
Demonstrating HuddleLamp: Spatially-Aware Mobile Displays for Ad-hoc Around-the-Table Collaboration

Roman Rädle

HCI Group
University of Konstanz
Roman.Raedle@uni-konstanz.de

Harald Reiterer

HCI Group
University of Konstanz
Harald.Reiterer@uni-konstanz.de

Hans-Christian Jetter

ICRI Cities
University College London
h.jetter@ucl.ac.uk

Yvonne Rogers

UCL Interaction Centre
University College London
y.rogers@ucl.ac.uk

Nicolai Marquardt

UCL Interaction Centre
University College London
n.marquardt@ucl.ac.uk

Permission to make digital or hard copies of part or all 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. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

Copyright is held by the owner/author(s).

ITS '14, Nov 16-19 2014, Dresden, Germany
ACM 978-1-4503-2587-5/14/11.
<http://dx.doi.org/10.1145/2669485.2669500>

Abstract

We present HuddleLamp, a desk lamp with an integrated RGB-D camera that precisely tracks the movements and positions of mobile displays and hands on a table. This enables a new breed of spatially-aware multi-user and multi-device applications for around-the-table collaboration without an interactive tabletop. At any time users can add or remove displays and reconfigure them in space in an ad-hoc manner without the need of installing any software or attaching markers. Additionally, hands are tracked to detect interactions above and between displays, enabling fluent cross-device interactions. The demo consists of the technical implementation of HuddleLamp's hybrid sensing and a Web-based architecture for installation-free ad-hoc collaboration. We demonstrate our implementation by showing a variety of possible interaction techniques.

Author Keywords

Mobile devices; cross-device interactions; depth camera; ad hoc collaboration; society of devices; ubicomp ecologies.

ACM Classification Keywords

H.5.2 [Information Interfaces and Presentation]: Input

Introduction

We are witnessing an explosive growth of the number and density of powerful mobile devices around us. However, their great majority are still blind to the presence of other devices and performing tasks among them is usually tedious [1] and lacks recognizable guiding principles [7]. We envision a future in which mobile devices can contribute their interaction resources (e.g., their multi-touch displays) to a community of devices in their proximity that then serves as one seamless multi-device user interface (UI). At any time, users can dynamically compose and reconfigure this UI according to their current needs and the task at hand.

In this vision, users can seamlessly add or remove devices from the community in an ad-hoc fashion without explicit setup or pairing. Instead this happens implicitly as a by-product of natural use in space, for example, by bringing multiple devices to the same room, placing them side-by-side on a table or desk, and moving them around as needed. Ideally users will experience these co-located cooperating devices and reconfigurable displays as one seamless and natural UI for ad-hoc co-located collaboration.

As the first incarnation of this vision, we present HuddleLamp, a modified desk lamp with an integrated low-cost RGB-D camera that detects and identifies multiple mobile displays (e.g., smartphones or tablets) on desks or meeting tables and precisely tracks their spatial configuration and movements over time. Thereby users can add or remove off-the-shelf web-enabled devices of many different sizes in an ad-hoc fashion without prior installation of custom hardware (e.g., radio modules, IR markers) or any additional

software. Due to the fully web-based architecture, no complex device pairing procedures are necessary; all a person needs to do to add a new device is opening the URL of the HuddleLamp on their device and putting it on the table so that it becomes visible to the lamp's camera. Additionally, HuddleLamp tracks users' hand positions above the surface to enable interactions between and above devices such as 'pick, drag, and drop'.

This enables a new kind of computer-supported around-the-table collaboration without interactive tabletops. Users can still sit around ordinary tables that can remain cluttered with non-digital objects (e.g., printouts, maps, notebooks) while their digital collaborations happen using spatially-aware mobile screens that blend into existing spatial and social practices. Unlike traditional interactive tabletops, HuddleLamp is portable and low-cost, so it can be used in improvised settings or where the costs for hardware and administration typically prohibit the use of large interactive tabletops, e.g., public libraries, schools, community centres.

Related Work

In order to enable cross-device interactions with co-located devices [5], the system first needs knowledge about the presence and position of other devices around it. Previous research has explored various techniques for establishing connections between co-located devices. Some techniques initiate a connection between devices after the user performs action sensed by the system [3]. Another approach is to use custom sensing hardware, for example infrared-, hall-, or radio-based position sensing [4,6]. More recently, researchers explored using the devices' built-in



Figure 1. HuddleLamp supports a device ecology of heterogeneous form factors and display sizes for ad-hoc around-the-table collaboration.




Figure 2. HuddleLamp setup with the modified desk lamp, the integrated RGB-D camera, and the PC for processing.

cameras to infer location information [9]. Alternatively, cameras positioned above a table can track the spatial layout of devices placed on the table below [8].

A special case of co-located mobile devices are (re)configurable tiled displays. For expressive interaction with digital content, the interaction space with digital surfaces can be extended beyond pure multi-touch interaction to include the space above the surface or around surfaces [2].

In HuddleLamp, we support a device ecology of heterogeneous form factors and display sizes that can be used in an ad-hoc fashion – all without requiring custom software as our platform uses entirely web-based standards. One important design goal of the HuddleLamp system is to enable similar cross-device interactions, around the ad-hoc assemblages of phones and tablets on a table (Figure 1).

HuddleLamp's Architecture

During the development of HuddleLamp, we generally used off-the-shelf hardware and free or open source software that can be easily acquired by other researchers for deployments outside our lab. In the following, we describe our technical setup and the essential architecture.

Hardware Components

The HuddleLamp uses a low-cost short range time-of-flight RGB-D camera which delivers a 1280x720 RGB colour image and a 320x240 depth image at 25 or 30 fps. It shares its technical specifications with a Creative Senz3D or SoftKinetic DepthSense 325. The camera is fixed to an Artemide Tolomeo Tavolo desk lamp (Figure 2) in which it replaces the light bulb (top right corner of

Figure 2). Using the lamp, users can conveniently move the camera into its downward-facing operating position that lies 80cm above the horizontal surface to track. This results in a rectangular tracking region of approximately 110x60cm from which the camera receives sufficient RGB and depth information to track mobile devices, their spatial configurations, and users' hands and to distinguish them from non-interactive objects.

The camera was chosen for its small form factor (11x3x2.5cm) and the low noise in its depth data. Future versions of HuddleLamp could use next generation high-resolution and wide-angle depth cameras such as Kinect v2 to increase the size of the tracked area. However, we recommend time-of-flight rather than structured-light depth cameras. An advantage of the particular camera and its Perceptual SDK by Intel is that its RGB and depth images can be aligned without further calibration. Therefore it is easy to retrieve RGB values from the colour image for each pixel in the depth image and vice versa. This greatly facilitates realizing our hybrid sensing approach.

For vision processing and for hosting a web server for communication with the mobile devices, HuddleLamp uses a Windows multicore desktop PC or laptop (Figure 2, left). We have considered using small single-board computers that can be integrated directly into the lamp (e.g. Raspberry Pi) to make HuddleLamp truly portable. However, at this stage, the necessary vision processing and algorithms are too computationally expensive for achieving our targeted tracking rate of 25-30 fps with single-board PCs.

The software for the processing of the RGB and depth camera data was implemented with Microsoft .NET Framework 4.5.1 (C#/WPF) and the Emgu CV wrapper for accessing OpenCV with C#. For finding and decoding fiducial markers in the RGB stream we use the glyph decoder of the AForge.NET library¹.

System Architecture

The camera data is processed by the vision processing and the locations of all devices and hands are streamed to the connected mobile client devices using JSON and a web socket server. The clients access this stream using JavaScript from state-of-the-art browsers on mobile devices like Safari Mobile, Chrome, and Internet Explorer. Therefore an installation of native software on the devices is not required, which is of great importance for enabling ad-hoc scenarios.

For the collaborative applications, a shared data model implemented with the Meteor web framework² was used to synchronize application data across multiple devices. Meteor's data-driven protocol compensates for latency to achieve a reactive user interface.

Conclusion

We have described HuddleLamp, a sensing system in the form of a desk lamp with an integrated RGB-D camera that tracks the movements of multiple mobile displays on a table for around-the-table collaboration. Our implementation based on a hybrid sensing approach detects mobile screens by exploiting their optical properties in the IR range and tracks fiducial markers to distinguish them from other objects or

users' hands. We also introduced HuddleLamp's web-based architecture for ad-hoc collaboration that enables users to add or remove displays and reconfigure them in space at any time without installing any software or attaching markers.

References

- [1] Greenberg, S., Marquardt, N., Ballendat, T., Diaz-Marino, R., and Wang, M. Proxemic Interactions: The New Ubicomp? *interactions* 18, January (2011), 42–50.
- [2] Hilliges, O., Izadi, S., Wilson, A.D., Hodges, S., Garcia-Mendoza, A., and Butz, A. Interactions in the air: adding further depth to interactive tabletops. *In Proc. of UIST '09*, ACM Press (2009), 139–148.
- [3] Hinckley, K., Ramos, G., Guimbretiere, F., Baudisch, P., and Smith, M. Stitching: pen gestures that span multiple displays. *In Proc. of AVI '04*, ACM Press (2004), 23.
- [4] Lucero, A., Holopainen, J., and Jokela, T. Pass-them-around: collaborative use of mobile phones for photo sharing. *In Proc of CHI '11*, ACM Press (2011), 1787–1796.
- [5] Lucero, A., Jones, M., Jokela, T., and Robinson, S. Mobile Collocated Interactions: Taking an Offline Break Together. *interactions* 20, 2 (2013), 26.
- [6] Merrill, D., Kalanithi, J., and Maes, P. Siftables: towards sensor network user interfaces. *In Proc. of TEI '07*, ACM Press (2007), 75.
- [7] Oulasvirta, A. When users “do” the Ubicomp. *interactions* 15, 2 (2008), 6.
- [8] Piazza, T., Zhao, S., Ramos, G., Yantaç, A.E., and Fjeld, M. Dynamic duo: phone-tablet interaction on tabletops. *In Proc. of CHI EA '13*, ACM Press (2013), 2803–2804.
- [9] Schmitz, A., Li, M., Schonefeld, V., and Kobbelt, L. Ad-Hoc Multi-Displays for Mobile Interactive Applications. *The Eurographics Association*, (2010), 45–52.

¹ <http://www.aforgenet.com> - AForge.NET library

² <http://www.meteor.com> - Open Source Web Platform