Available online at www.sciencedirect.com
ScienceDirect
Computers & Graphics 30 (2006) 681–691
COMPUTERS & GRAPHICS
www.elsevier.com/locate/cag

Mobile Computing and Ambient Intelligence

Supporting content scheduling on situated public displays

Oliver Storz*, Adrian Friday, Nigel Davies

Computing Department, Infolab 21, Lancaster University, Lancaster, UK

Abstract

There is increasing interest in creating networks of situated public displays that offer novel forms of interaction and rich media content—often as work towards a vision of ubiquitous computing or ambient multimedia. In this paper, we present an infrastructure developed as part of the e-Campus project that is designed to support the coordinated scheduling of rich media content on networks of situated public displays. The design of the system was informed by an iterative process of developing, deploying and evaluating a set of three technology probes. The resulting system provides flexible support for the construction of domain-specific scheduling approaches on top of a common, domain-independent API. Using this approach, we are able to support a combination of both statically scheduled content and interactive content across multiple displays. The API provides support for transactional semantics, allowing developers of schedulers to reliably schedule content across displays in the presence of conflicts and failures without negative impact on running applications.

Keywords: Digital signage; Public displays; Distributed systems infrastructure; Ubiquitous computing; Coordination

1. Introduction

Increasingly we are witnessing the penetration of display technologies into urban environments and public spaces. In the UK, the BBC have deployed several large situated displays (the ''big screens'') in urban centres.[1] In airport terminals, train stations and even on public transport itself, physical posters have been replaced with LCD or plasma displays. Projectors and steerable projections turn conventional surfaces such as floors and walls into displays. These displays present us with a combination of digital signage, information, traditional broadcast media and advertising. This trend will only continue as novel low-cost display technologies such as printable displays and e-paper move toward production.

In the future we believe the situated and dynamic nature of these displays will mean that new interactive and context-aware applications become possible—extending such systems well beyond the current state-of-the-art in digital signage. For example, the displays could:

- select content or adapt the presentation of content, e.g. to target the interests of passers by.
- enable more complex forms of interaction (e.g. using mobile devices [1] or providing steerable content[2]).
- form a network of multiple displays that can be exploited simultaneously or in sequence over time to present information in novel ways, e.g. supporting navigation, games, visualisations, art and community interaction.

We are seeking to prototype such an environment in the e-Campus project. In e-Campus we are deploying a series of display installations in public spaces around our campus. One of the principle challenges in supporting these deployments is to create a software infrastructure that allows multiple concurrent applications to schedule content across the shared display network. This is crucial to enable scheduling requests to be generated from a wide variety of sources including interactive applications and domain specific schedulers. For example, scheduling requests might be generated from multimedia presentation software or multimedia document viewers. In this paper, we discuss the design and implementation of such a software

*Corresponding author.
E-mail addresses: oliver@comp.lancs.ac.uk (O. Storz), adrian@comp.lancs.ac.uk (A. Friday), nigel@comp.lancs.ac.uk (N. Davies).

[1] http://www.bbc.co.uk/bigscreens.

[2] http://www.blendomedia.com.

0097-8493/$ - see front matter © 2006 Elsevier Ltd. All rights reserved.
doi:10.1016/j.cag.2006.07.002

infrastructure—focusing specifically on the computational model and API that we provide to programmers wishing to develop schedulers for the network of displays. To understand the requirements and refine our design to better support this domain we have iterated our software through three public display deployments. The resulting design offers a simple API for supporting the creation of heterogeneous content schedulers and provides fine-grain control over the mapping of content to one or more displays. Importantly, the API supports a transaction like concept that ensures content only becomes visible if all needed display and content resources are available. We have built a working implementation of this API and it is currently being used to support an ongoing deployment on our University campus.

The paper is structured as follows: in Section 2 we describe the iterative process and technology probes we have used to refine the design of our software and reflect on the design impact on our API at each stage. In Section 3 we describe the computational model, API and associated system support that we provide to programmers for scheduling content on a network of situated displays. Section 4 describes and evaluates our current implementation, reflecting on our use of the API in supporting the most recent of our deployments. We discuss related work in Section 5 and finish with our concluding remarks in Section 6.

2. Initial technology probes

To date we have deployed three situated display prototypes: a digital signage system at a conference, an installation at a local gallery, and a larger, more permanent installation in an underground bus station on campus (Fig. 1). These prototypes have served as technology probes, generating requirements for, and allowing us to experiment with and refine the design of our software. In this section, we review each of the installations and discuss the impact and refinements to our design.

2.1. Installation 1: WMCSA 2004 conference signage

2.1.1. Overview

Our first technology probe deployed a digital signage solution at the 6th IEEE Workshop on Mobile Computing Systems and Applications, WMCSA 2004. The WMCSA system consisted of four public displays stationed outside each of the entrances to the main auditorium and demo room. The displays provided a rolling display of information for delegates tailored to the display's location (proximity to ongoing conference activities) and the time of day. Each display was able to show information relevant to the talks being presented in the adjacent rooms, about activities in the wider locale, and navigation symbols directing delegates to refreshments at appropriate times of the day. The displays were interconnected via a local network, allowing us to synchronise content across the displays on a per content item basis.

One of the key issues we sought to explore with the WMCSA deployment was how to simplify the process of injecting content into the system and of mapping that content to displays. We did this by exploiting a separation of concerns: authors could create content items (images,

Fig. 1. Opening of the Metamorphosis installation using the public displays in the Underpass. Inset: members of .:thePooch:. arts collective keeping a close eye on the system console.