

# Web-based Framework for Spatiotemporal Screen Real Estate Management of Interactive Public Displays

Tomas Lindén, Tommi Heikkinen, Timo Ojala, Hannu Kukka, Marko Jurmu

MediaTeam Oulu, Department of Electrical and Information Engineering, University of Oulu  
P.O.Box 4500, 90014 University of Oulu, Finland

{firstname.lastname}@ee.oulu.fi

## ABSTRACT

In this paper we present a web-based framework for spatiotemporal screen real estate management of interactive public displays. The framework facilitates dynamic partitioning of the screen real estate into virtual screens assigned for multiple concurrent web applications. The framework is utilized in the implementation of so-called UBI-hotspot, which provides various information services via different interaction modalities including mobile. The framework facilitates seamless integration of third party web applications residing anywhere in the public Internet into the UBI-hotspot, thus catering for a scalable and open architecture. We report the deployment of a network of indoor and outdoor UBI-hotspots at downtown Oulu, Finland. The quantitative data on the usage of the UBI-hotspots implicitly speaks in favor of the practical applicability of the framework.

## Categories and Subject Descriptors

H.5.2 [Information Interfaces and Presentation]: User Interfaces – *user interface management systems, windowing systems.*

## General Terms

Experimentation, Human Factors

## Keywords

Ubiquitous computing, urban computing, UBI-hotspot

## 1. INTRODUCTION

Contemporary advances in display technology, along with the rapidly falling costs of high-end panels suitable for outdoor use, have made modern-day urban areas increasingly saturated with large public displays. From the extravagant wall-size displays in Times Square to information screens in airports, digital signage systems in malls, or even small touch-screen enabled panels inside taxi cabs, broadcast information is becoming increasingly digital. While the advantages of digital signage system over analog paper-based advertising are obvious, a recent trend has seen growing interest in also embedding interactive services in the public displays. Few commercial installations have been carried out, the most prominent being perhaps the BBC Big Screens in the UK (<http://www.bbc.co.uk/bigscreens/>). These giant displays installed in highly public locations provide broadcast information, and also rudimentary interactivity based on, for instance, crowd movement.

Copyright is held by the International World Wide Web Conference Committee (IW3C2). Distribution of these papers is limited to classroom use, and personal use by others.

WWW 2010, April 26–30, 2010, Raleigh, North Carolina, USA.

ACM 978-1-60558-799-8/10/04.

Research organizations have taken an active role in studying different types of interaction possibilities with large displays. The CityWall [6] deployed at downtown Helsinki employed a large multi-touch enabled display with simple web-based multimedia applications. The purpose of the experiment was to study the different roles people take when presented with a multi-user display in a public setting, and also to study the so-called ‘honeypot’ phenomenon, where ongoing interaction on a public display serves as an attention incentive for others to also approach the display and shift from passive on-lookers to active participants. The iDisplays project [4] focused on smaller displays placed in public phone booths to support pedestrian navigation in urban spaces.

The multi-user and multi-purpose use of public displays has also been a topic of interest in several projects. Vogel and Balakrishnan [7] presented a dynamic display partitioning system based on proximity-related contextual cues. Their interaction model utilizes the different affordances of a large display in relation to the user’s perceived interaction intent, ranging from ambient broadcast information to personal direct-input use, based on the user’s physical orientation in relation to the display. The Dynamo system [1] enables ‘carving out’ personal partitions of a public display for individual programs or sharing media.

While these projects have been successful in demonstrating the different capabilities and affordances of interactive public displays in laboratory or semi-fixed public settings, they have been rather limited when considering a large-scale deployment in an actual urban context. In order for a system to be applicable in a real-life setting for a long period of time, issues such as openness, scalability, reliability and maintenance have to be considered.

In this paper we present the design and implementation of a web-based framework for spatiotemporal screen real estate management of interactive public displays. We report how we have successfully utilized the framework in the implementation of a network of interactive touch-screen enabled UBI-hotspots installed at downtown Oulu, Finland, providing a diverse range of information services. Our hotspots are open for anyone and everyone to use, and currently we are logging thousands of usage sessions each week. We demonstrate how the proposed framework enables our hotspots to run several concurrent applications on a single display, and allows switching between applications based on interaction events. Since our framework is based on web technology, it allows seamless integration and deployment of third party services residing anywhere in the public Internet, thus catering for openness and scalability.

A comprehensive description of the motivation, design, implementation and evaluation of our UBI-hotspot infrastructure is presented in [5]. The functionality of an UBI-hotspot is illustrated by the video found at <http://www.ubioulu.fi/node/133>.

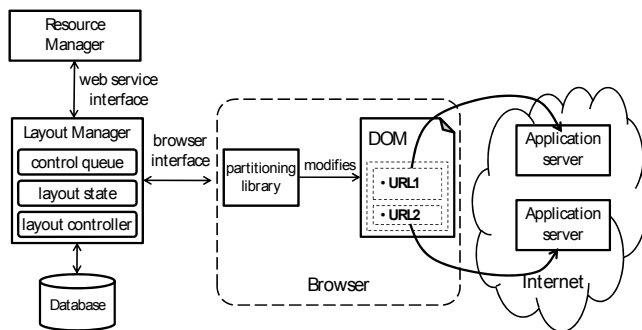
## 2. DESIGN

We set the following functional requirements to our screen real estate management framework:

- It must support dynamic run-time partitioning of the screen real estate between multiple concurrent applications both in spatial and temporal dimensions;
- It must allow integration of external third party applications as such;
- It must be OS independent.

Given these requirements, the web appears as the most suitable platform. With a pure Java based solution the integration of third party applications would have been challenging. The last requirement in turn eliminates all OS specific virtual display or desktop managers such as iShadow (<http://www.ishadow.com>).

Fig. 1 illustrates the conceptual architecture of the proposed framework. The Resource Manager controls the temporal and spatial allocation of the screen real estate between applications. The Layout Manager controls the actual layout of the screen implemented by the Browser running in full screen mode. The full screen is partitioned into virtual screens assigned to different applications which are identified by their URLs.



**Figure 1. Conceptual architecture of the proposed framework.**

The Layout Manager contains a control queue which buffers the incoming requests from the Resource Manager. The layout state component is effectively a state machine that contains the specification of allowed virtual screen layouts and possible transitions between them. The layout controller component manages the current layout, i.e. the dimensions, absolute position and the application URL of each virtual screen.

The interface between the Resource Manager and the Layout Manager provides functions for changing the layout state and setting the URL of a virtual screen. Any incoming calls that are received during a state transition are placed in the control queue, and any requests in the queue are processed in a FIFO manner until the queue is empty. If a requested transition is illegal according to the layout state, the request is discarded.

The partitioning library provides functionality for adding, removing, resizing, moving and defining the URLs of individual virtual screens. The database provides persistent storage for the run-time layout state.

## 3. IMPLEMENTATION

The Resource Manager is implemented as a threaded server, which receives events from different context sources and instructs

the layout manager according to the management policies set forth in the configuration of the system. A context source can be a touch screen foil, an RFID reader or a camera, for example. For negotiating application compositions with mobile devices, the Resource Manager employs the notion of a lease, which is a negotiated contract between the mobile client and the system, governing the execution of a chosen application. This agreement can be further specified with spatiotemporal constraints, if the management policy of the system so dictates. [2]

The Layout Manager is effectively an HTTP server implemented atop Apache Tomcat (<http://tomcat.apache.org>). The internal components of the Layout Manager are in-house Java implementations. The web service interface between the Layout Manager and the Resource Manager is implemented as SOAP API (<http://www.w3.org/TR/soap12-part1>). The database is implemented using MySQL (<http://mysql.com>).

The Browser is implemented using Mozilla Firefox (<http://firefox.com>), which renders a skeleton HTML page in full screen mode. We have also tested the framework with Google Chrome and Apple Safari. The partitioning library is an in-house JavaScript implementation, which modifies the DOM of the skeleton page to create the virtual screens. They are sandboxed as HTML *iframes*, which have specific *dimensions*, *position* and *URLs*. The two-way interface between the Layout Manager and the Browser is based on the DWR library (<http://directwebremoting.org>). It enables the Java components in the Layout Manager and the JavaScript in the partitioning library to interact and call each other. Moving and resizing of virtual screens are smoothly animated using the *setTimeout* JavaScript function to give an application enough time to gradually scale and adjust as it is being moved or transformed.

Using a standard web browser for running applications on a public display requires careful navigation control. First and foremost, we want to prevent the user from navigating outside the applications. Therefore, the applications should not contain any hyperlinks pointing to external websites. As an additional safeguard we employ a Mozilla add-on called BlockSite (<https://addons.mozilla.org/en-US/firefox/addon/3145>). Together with some tweaking of the browser configuration it allows for filtering out unwanted link clicks. The browser is further protected against tampering by using another add-on called R-kiosk (<https://addons.mozilla.org/en-US/firefox/addon/1659>), which also enables true full screen functionality.

## 4. DEPLOYMENT IN UBI-HOTSPOT

We have deployed the proposed framework in the implementation of the so-called UBI-hotspot (Fig. 2, ‘hotspot’ from now on). It is effectively a large public display which offers a wide range of services via different interaction modalities including mobile. The hotspot comes in two different versions for outdoor and indoor use. They both are customized versions of the Icon57 DID (Digital Information Display) product manufactured by Symbicon Ltd. (<http://www.iconone.fi>). While the indoor version has one 57” TFT LCD full HD panel controlled by industrial grade quad core PC, the outdoor version is two-sided with two LCD panels back-to-back and separate control PC’s for each panel. The panel is protected by a 6 mm hardened safety glass. Behind the glass is a projected capacitive touch screen foil. The hotspot is equipped with a loudspeaker, a RFID reader, two overhead cameras, a Bluetooth access point and a WLAN access point.

The interaction model of the hotspot is based on a slightly simplified version of Vogel's [6] framework for interaction phases on public displays. The hotspot alternates between a passive broadcast mode (Vogel's *ambient display phase*) and an interactive mode (Vogel's *personal interaction phase*). The transition between the modes is triggered by the following interaction events: a user touches the touch screen, a face is detected from the video feed of the two overhead cameras (Vogel's *subtle interaction phase*) or a user presents an RFID tag to the RFID reader. The Resource Manager subscribes to these events and instructs the Layout Manager to manage the screen real estate accordingly. The end of the interaction is determined with a timeout since the last interaction event, after which the hotspot returns to the broadcast mode. We refer to these interaction periods commonly as *sessions*.

In the broadcast mode the whole screen constitutes a single virtual screen allocated to a digital signage service called UBI-channel (Fig. 2(a)). In the transition to the interactive mode it is smoothly squeezed into the upper left hand part of the screen and two additional virtual screens are created, one assigned for a touch screen portal called UBI-portal and another assigned for mobile services (Fig. 2(b)).



(a)



(b)

**Figure 2. UBI-hotspot in the passive broadcast mode (a) and in the interactive mode (b).**

The UBI-channel repeats a playlist comprising of 'spots', which are either images in full HD resolution or videos either in HD ready or full HD resolution. The playlist is generated and uploaded to the hotspot together with the actual media files by an external digital signage system. The playlist is configurable on per hotspot basis. The application URL assigned to the virtual screen corresponds to the URL of the open source JW FLV player (<http://www.longtailvideo.com>), which is executed on the local Tomcat server.

The UBI-portal is effectively a web portal of various information and leisure services (web pages). The default view of the UBI-portal is stored in an URL provided by the Resource Manager and configured on per hotspot basis. The web pages aggregated to the portal are referenced by URL's and can reside on any web server in the public Internet. The generic functionalities include language selection, a service specific help page and two buttons for giving "thumbs up" or "thumbs down" for the currently selected page.

The UBI-portal employs a two-level navigation hierarchy, where the 1<sup>st</sup> level pages are *Oulu Today*: up-to-date news, events and weather information, hosted by the local main newspaper, demonstrating the openness and extensibility of our architecture; *Map*: a map-based service and information directory based on Google Maps, including bus stops with real-time schedules, restaurants, cafés, hotels, libraries, restaurants, and UBI-hotspots; *Multimedia*: multimedia content such as images and video, including user-created content uploaded from personal mobile phones; *Fun and Games*: entertainment and leisure services; *Help*: usage instructions and relevant information; and *Survey*: online questionnaire for collecting user feedback.

A 1<sup>st</sup> level tab may have a second level of tabs corresponding to services within the first level category. The 2<sup>nd</sup> level tabs under *Multimedia* are *Photos*: a collection of images with the tag *Oulu* retrieved from Flickr and displayed with a customized version of the *TagGalaxy* (<http://www.taggalaxy.de>) 3D browsing interface; *UBI-photos*: a collection of photos uploaded by users from their mobile phones. The uploaded images are tagged with appropriate tags and stored in Flickr. The photos are also displayed using the TagGalaxy interface.; *Videos*: a collection of general videos about Oulu provided mostly by the City of Oulu and stored on our own central media server. *UBI-videos*: videos uploaded by users from their mobile phones. Videos are stored in YouTube and displayed on a customized interface; and *Street Gallery*: an interactive art gallery for presenting exhibitions provided by and coming new media artists. Inappropriate user uploaded content in UBI-photos and UBI-videos can be flagged and temporarily removed by the community for later inspection by a moderator.

The 2<sup>nd</sup> level tabs under *Fun and Games* are *UbiPostCard*: a social application where using the integrated cameras people can take a photograph of themselves in front of a hotspot and send it to a friend by email. *Hangman*: the traditional game of hangman; and *UbiFridge*: a social word-game application where users can form short sentences by moving individual words represented by refrigerator-door magnets, and can send the sentence to either a mobile phone or to a friend over email. The word magnets are display-panel-specific and the state is maintained between user sessions. Thus a subsequent user can continue the sentences made by a previous user.

To use the mobile services of the hotspot the user has to first authenticate himself/herself by presenting a so-called personal UBI-key to the RFID reader. The UBI-key is effectively an RFID

tag paired with the personal mobile phone. When the user selects a service from the mobile client, the Resource Manager fetches the corresponding metadata description from a database. The metadata includes the URL of the corresponding application, which is rendered to the virtual screen assigned for the mobile applications.

The mobile services include *UbiPoker*: a multi-player Texas hold'em game, demonstrating the multi-user capabilities of the hotspot.; *UbiAlbum*: an application for users to manage the photos they had uploaded to the displays (add/remove tags from photos, remove photos or videos etc.); *BlueInfo*: a collection of information services including bus schedules, weather information, news from Kaleva and Reuters, and TV programs, all of which can be pulled to a registered mobile device over a Bluetooth connection from the hotspot [3]; and *PlaceMessaging*: a bulletin-board service allowing people to post notes containing text and images from their mobile phone to the display. Each hotspot has its own PlaceMessaging board, thus enabling place-based messaging.

## 5. EVALUATION

We first characterize the performance of the proposed framework in terms of its memory and CPU load. When the hotspot is in the interactive state, the Firefox browser process uses ~350 MB of memory and the Layout Manager process ~50 MB. The Firefox process imposes ~30% load on the 3 GHz CPU, while the load of the Layout Manager is negligible.

We have deployed a network of 12 hotspots at key outdoor (6) and indoor (6) locations around downtown Oulu. We have systematically logged the usage of the hotspots since July 17, 2009. We have had hardware failures and other issues, which have temporarily rendered a particular hotspot or a panel unavailable. When a panel is considered to have been available for a given day if it had at least one successful session that day, the average availability of the 18 panels has improved from 78% in July 2009 to >99% from October 2009 onwards.

The Layout Manager logs the start of each session and the clicks in the UBI-portal. The Resource Manager logs the interaction events of a face being detected, the screen being touched, or an UBI-key being read by the RFID reader. During the 199-day period from July 17, 2009 till Jan 31, 2010, the hotspots have had 180370 sessions (59525 (33 %) indoors, 120845 (67 %) outdoors) in total. This corresponds to on average of ~906 sessions per day or ~50 daily sessions per panel. The average session duration has been 187 s (~3.1 minutes). On average in 26 % (30 % indoors, 24 % outdoors) of the sessions the user clicked any 'control button' in the UBI-portal at least once after the transition to the interactive state (i.e. the first touch triggering the interactive state is not taken into account). The proportion of such sessions has dropped from 35 % (28 % indoors, 39 % outdoors) in July 2009 to 20 % (28 % indoors, 12 % outdoors) in January 2010, as the arctic conditions in Oulu during the winter months do not invite the users to fiddle with the outdoor hotspots. The remaining sessions correspond to the subtle interaction phase, e.g. someone in a group of people facing the hotspot sufficiently close or just briefly glancing at an interesting spot in the UBI-channel, which does not lead to further personal interaction. The average number of clicks per session has been 6.8 and median 4 clicks.

We acknowledge that these statistics merely quantify the voluminous usage of the UBI-portal instead of providing direct evidence of the quality and performance of the proposed

framework. However, the extensive successful usage of the UBI-portal implicitly testifies in favor of the practical applicability of the framework.

## 6. CONCLUSION

We presented a web-based framework for spatiotemporal screen real estate management of interactive public displays. We have successfully deployed the framework in the implementation of UBI-hotspots, which provide a wide range of web services. We have by now installed a dozen UBI-hotspots in public outdoor and indoor locations around downtown Oulu, Finland. The ongoing active usage of the UBI-hotspots implicitly speaks in favor of the practical applicability of the proposed framework.

The web-based implementation of the framework has its pros and cons. The main advantage is that in principle any third party web application residing in the public Internet can be seamlessly incorporated into the UBI-hotspot. And vice versa, only web-based applications can be incorporated. Further, the web applications have to be designed to be scalable to accommodate for different layout configurations.

## 7. ACKNOWLEDGMENTS

The financial support from the European Regional Development Fund, the Finnish Funding Agency for Technology and Innovation, the City of Oulu and the Urban Interactions consortium is gratefully acknowledged.

## 8. REFERENCES

- [1] Izadi, S., Brignull, H., Rodden, T., Rogers, Y., and Underwood, M. Dynamo: a public interactive surface supporting the cooperative sharing and exchange of media. In Proceedings of UIST'03 (Vancouver, Canada, November 2003), 159-168.
- [2] Jurmu, M., Boring, S., and Riekkilä, J. ScreenSpot: Multidimensional resource discovery for distributed applications in smart spaces. In Proceedings of MobiQuitous'08 (Dublin, Ireland, July 2008), article no. 41.
- [3] Kukka, H., Kruger, F., and Ojala, T. BlueInfo: Open architecture for deploying web services in WPAN hotspots. In Proceedings of ICWS 2009 (Los Angeles, CA, USA, July 2009), 984-991.
- [4] Müller, J., Wilmsmann, D., Exeler, J., Buzeck, M., Schmidt, A., Jay, T., and Krüger, A. Display blindness: the effect of expectations towards digital signage. In Proceedings of Pervasive 2009 (Nara, Japan, May 2009), 1-8.
- [5] Ojala, T., Kukka, H., Lindén, T., Heikkinen, T., Jurmu, M., Hosio, S., and Kruger, F. UBI-hotspot 1.0: Large-scale long-term deployment of interactive public displays in real-world setting in a city centre. In Proceedings of ICIW 2010 (Barcelona, Spain, May 2010).
- [6] Peltonen, P., Kurvinen, E., Salovaara, A., Jacucci, G., Ilmonen, T., Evans, J., Oulasvirta, A., and Saarikko, P. It's Mine, Don't Touch!: interactions at a large multi-touch display in a city centre. In Proceedings of CHI'08 (Florence, Italy, April 2008), 1285-1294.
- [7] Vogel, D. and Balakrishnan, R. Interactive public ambient displays: transitioning from implicit to explicit, public to personal, interaction with multiple users. In Proceedings of UIST'04 (Santa Fe, NM, October 2004), 137-146.