Speakeasy: A Platform for Interactive Public Displays

Julie A. Black[1], Jason I. Hong[2], Mark W. Newman[3],
W. Keith Edwards, Shahram Izadi, Jana Z. Sedivy, Trevor F Smith

[1] Stanford University
Computer Science Department
Stanford, CA 94309
jblack@cs.stanford.edu

[2] Group for User Interface Research
University of California, Berkeley
Berkeley, CA 94720-1776
jasonh@cs.berkeley.edu

[3] Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA, 94309
mnewman@parc.com

ABSTRACT
This is our vision: display and interact with any media type on any kind of device that you want to turn into an electronic public display. We desire to support all media types and devices that exist today as well as those that will exist tomorrow. To explore this vision, we developed two electronic public displays using Speakeasy, a foundational platform that supports ad hoc communication between devices and services with little or no *a priori* knowledge of one another. We discuss the design, implementation, and our experiences with these two electronic public displays.

Keywords
Public display, Speakeasy, content, content management

INTRODUCTION
Consider the following scenario. A group of researchers is enjoying lunch together in a public cafeteria. During the course of the meal, one of the researchers, Paul, notices a flyer posted near their table that advertises a corporate softball game. He mentions the flier to the other members of the lunch group and they all appear interested in joining one of the teams. Paul takes note of the game's date and time and plans to email this information to the others later.

As the conversation continues, another member of the group, Suzie, remembers a funny picture from the last corporate softball game that she would like to share with the rest of the group. Because no computer is directly accessible from the lunch table, Suzie will have to email the picture to everyone at the table later. After lunch, *if* Paul and Suzie remember to send their messages to the group, they must also remember *who* they ate lunch with today.

We interact with many kinds of non-electronic public displays on a daily basis, such as posters, calendars, bulletin boards, and flyers. People often rely on non-electronic displays to provide information about upcoming events and to share different media with other members of their community. For example, one person might post recent vacation photographs while another might share a funny news clipping or, as in our scenario, someone might post an advertisement or flier.

However, non-electronic public displays have several shortcomings. First, content cannot be added or modified remotely; in order to show her softball photograph on a non-electronic public display such as a bulletin board, Suzie would have to walk to the display itself and post a hardcopy. Similarly, it can also be difficult to update content; if there is a change in the softball game's date and time the poster will have to create entirely new fliers and manually replace all of the previously posted ones. Second, the display cannot be viewed remotely. If an individual sitting in his office remembers that he saw an interesting flyer on the cafeteria bulletin board, he has to walk from his office to the cafeteria to view the posting. Third, non-electronic public displays provide limited on-site interaction with postings. If an individual would like more information about an upcoming event, she has to write down the contact information and remember to email the organizer later. Finally, because of the physical size limitations of non-electronic public displays, there is a maximum number of items which can be posted at any

Figure 1: Two of the authors interacting with a mockup of the Speakeasy public display.

given time. In the case of a bulletin board, posters only add content where there are no pre-existing posts.

Many researchers have been looking to electronic public displays as a solution to these problems. With electronic displays, members of a community can share many different types of media easily, effectively, and remotely. Like non-electronic displays, electronic public displays act as a portal between the individual viewing the content and the community of individuals who post information.

However, electronic public displays introduce two new problems. First, it is difficult to add arbitrary new content to public displays. For example, a public display might know how to show MPEG and AVI movie files, but it would be difficult to deploy an electronic public display that knows how to play *all* kinds of movie file formats, especially those not yet invented. Second, it is difficult to make arbitrary devices into public displays. For example, even if one device you want to turn into a public display has all of the software it needs, it is difficult to have *all* of the devices have all of the right software.

Our first steps in exploring our vision of *ad hoc* interaction with any kind of display device showing any media type was to develop two different electronic public displays that can show different kinds of media. One of these is shown in Figure 1. We developed these displays using Speakeasy, a platform that allows end-users to connect arbitrary services and devices together in an *ad hoc* fashion. In the rest of this paper, we discuss the design, implementation, and our experiences with these displays.

RELATED WORK

The Notification Collage [4] provides an instructive look at the integration of a display system into an office setting. In this work, researchers have selected a limited number of acceptable interactions such as viewing a slideshow of photographs, a web page, and live streaming video, and have focused on the layout of viewable information upon the public display. Because such display composition has been restricted to a small set of media types and each type has been specially integrated into the system, supporting new data types would seemingly require moderate programming effort. Because our work leverages the Speakeasy framework (see following section for added detail), new data types and interaction mechanisms can be added more easily to the pre-existing display system.

Additionally, the Campiello project [1] provides a glimpse into the use of public displays as a community oriented means to share historically significant information. Schuler's work on community networks [6] provides insight into possible uses of such a community display system and the technical and social constraints placed on such a system as a result.

The aforementioned research projects investigated issues such as content layout and social implications of public displays within communities. In contrast, the primary focus of our work is to examine how to successfully integrate support for communication between a varied, ever-growing set of devices and components.

SPEAKEASY SYSTEM OVERVIEW

The need to support interactions among a wide variety of devices and services led us to build our public display system on top of the Speakeasy framework for recombinant computing [2, 3]. Speakeasy is a mobile code-based software architecture that allows arbitrary devices and services (collectively referred to as *components*) to interoperate with very little *a priori* knowledge of one another. Components that provide data or a service are known as *sources*. Components that receive data, whether to render, store, print, or perform other tasks, are as known as *sinks*. The primary advantage of using Speakeasy for this project was that we were able to quickly and easily integrate data from a wide variety of different sources—such as web pages, files of various data types, and streaming media—into our display without having to write protocols to deal with each source type independently. In other words, once a new sink was created, it could be easily connected to all of the sources that already exist.

Additionally, Speakeasy defers the semantics of connecting components to users. For example, the computer does not know what a digital camera or printer is, but Speakeasy allows the end-user to connect the two together and allows the user to determine what that connection means and how it should be controlled and managed. Because Speakeasy leaves the semantics of connections to users, new components can be introduced to the Speakeasy network without introducing additional communication protocols. Thus, for the public display system, Speakeasy allows users to determine which components should be connected to the display, how they should be connected, and how such connections should be managed.

WORK TO DATE

We are in the process of deploying a modular display system for use in the Computer Science Lab at PARC. To date, our work is in an early phase with emphasis on integrating the variety of Speakeasy sources and their associated data types into an extensible, modular public display system that can be evolved as we learn more about the community's desired applications for such a display system. This section describes in more detail our work with two different public displays, a text-based tickertape LED sign and a large format graphical display.

Text-Based Tickertape LED Sign

As shown in figure 2, the text-based tickertape sign was built using a simple tri-color one line LED display.[1] The specific sign that we chose had a display area of 27 inches (width) by 2.1 inches (height), could display approximately 15 characters a time, and could store approximately 7,000 characters in memory. By default, messages displayed on this LED sign would scroll from right to left and color

[1] We used the Alpha 215 Sign built by Adaptive Micro Systems.

Figure 2: The Speakeasy tickertape LED sign scrolling a Slashdot headline.

combinations would be randomly selected (combinations of red, green, and orange).

Members of the community could add content to this display via an online Java applet or through the Speakeasy network. Additionally, each poster could specify a lifetime for their message, after which, the message would be removed from the queue of displayable content. We also added a RSS news feed component to the Speakeasy network which parsed and streamed news updates from web sites such as Slashdot and NewsForge to the LED sign. Added functionality enabled the sign to export a list of all messages currently stored in the display queue.

After the sign was integrated into the Speakeasy network, it was hung in a common area and community members were notified that they could add content. Not surprisingly, most messages that users chose to add were comical in nature; messages included Jenny Holtzer's truisms, quotations from Star Wars, and musings at individual researchers. Several viewers identified the Slashdot and NewsForge streams as particularly useful and desirable.

This sign was relatively easy to develop because of the underlying Speakeasy platform. Supporting Speakeasy enhanced networked communication with the sign only required two new pieces of code: one for communicating between Speakeasy sources and the message queue of content, and another for communicating between the message queue and the sign. We did not have to build any backend support for the direct interaction between components because Speakeasy provides real-time component discovery and component connection mechanisms. Therefore, development of this sign clearly illustrated that Speakeasy would provide a sturdy, yet flexible foundation for a graphical display system. Because it was also evident that community members would desire an alternative means by which they might add content to displays without installing Speakeasy software, we were motivated to deploy a web applet that allowed anyone with access to a suitable web browser to add content to the sign.

Graphical Display System

One problem with the LED sign is that it only allowed text-based content, which overlooks the main advantage of building on top of Speakeasy–the ability to take advantage of heterogeneous sources and data types. To address this problem, we began the deployment of a graphical display system in public and private areas throughout our lab.

As is illustrated by the schematic diagram in figure 3, this system consists of three major pieces: an unstructured "media soup" of data that community members can add data to or retrieve data from; a set of public displays, located in common areas such as break rooms, hallways, printer rooms, and meeting rooms; and a set of "private-public" displays located in individual offices. Such "private-public" displays act as a user's portal to this display system providing a personalized view of the shared data. Additionally, in order to support spontaneous information sharing (such as Suzie sharing pictures with her lunch group), our public displays can be "taken over" for temporary interactive use. Users can access the media soup directly (such as following a link in an article, contacting the poster of something interesting, or examining a post more closely) or perform focused tasks such as migrating information from small to large displays, or accessing other media shared on the Speakeasy network. Similarly "private-public" displays, which are normally dedicated to displaying personalized versions of the soup contents, can be used for tasks such as videoconferencing with another community member.

The "media soup" accepts both static content such as web pages or a copy of a movie file as well as dynamic, live content such as a live streaming web camera or a URL which points to a web page that is constantly updated. The soup can exist on any computer belonging to the Speakeasy network and remain accessible to all display devices. As a result, data can then be accessed and displayed in a variety of locations within the display system ranging from large format displays in common areas to more personalized "private-public" displays in offices or other private locations.

Content in the soup is displayed using a screen-saver-like application. This application can be installed on any of the types of displays supported by the system. Once data has been shipped from the soup to the screen saver, the screen saver cycles through data and rotates which postings are shown at any given time. Periodically, the screen saver checks the soup to see if new content has been added. This modular approach to building the backend infrastructure for the graphical display makes it possible to view and interact with displayed content from any location on the Speakeasy network.

We have also developed basic interactive support with the soup. Using an external wireless keyboard-mouse unit, a user can select an item in the soup and render it in a larger window for closer examination. Other operations will be

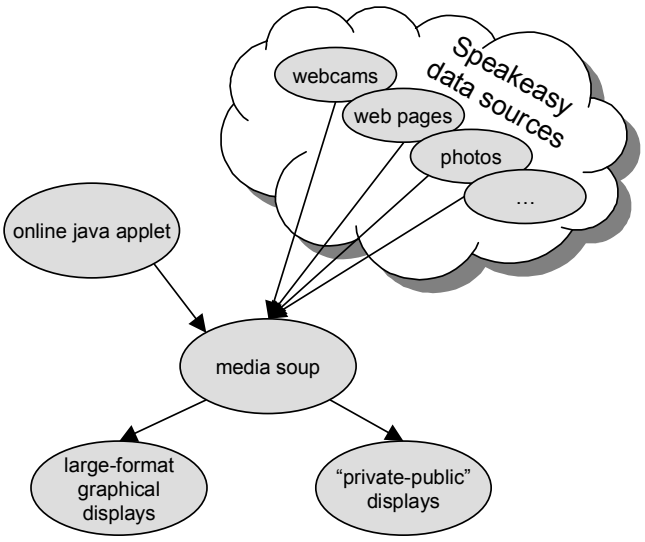

Figure 3: Users can add content to the soup by connecting Speakeasy sources to it or via the online Java applet. Data is extracted from the soup for viewing on public or private-public displays.

supported in the near future such as printing and emailing content or contacting a poster for more information.

Additionally, in order to support contributions from non-Speakeasy community members, we provide, as with the text-based LED sign, an online Java applet that allows users to add web page URLs to the soup. We leverage Speakeasy's pre-existing ability to render various media such as URL's and other documents, as well as video streams.

Our current screen-saver implementation is very simple: it renders the soup contents to a set of dynamically changing "tiles," rendered in a grid layout. The location and dimensions of this grid can be set at initialization time. We anticipate exploring more interesting and customizable display clients in the future.

CONCLUSION

Our experience has shown that the Speakeasy infrastructure provides a strong foundation for the development of Public Display technology. Leveraging Speakeasy allows us to support numerous data types and allow ad hoc interactions between components and devices exporting and accepting viewable data. Additionally, our Speakeasy based system, by supporting flexible interactions between various data sources, the soup, and arbitrary displays, enables a modular and portable public display system in which individuals can view the content of the public display in numerous locations.

REFERENCES

1. Campiello: Interacting in collaborative environments to promote and sustain the meeting between inhabitants and tourists. http://klee.cootech.disco.unimib.it/Campiello/index.html
2. Edwards, W.K., Newman, M.W., Sedivy, J.Z., Smith, T.F. and Izadi, S. Challenge: Recombinant Computing and the Speakeasy Approach. *Proceedings of The Eighth ACM International Conference on Mobile Computing and Networking (Mobicom 2002)*, (Atlanta, GA USA, 2002).
3. Edwards, W.K., Newman, M. Sedivy, J. Smith, T. Using Speakeasy for Ad Hoc Peer-to-Peer Collaboration. *Proceedings of CSCW 2002* (New Orleans, LA. November 2002) (to appear).
4. Greenberg, S., Rounding M. The Notification Collage: Posting Information to Public and Personal Displays. *Proceedings of the SIGCHI conference on Human factors in computing systems* (Seattle, WA. 2001).
5. Koch, M., Rancati, A., Grasso, A., Snowdon, D. Paper User-Interfaces for Local Community Support. *Proceedings of HCI International 99 Vol. 2*. London: Lawrence Erlbaum Publishers, 1999.
6. Schuler, D. Community Networks: Building a New Participatory Medium. *Communications of the ACM*, Vol. 37, no.1. January, 1994.