and senses full 6 degree-of-freedom information. Fixing a pen shaped tracker to the pen also gives us an initial understanding of how a pen shaped projector fixed to the pen would look and feel.

Projection Image

Instead of actually projecting from the pen, we use a top mounted projector (Mitsubishi XL4U, 1280x960 px), which projects downwards onto the paper (Figure 3 right). It is mounted 120 cm above the table projecting a maximum area of 90 cm by 66 cm. The 3D tracker is used to calculate the frustum for the simulated projector, as if it were mounted on the pen. The simulated location of the miniature projector is 1 cm above and 5 cm away from the pen tip on its front side. The simulated angle between the pen and the projector is 7° , and the field of view angle is 30° with an aspect ratio of 4/3. This configuration creates a 3.5 cm x 2.5 cm projected image when the pen tip is 5 cm above the display surface (Figure 4). The actual top-mounted projection image projects only into this simulated display region (Figure 3).

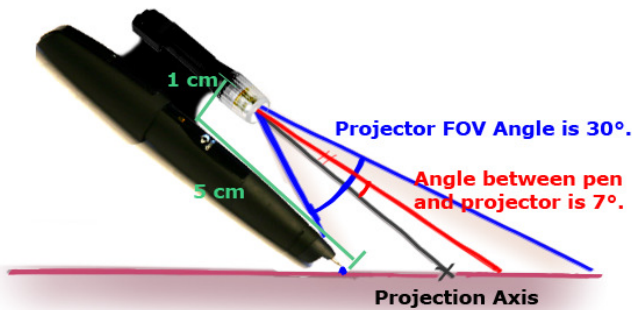


Figure 4. Axis and angles of pen and projector.

Software Application Domain

In improving the visual feedback provided by digital pens, we believe that PenLight will have several interesting usage scenarios for paper-based interaction. However, for the purpose of our explorations, we focus our implementation on a single application domain, allowing us to develop a working application supporting specific tasks. Many of the core concepts will easily generalize to other domains.

The architecture profession has one of the most paper intensive workflows as paper is the common medium to distribute designs among different parties and it represents the actual contract commitment (Figure 5). We consulted a practicing architect to discuss how paper is frequently used in the architecture domain and the potential practices in architecture for which the PenLight system could be useful.

While paper drawings are ubiquitous in each stage of architecture practice, they have limited capabilities. In particular: (1) it is difficult to access additional information related to the printout. During a discussion between architects and their clients in a meeting room, it is often the case that customers want to see a 3D rendering of the design. This normally requires a computer nearby and real-time applications to simulate the walkthrough. (2) Levels of detail are spread across many different drawings; manually tracing one layer of information and overlaying it on top of another printout is a common practice that architects use for this problem. (3) It is difficult to coordinate different

versions of a document as well as between remote collaborators.

The above mentioned problems have a close relationship to the management of multiple layers of data and input which we discussed in the design space section above. We thus chose to implement an architectural application which allows users to query and augment physical architectural sketches, addressing the limitations of the current practices.



Figure 5. Use of paper in architecture and construction.

PENLIGHT INTERACTION TECHNIQUES

The multiple input and display layers which PenLight introduces bring forth new interaction techniques that, in combination, have not been explored in previous digital pen interfaces. Our system allows users to navigate among different virtual ink and content layers, perform operations on physical and virtual content, extract and display different representations of the printed content, and access functionality through a menu system.

Menu Design

We designed a hierarchical radial menu system which can be used to access the various functionality of the system. The radial distribution of menu items simplifies its use, since users only need to remember what direction to move.

Users can access the menu system by clicking the barrel button on the digital pen. This causes the top level of the menu to be projected (Figure 6). Displaying the menu on the virtual layer addresses one problem with current digital pen menu systems – they cannot be displayed to the user, unless they are preprinted on every page. Another problem which we want to address is that physical ink marks created from command activations result in undesirable clutter.

We present three menu selection techniques which explores different 1) input layers, 2) display metaphors of virtual display layers and 3) reliance on visual feedback.

In the “crossy” [2] menu, once it is activated by pressing the button, its position is locked on-surface using a peephole *display metaphor* the entire time, and then pen motion controls a virtual cursor in the *hover input layer*

which begins in the center of the menu (Figure 6a). A menu item is selected when the virtual cursor crosses the radial menu's outer border. The menu remains locked and the process is repeated on the next level if the selected item is hierarchical.

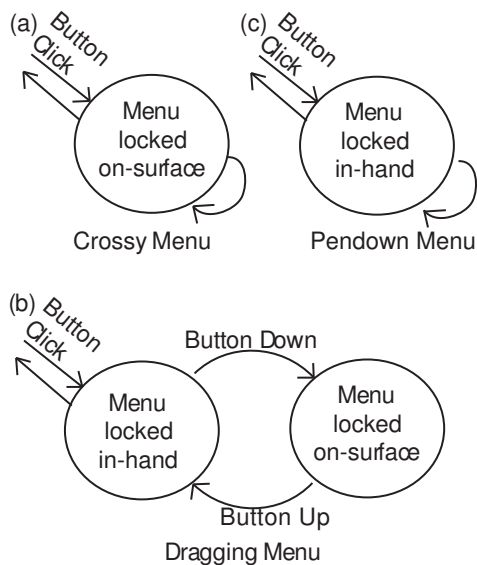


Figure 6. State transition diagram for menu interaction. a) Crossy menu. b) Dragging menu. c) Pendown menu.

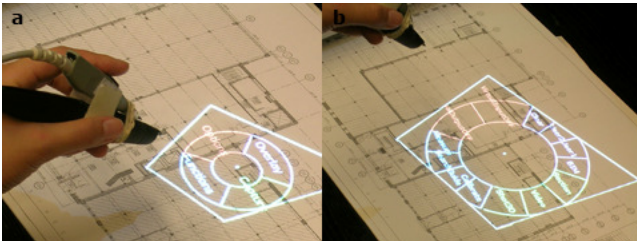


Figure 7. Lifting the pen reveals the next level of the menu.

The second “dragging” technique utilizes both display metaphors (locked on-surface and locked in-hand) in the *hover input layer*. Once the menu is activated, both menu and virtual cursor is locked in-hand and menu items cannot be selected because the virtual cursor remains in the center of the menu (Figure 6b). To lock the menu to the surface, the user holds the button down, makes a mark in the appropriate direction, and then releases the button. If the menu is hierarchical, the next level of the menu would then be displayed and the process is repeated. This technique would be appropriate if the user only wants to find their desired menu item while the pen is stationary, but then make their selection without having the menu displayed.

The third “pendown” technique is similar to the “dragging” technique, but the marks are made on the *surface input layer* (Figure 6c). Unlike the previous two techniques this could leave an ink trail from the menu use. If the user did not want this to occur, the mark could be made on the virtual surface input layer, by using the physical transparency under the input location.

Since PenLight has a *spatial input layer*, we explored using the height information to control the semantic scale of the menu. When the user lifts the pen above the hover layer, two levels of menu items are shown around the ring, allowing the user to see more items at once (Figure 7). Although the menu items are bigger, the motor space is smaller [39], making them difficult to select. This technique is similar to previously developed multi-scale widgets [30].

Ink Management

The most basic functionality of digital pens is creating and managing ink. In PenLight, creating physical ink is not different from sketching in the physical realm with pen and paper. In addition to the physical ink, PenLight allows users to create and manage virtual ink that users can make use in different functions: tracing, and virtual guides.

Virtual Ink

An ideal hardware implementation of enabling virtual ink would be to use a mechanical button that would change to a pen tip with no physical ink. We use a transparency instead, so the user has to select a menu item to enable the virtual ink input when using the transparency. When enabled, all strokes are added to the virtual ink layers, in the location of the paper which they are created. By creating the strokes in the virtual surface input layer, the annotations can be added to only the virtual layer. This allows a user to annotate a blueprint without altering the original document.

Tracing

Users can trace over both physical and virtual content and then apply the trace data to different spatial locations. Users can also load existing virtual templates to trace out with physical ink input. Tracing is different from previous guided sketching projects [12], as PenLight requires users to rely on limited field of view that changes its resolution and size depending on the location of the input device.

Virtual Guides

Instead of tracing, virtual guides can be created to aid a physical sketch. Such grids and guides are widely used in image editing applications, but unavailable when working on physical paper. To create a geometric guide, the user can select one of the menu items; line, circle, rectangle, or grid. Instead of entering points that define the geometry, the user can draw a similar shape and the system approximates the selected shape, similar to Arvo's approach [4]. For example, the user can draw a circle and the system figures out the center point and the radius. In grid mode, users can draw a rectangle that serves as the unit rectangle shape of the grid. Once the pen is lifted, the entire virtual layer is covered with a self replicating grid layout.

Overlaid Content

One of the main benefits of PenLight is being able to present overlaid content. This can, for example, be an important operation when working with physical blueprints. In the architecture domain, managing the various aspects of building data in a single 3D model is a recent trend called Building Information Modeling [9]. PenLight stores these

various layers of virtual content which can be overlaid onto the physical image. In our application, the users can select to view the ventilation, pipeline, or lighting layer, which is overlaid on top of the physical image (Figure 8).

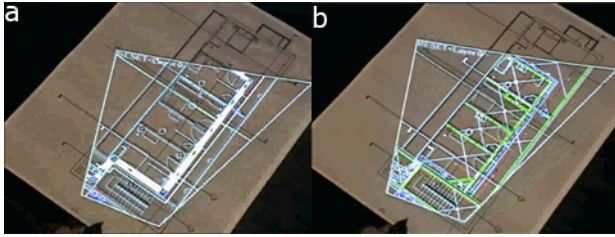


Figure 8. Overlaid Content: On top of the floor plan, different types of information are overlaid: (a) heating, ventilation and air conditioning; (b) mechanical, electrical pipeline.

Copy and Paste

The overlaid content or the original physical content can be copied to another location to be overlaid. The user enters a copying mode from the menu, and circles an area using the pen to specify a contextual parameter of the image on the surface. The user then enters a pasting mode from the menu, and the copied content is displayed using the *locked-in hand* metaphor, and copied when the user clicks the button.

Overlaid Computations

The virtual display layer of PenLight allows digital pens to carry out computations, and directly display the results of those computations in the context of the user's workspace.

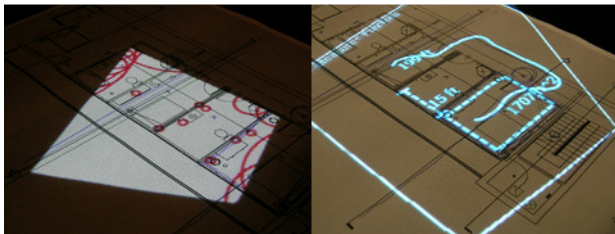


Figure 9. Search and Measurement: (left) search results highlighted using Halo. (right) measurement query output (straight line, path, area).

Search

The search command allows users to search for an item that exists on the physical display layer. The user can perform the query in two ways. They can choose from a list of query objects in the search menu using a virtual cursor, such as sprinklers, or they can directly circle an instance of an object on the printout, such as a power outlet using the pen tip. Once the query is performed, all instances of the objects are highlighted in the virtual display layer. Instead of having users performing a linear search [24], we use the Halo technique to guide the user to instances of the object [5] (Figure 9, left). Users can raise the pen to see a larger portion of the display to navigate to the object, so that the items of interest can be found faster.

Dimensions

The dimension tool is another tool which overlays the result of a computation on the physical workspace. Using the menu the user can choose to measure a distance, path length, area, or volume. The user then makes an appropriate stroke on the desired location of the image, and the computation result is overlaid on top of the paper as part of the virtual display layer (Figure 9, right).

Creating Alternate Views

In PenLight, the printed content on the paper is actually only one abstract view of a larger electronic file that is stored in the system. For example, when a 2D floor plan is printed out on paper, the digital pen could directly store the highly detailed 3D model. This type of imagery could be displayed, possibly in an empty area of the page, or on a nearby blank surface.

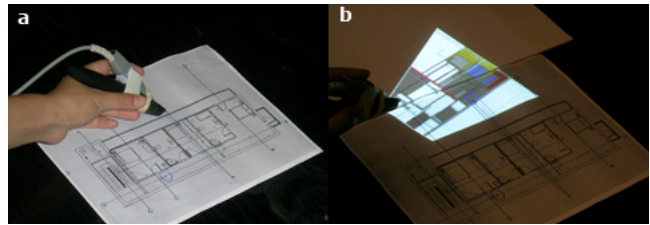


Figure 10. 2D Section View: (a) An ink trail along the section line determines the contextual input. (b) The result can be displaced.

2D Section View

When in 2D section view mode, the user draws a line on the paper to define a cutting surface to extract a 2D section of the current 3D building based on the position and orientation of the line (Figure 10). The temporary view is locked-in-hand and can be dropped on-surface when the pen button is clicked.

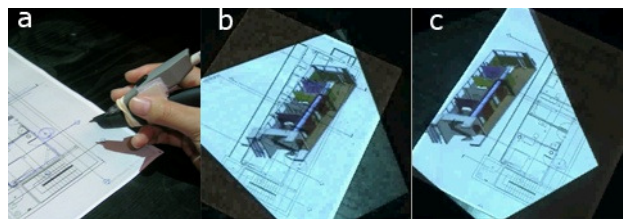


Figure 11. 3D Snap Shot: (a) The angle and position of the pen defines the location of the camera. (b) The 3D view is being displaced (c) The 3D view is locked in a convenient spot.

3D Snap Shot

By using the *spatial input layer* of PenLight, users can extract a 3D snap shot of the model. When choosing this operation, the user can use the location and direction of the pen in reference to the paper to specify the camera location and the viewing vector into the digital model. Varying the pen height determines what building view is to be captured: the interior view (when the pen is near the physical surface) or exterior view (when the pen is high above the physical surface). As with the section view, the 3D snapshot can then be displaced and locked on-surface nearby (Figure 11).

2D Walkthrough

A similar operation can be used to create a 2D walkthrough of the building. When using this operation, the user can draw a path through and along the floor plan (The current implementation only supports several predefined paths). When the pen-up event is detected, the system locks the video under the pen (in-hand), and clicking the barrel button triggers the system to play the video and lock its location on-surface. As the video is being played, a red marker dynamically moves along the path which indicates the current location of the video (Figure 12).

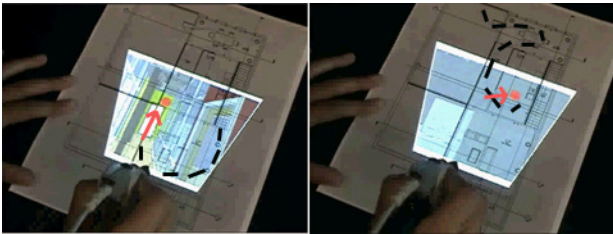


Figure 12. Walkthrough Video: After the user draws the walkthrough path, a video is rendered and projected. A red dot shows the user the location of camera for the rendering.

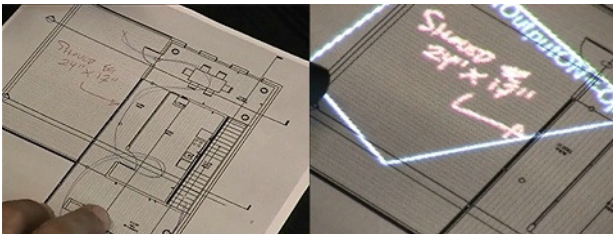


Figure 13. Remote Collaboration: While one user is writing on their floor plan, the annotation can be transferred to a remote user's digital pen and hence projected for synchronization.

Remote Collaboration

The system supports illustrative communication between remote collaborators, such as a designer and a fabricator. We briefly explored this scenario by introducing a second Anoto pen to our system, and printing out a second copy of a floor plan. When annotations are made by one user with the standard Anoto pen, they can be displayed in real time as virtual ink by the user of the PenLight system (Figure 13). Annotations from the remote user are displayed in a different color. We implemented this locally only, but in practice it could be implemented for remote scenarios by connecting the pens over a network.

EXPERT INTERVIEWS

PenLight was demonstrated to three professional architects to assess the usefulness and the potential of each interaction technique. The first 15 minutes was used as a demonstration session to provide the participants with an overall understanding of the PenLight system and its functionality. The next 45 minutes was used as a semi-structured interview. During this interview, participants were asked to comment on the features of our system, including the applicability of each feature to their own everyday practices. In general, their feedback validated that our

design choices would apply well with paper intensive practices such as the architecture field.

The most positive response received during these sessions was for overlaying additional building information on top of the blueprint. The architects felt that this tool would be extremely useful and easy to use. Furthermore, all three architects also felt that the ability to capture and subsequently display a user's annotations to a second user could be useful, as miscommunications due to the absence of such abilities in current practices end up increasing the time and cost to complete a project.

One of the architects also suggested the support for "consistency checks". Such functions inform the user of potential problems when a layout is modified, such as inserting a new pipeline. It would be useful if these consistency checks could be performed at an earlier stage of the design process, taking place on the paper while it is being marked up.

Overall, participants liked the various interaction techniques that were implemented. Searching for an item of interest in a large-sized blueprint was mentioned as being "priceless". However, the participants did not see an immediate benefit of the dynamic measurement computations for their work activity.

Participants also commented on the configuration of the hardware. One issue discussed was the location of the simulated projector. Users were satisfied with the location and size of the projected image, and liked the ability to raise the pen to view a larger area of the virtual image. They especially liked this type of interaction when performing a search to get context. However, they did comment that sometimes the size of a blueprint can be quite large (e.g. A0) and the current field of view might not be wide enough to obtain a full overview, regardless of how high the pen is.

DISCUSSION

Here we discuss some of the issues relating to our implementation and a future vision for the PenLight system.

3D Tracking

Our assumption is that pen size projectors will emerge in the near future and high accuracy 3D tracking will be made possible. However, the stability that we were able to achieve using an overhead projector may not be immediately replicable with a pen-mounted projector. Here we discuss how the tracking can be improved in hardware components, and in software techniques.

Improving optical pattern tracking

Today's Anoto technology only provides 2D location information when within 0.5 cm of the paper. However, there are other tracking solutions to improve long range, 3D optical tracking. Traceable patterns can be added to retrieve camera calibration parameters [17], similar to ARTags to detect 3D location and orientation. Another approach will be to change the pattern to a hierarchical encoding [21], which will allow the camera to cover a wide range of

depths over the surface of the paper. Additional patterns can be printed in infrared ink to be less distracting. Given these existing research results, it is reasonable to assume that a pen-integrated camera-based tracking solution will be available in the future.

Improving Image Stability

Previous spatially-aware projection systems have provided a continuous virtual display, projecting imagery at all times, regardless of the movement of the device [7, 14]. There is a significant technical barrier to keeping the virtual image stable if it is meant to be displayed with a peephole metaphor. With a pen, the problem would only be exasperated, due to its high frequency of movement.

There are hardware solutions that can alleviate this issue. One such technique is image stabilization, which is a feature of many commercial cameras.

In the interaction design space section, "Locked-in-hand" projection (displaying content that does not require spatial information) is a solution that we already make use of in our interaction techniques. Another alternative interaction paradigm is a "discrete display mode" which only projects imagery at discrete intervals, when the pen is in a relatively stable location. Once the pen begins moving faster than a threshold value, the imagery would fade out. This introduces a unique interaction style, where the user may be able to see the virtual imagery when viewing it, but have to rely on their persistence of vision [10] to interact with it.

Projection Image

Projector Location

The location of the miniature projector must be carefully considered, as it has a number of implications. The location of the projector on the pen determines the size of the projected image and the pen's center of mass. Furthermore, the angle of the projector will determine where the tip of the pen is in reference to the projected image. This is an important consideration for any technique which requires the user to rely on visual persistence to interact with virtual imagery, such as tracing. The angle of the projector could also determine if any "finger shadows" will exist on the projected image. One of the participants in our interviews commented that the task may have been easier if the display size was bigger. Mounting the projector with a wider angle lens or a redirection mirror may assist this issue.

Dynamic resolution and brightness

Hand-held projectors provide a dynamic resolution and brightness. In terms of dynamic resolution, focus will be an issue for a lens based projector. For this problem, a laser based projector will keep the image in constant focus at different distances. The dynamic brightness could also be accommodated, using a projector that modulates the brightness based on its distance and rendering software that takes the dynamic dpi into account.

When an actual pen-mounted projector is close to the paper, the resolution will be higher than our simulator, making it

possible to display more details. Hence, PenLight's UI elements and projected content will naturally transfer to the actual setting. In general, we intend to explore how dynamic resolution and brightness would impact our explored interaction paradigms in the future.

Different Hardware Configurations

PenLight simulates a miniature integrated projector, instead of having a separate pen and projector. This decision was made with mobile usage scenarios in mind, where fewer and more lightweight hardware components are preferred. Furthermore, there are interactions that are not possible with a separate projector and a pen. For example, a pen mounted projector introduces a dynamic display area, which is useful in selecting and moving virtually overlaid content. This large dynamic display area with varying resolution can be used to display different focus+context information [7].

However, a separate projector configuration, such as a paper-mounted projector or even a removable "pen cap projector", would be interesting to explore and compare to the current configuration.

CONCLUSIONS & FUTURE WORK

In this paper, we have initiated the exploration of augmenting digital pens with miniature spatially-aware projectors, and defined and explored the main aspects of a design space that this introduces.

Novel aspects of this design space can be narrowed down to three items. First, ink no longer has to be represented in a physical form. Virtual ink benefits users in many ways. For instance, users can get visual feedback without permanently modifying the physical surface, and virtual strokes can be used to communicate with a remote user. Second, we showed that the interaction space is not merely locked to the surface input layer but extends to the space above the paper. Third, a spatially-aware pen and projector allows a user to visibly correlate information that is stored inside the pen or on any connected resource with the document. As a result, paper is no longer just a static source of data, but it can be used as a dynamic workspace. In essence, PenLight illuminates information that was hidden due to the static nature of physical paper, just as a traditional penlight lights up unseen parts of a document in the dark.

An obvious line of future work is the development of a working prototype with the projector mounted on the digital pen. The smallest miniature projectors developed to date are almost adequate for such a prototype. Significant issues remain to be researched including: providing mobile 3D location sensing; providing projector power; continued miniaturizing of pen computation and mass storage; ergonomic considerations of the pen shape; and, technical issues covered in the discussion section.

ACKNOWLEDGMENTS

This research was supported in part by the National Science Foundation under Grants IIS-0447703, IIS-0749094 and by a gift of Autodesk Research to the University of Maryland

REFERENCES

1. Anoto, Development Guide for Service Enabled by Anoto Functionality. 2002, Anoto.
2. Apitz, G. and F. Guimbretiere. CrossY: A Crossing-Based Drawing Application. *ACM UIST '04*, p. 3 - 12.
3. Arai, T., D. Aust, and S.E. Hudson. PaperLink: a technique for hyperlinking from real paper to electronic content. *ACM CHI '97*, p. 327 - 334.
4. Arvo, J. and K. Novins. Fluid sketches: continuous recognition and morphing of simple hand-drawn shapes. *ACM UIST '00*, p. 73 - 80.
5. Baudisch, P. and R. Rosenholtz. Halo: a technique for visualizing off-screen objects. *ACM CHI'03*, p. 481-488.
6. Bier, E.A., M.C. Stone, K. Pier, W. Buxton, and T.D. DeRose. Toolglass and magic lenses: the see-through interface. *ACM SIGGRAPH '93*, p. 73 - 80.
7. Cao, X. and R. Balakrishnan. Interacting with dynamically defined information spaces using a handheld projector and a pen. *ACM UIST '06*, p. 225-234.
8. Cao, X., C. Forlines, and R. Balakrishnan, Multi-user interaction using handheld projectors. *ACM UIST'07*, p. 43 - 52.
9. Eastman, C., P. Teicholz, R. Sacks, and K. Liston, *BIM Handbook: A Guide to Building Information Modeling for Owners, Managers, Designers, Engineers and Contractors*. 2008: Wiley.
10. Erwin, D.E., Further Evidence for Two Components in Visual Persistence. *Journal of Experimental Psychology: Human Perception and Performance*, 2(2): p. 191 - 209, 1976.
11. Fitzmaurice, G.W., Situated information spaces and spatially aware palmtop computers. *Communications of the ACM*, 1993. 36(7): p. 39 - 49.
12. Flagg, M. and J.M. Rehg, Projector-guided painting. *ACM UIST'06*, p. 235-244.
13. Flypentop Computer. <http://www.flypentop.com/>
14. Forlines, C., R. Balakrishnan, P. Beardsley, J.v. Baar, and R. Raskar. Zoom-and-pick: facilitating visual zooming and precision pointing with interactive handheld projectors. *ACM UIST '05*. p. 73 - 82.
15. Grossman, T., K. Hinckley, P. Baudisch, M. Agrawala, and R. Balakrishnan. Hover widgets: using the tracking state to extend the capabilities of pen-operated devices. *ACM CHI '06*, p. 861-870.
16. Guimbretiere, F. Paper Augmented Digital Documents. *ACM UIST '03*, p. 51 - 60.
17. Heikkila, J. and O. Silven. A Four-step Camera Calibration Procedure with Implicit Image Correction. *IEEE CVPR '97*, p. 1106 - 1112.
18. Heiner, J.M., S.E. Hudson, and K. Tanaka. Linking and messaging from real paper in the Paper PDA. *ACM UIST'99*, p. 179 - 186.
19. Hinckley, K., R. Pausch, J.C. Goble, and N.F. Kassell. A survey of design issues in spatial input. *ACM UIST '94*, p. 213 - 222.
20. Holman, D., R. Vertegaal, M. Altosaar, N. Troje, and D. Johns. Paper windows: interaction techniques for digital paper. *ACM CHI '05*, p. 591-599.
21. Keisuke, T., K. Itaru, and O. Yuichi. A nested marker for augmented reality, *ACM SIGGRAPH '06*.
22. Lee, C.J., H.P. Dietz, D. Leigh, S.W. Yerazunis, and E.S. Hudson. Haptic pen: a tactile feedback stylus for touch screens. *ACM UIST '04*, p. 291 - 294.
23. Liao, C., F. Guimbretière, and C.E. Loeckenhoff. Pen-top feedback for paper-based interfaces. *ACM UIST '06*, p. 201 - 210.
24. Liao, C., F. Guimbretière, K. Hinckley, and J. Hollan. Papiercraft: A gesture-based command system for interactive paper. *ACM Transactions on Computer Human Interaction*, 2008, 14(4): p. 1-27.
25. LiveScribe. <http://www.livescribe.com/>
26. Light Blue Optics. <http://www.lightblueoptics.com/>
27. Mackay, W.E., D.S. Pagani, L. Faber, B. Inwood, P. Launiainen, L. Brenta, and V. Pouzol. Ariel: augmenting paper engineering drawings. *ACM CHI '95*, p. 421 - 422.
28. Mackay, W.E., G. Pothier, C. Letondal, K. Bøegh, and H.E. Sørensen. The missing link: augmenting biology laboratory notebooks. *ACM UIST '02*, p. 41 - 50.
29. Norman, D., *Psychology of everyday things*. 1988: Basic Books.
30. Perlin, K. and J. Meyer. Nested user interface components. *ACM UIST '99*, p. 11 - 18.
31. Phelps, T.A. and R. Wilensky. Multivalent Annotations. *Proceedings of ECDL '97*, p. 287 - 303.
32. Signer, B. and M.C. Norrie. PaperPoint: a paper-based presentation and interactive paper prototyping tool. *Proceedings of TEI '07*, p. 57-64.
33. Subramanian, S., D. Aliakseyeu, and A. Lucero, Multi-layer interaction for digital tables. *ACM UIST '06*, p. 269 - 272.
34. Tsang, M., G. Fitzmaurice, G. Kurtenbach, A. Khan, and B. Buxton. Boom chameleon: simultaneous capture of 3D viewpoint, voice and gesture annotations on a spatially-aware display. *ACM UIST '02*, p. 111 - 120.
35. MicroVision <http://www.microvision.com/>
36. Wellner, P., Interacting with paper on the DigitalDesk. *Communications of the ACM*, 1993. 36(7): p. 87 - 96.
37. Yee, K. Peephole displays: pen interaction on spatially aware handheld computers. *ACM CHI '03*, p. 1 - 8.
38. Yeh, R.B., C. Liao, S.R. Klemmer, F. Guimbretière, B. Lee, B. Kakaradov, J. Stamberger, and A. Paepcke. ButterflyNet: A Mobile Capture and Access System for Field Biology Research. *ACM CHI'06*, p. 571 - 580.
39. Zhai, S., S. Conversy, M. Beaudouin-Lafon, and Y. Guiard, Human on-line response to target expansion. *ACM CHI '03*, p. 177 - 184.