

Opportunities and Challenges for Cross-Device Interactions in the Wild

Steven Houben, Lancaster University

Nicolai Marquardt, University College London

Jo Vermeulen, University of Calgary

Clemens Klokmoose, Aarhus University

Johannes Schöning, University of Bremen

Harald Reiterer, University of Konstanz

Christian Holz, Microsoft Research

People are increasingly using multiple computing devices in their daily lives as portals into a shared online information space. We can select devices based on their form factor and affordances to match our task and context of use. Moreover, we are often using multiple devices at once, for example when sharing content, playing games, controlling other devices such as a smart TV, or collaboratively editing and presenting content in a meeting. The seamless use of multiple devices to work toward achieving the same goal is known as *cross-device interaction*. Cross-device interaction opens up new opportunities for how we interact with digital technology, but it also presents several fundamental challenges.

To uncover some of these challenges, problems, and opportunities with cross-device interaction, we organized the well-attended Cross-Surface workshop series (<http://www.cross-surface.com/>). We discussed how such technologies could be used in the wild, supporting new domains and use cases at ACM Interactive Tabletops and Surfaces (ITS) 2015 [1]; how we could move away from Weiser's vision of ubiquitous devices to a "bring your own device" approach at CHI 2016 [2]; and finally, how space and spatial relations between people and devices could be used to support better device awareness at ACM Interactive Spaces and Surfaces (ISS) 2016 [3].

In 1991, Mark Weiser outlined his vision for *ubiquitous computing* [4], in which people interact with multiple computing devices in different form factors. In his vision, people have ready access to a plethora of devices to pick up and use seamlessly, allowing them to interact with content through cross-device interactions. This work accelerated academic research into providing new technologies, conceptual models, and interaction techniques that support interactions across device *ecologies* [5,6]. However, despite the availability of such a range of devices, technologies, and techniques, has this vision of cross-device ecologies really materialized?

Weiser introduced three types of devices: inch-scale pads (similar to today's smartphones), foot-scale pads (similar to today's tablets), and yard-scale boards (similar to today's large displays). Two decades later, this aspect of Weiser's vision concerning interaction with multiple devices in

different form factors has certainly become a reality. We now have access to more types of computing devices than ever before. Advances in the consumer market have led to a wide range of static and mobile interactive devices such as smartphones, tablets, laptops, interactive televisions, smart watches, large multi-touch surfaces, and smart glasses. These devices are packed with sensors and diverse input/output modalities, the combination of which enables many different interesting interactions. Moreover, the way that people are using these computing devices is changing rapidly. People no longer interact only with files, icons, and documents, but now also have access to real-time online interactive media, tools, services, games, and social media. They've moved from using a single device to using multiple devices.

Nevertheless, handling information and services across different devices and people is still challenging on both a technical and a conceptual level. Creating meaningful, seamless connections across a device ecology is often very difficult unless such connections are pre-configured and the devices are compatible. For example, distributing a game across two tablets to provide more screen real estate is difficult to achieve with current commercial technology platforms. While moving an ongoing telephone call from your smartphone to a nearby TV or laptop might be possible, it often remains challenging and requires compatible software and platforms.

The central cause for this structural failure of easy interaction across devices can be traced back to the fact that devices and their containing software are still fundamentally designed for isolated personal use. Current technology is geared toward individual users. It encourages people to create personalized, private relations with their devices through accounts, subscription methods, application models, and privacy and authentication mechanisms. This highly personal relation between people and their devices is in stark contrast to Weiser's original vision, in which "for each person in an office, there should be hundreds of tabs, tens of pads" [4]. Furthermore, Weiser was already arguing in the '90s that "making pads, tabs, and boards work together seamlessly will require changes in the kinds of protocols by which applications, programs, and their displayed windows communicate." However, the reality of today is modern tools, applications, and the use of the cloud as a "catchall solution" for content and data. While this enables new opportunities for shared interactions, at their core, devices and operating systems are still designed (like in the 1990s) for a single user interacting with a single device. They have no notion of other devices, their interactive input/output capabilities, or their roles within users' lives. As summarized by Michel Beaudoin-Lafon, one of our panelists at the CHI 2016 Cross-Surface workshop, "We need to work toward ending the tyranny of personal computing."

Opportunities and Challenges

During the workshops, we uncovered a list of opportunities and challenges for cross-device interaction research and design. They fall into seven key topics, summarized here.

Enabling easy configuration of device ecologies. The central premise of cross-device interaction is that people can easily combine, pair, attach, and stack devices to blend the input and output space of these devices into one seamless *meta-device* composed of different distributed physical devices. However, despite decades of research, there is still a lack of fundamental understanding in how people perceive such device ecologies, as well as a lack of technologies, concepts, and established interfaces that enable people to create such connections. Ad hoc combinations and configurations of devices are fundamentally problematic. There are few discovery mechanisms for usable devices in the environment, a lack of control mechanisms that enable users to contribute or retract a device from the ecology, and little transparency for the users about what functions devices are providing and sharing.

Designing for scale and interoperability. Because devices and their operating systems are still largely designed from the perspective of a single user interacting with a single device, interconnecting different or large numbers of devices tends to be complicated. Without certain pre-programmed behavior or a priori knowledge about all possible configurations, cross-device configurations of systems are simply not possible. This makes ad hoc federations of devices that have never before been connected very difficult. Moreover, users often want to combine devices for specific tasks and activities. Most current approaches expose users to infrastructure problems where they need to connect, pair, or install additional software to make devices work together. Sharing content across devices and people is simply not an inherent part of the current computing paradigms, which are still based on the desktop metaphor and windows, icons, menus, pointer (WIMP) interaction. To move forward, devices and their operating systems need to provide better tools, systems, methods, and frameworks to exchange data as well as application and device information. And perhaps more important: Designers need to make cross-device configurations a central feature of a system's design philosophy.

Addressing cross-device interaction challenges. Although many novel cross-device interaction techniques have been proposed in literature, there is a lack of fundamental guiding principles on how to represent content, interactions, and capabilities across different heterogeneous devices. Like the UI design principles for individual devices and the guiding principles for responsive design, a set of core design principles for cross-device interaction is needed. These should cover methods to convey (i) device connections and action possibilities, (ii) distributed content and information spaces, (iii) cross-device shareable UI components, and (iv) social or spatial relations between device interfaces. Cross-device interaction poses fundamental HCI challenges due to the complexity and distributed nature of its user interface. This includes issues with discoverability of functionality, feedback, affordances, and knowing how and on which device to provide input and perceive output. How does a person know what they can do with a cross-device interface, which devices they can combine, whether these devices are effectively connected, where information is presented, and how content will be distributed? What data is shared when devices form a cross-device ecology, what happens to the data when devices are

disconnected, and how do people opt out? How do we, as designers, convey all this information, and how can we allow people to recover from potential mistakes? These and many other questions are core design challenges that need to be addressed to make cross-device interaction a reality.

Mitigating privacy and authentication challenges. Opening up interactions and the exchange of information across devices introduces a whole range of challenges with respect to the privacy of not only individual users but also larger groups of people and devices. More specifically, users need to be able to determine boundaries between personal user space on devices (non-shareable content) and semi-public content and information about the other devices in the ecology that is shared with all devices in the configuration. Users also must be able to opt in or opt out of cross-device setups depending on the situation or context of use. Imagine walking by a smart bus stop that could provide personalized travel information, but that also contains an advertising display. You may want to connect your device to the bus stop but not to the advertisements. Especially considering that many cross-device interaction systems use a tracking infrastructure to maintain the relations between the devices, it is important that users are given *some level* of control over how their devices are tracked. This could be done by setting contextual rules for explicitly or implicitly joining device ecologies but also through temporally or spatially constrained authentication mechanisms.

Building supporting infrastructures and toolkits. One major challenge in the evolution of cross-device systems is the lack of appropriate development tools, supporting infrastructures, and prototyping tools needed to build such systems. Designing cross-device systems currently requires a full-stack development process that includes everything from networking and distributed systems to sensor processing and distributed interfaces. New software frameworks and hardware are needed that enable device connections, discovery of devices and services, and data handling. We also need programming environments that have intrinsic support for distributed interfaces. Currently, most collaborative and sharing interfaces rely on Web servers in the *cloud* to support distributed functions. We argue that this is not necessary in all application contexts and that a proper division of labor across ad hoc local infrastructures and the cloud is needed.

Addressing further application domains and use cases. We believe that the domain of cross-device interaction spans the entire spectrum of computer use. Most of the commercially emerging cross-device interactions are for *personal* computing. We already see mechanisms for easing the transitions between devices in our own personal device ecology, as in, for example, Apple's Continuity, where calls can be picked up, work continued, and clipboards shared across any personal device. However, cross-device interaction and *collaboration* are two sides of the same coin. In our workshops and in related research, we have seen applications of cross-device interaction for collocated users collaborating through a mix of personal and shared devices. This includes, for example, using personal handheld devices with a shared wall-mounted display or a tabletop for collaborative sense-making in large datasets. Other examples are using personal

devices as mediators for interaction with public displays, or creating an ad hoc shared surface by combining personal devices on a table for collaboration.

There are domains where cross-device interaction has the potential to solve common everyday challenges. For example, in meeting situations, participants could use their connected personal devices for sharing, presenting, or collaborating on documents and materials. Similarly, the use of cross-device interaction in the classroom could become a resource for engagement and interaction with one another rather than a source of distraction. In addition to academic research, we have also seen cross-device systems in commercial contexts. Nintendo's Wii U is an example of cross-device interaction for gaming, where the display on the Wii U controller is used to extend or augment what is being shown on the TV screen. One of the known challenges for bringing cross-device interaction into "the wild" is to bring down what Kenton O'Hara from Microsoft Research calls the *transaction cost*. Cross-device interaction will not become a reality for the hospital doctor if it takes more than 20 seconds to transfer her assigned patients for the day's ward round from the large display in the morning conference room to her personal tablet.

Leveraging enabling technologies. Another recurring key topic at our workshops was about technologies that allow researchers and designers to build robust and scalable cross-device applications. Challenges for such technologies include: reliably and accurately tracking the positions and orientations of devices—and often people—in space; minimizing the user's manual effort to discover and pair devices; and addressing the trade-offs between local, on-device tracking to enable truly ad hoc scenarios and external tracking (e.g., through distributed sensors in the environment, such as cameras or radio-frequency beacons).

In the current state of the art, camera-based systems remain the most common method for tracking devices and people in small-space environments. Such systems use RGB, IR, and/or depth-image video streams to detect visual features to track the location and orientation of devices. However, they are designed mostly for *stationary* rather than mobile use, requiring external cameras (e.g., above a table or in the ceiling) that limit their suitability for ad hoc, real-world deployments. More recent approaches break away from external tracking and leverage on-device cameras for inside-out location tracking (e.g., unmodified front- or rear-facing cameras, or modified cameras with wide-angle lenses or filters). Challenges for camera-based tracking remain, such as minimizing occlusions (e.g., with multi-camera setups, varying angles, and adding other sensors), increasing robustness in diverse environments, and investigating the shift from outside-in to inside-out tracking to better support ad hoc cross-device application scenarios.

Another common technology for device tracking is Inertial Motion Units (IMUs), sometimes also used in combination with camera-based tracking. IMU sensors are responsive, built-in sensors in mobile devices, which makes them suitable for gesture recognition. However, most mobile devices are not equipped with high-quality IMUs, causing sensor drift and thus rendering dead reckoning unusable for ad hoc 3D tracking.

In the workshops, participants actively discussed the recent advances in tracking technologies, particularly audio processing for positioning and the advances of mobile radio tracking. First, in audio processing for ad hoc mobile tracking, commodity mobile processors can now process audio signals captured from internal microphones in real time to determine device locations (e.g., through measuring round-trip distances between devices) and establish 3D tracking across devices based on multiple such audio units. Key challenges are how to make these tracking systems robust in noisy environments and how to increase the currently limited interaction ranges. Second, radio-based tracking using components in today's mobile devices, such as WiFi, cellular signals, or Bluetooth Low Energy (BLE), estimate the distance to remote devices based on signal strengths. Such predictions are often not very accurate, since radio sensors are not designed for ranging purposes. More advanced radio units that account for spatial tracking could substantially increase this accuracy.

While researchers have examined ad hoc tracking approaches for a long time, it is the new technologies and advances in commodity hardware that let us think about mobile tracking in fresh ways. Recent promising developments include:

Better-quality built-in sensors. Future devices will feature better-quality sensors, such as multiple speakers for communication. An example is the inclusion of four speakers in the current iPad Pro, which will benefit ad hoc mobile tracking using acoustic signals by bringing 3D tracking to devices with fewer sensors (including watches and other wearables with only one microphone).

Commoditization of powerful sensors. Mass production and adoption will decrease the cost of novel emerging sensors, such as mobile depth cameras that are propelled by their use in VR headsets (e.g., Zcam, Occipital) and mobile devices (e.g., Google Tango).

Better inside-out tracking on mobile devices. Unlike with custom optical tracking devices (e.g., HoloLens), cameras on current mobile devices are under-utilized for tracking due to their position and limited field of view. Extending optical tracking to peripheral areas around the device will advance inside-out tracking on commodity mobile devices. Promising approaches include leveraging cameras with wide-angle lenses or even extreme fisheye lenses (up to 180 degrees field of view) that provide more detailed information about device periphery.

Ubiquitous tracking infrastructure. BLE has recently emerged as a fundamental technology for indoor tracking systems, showing promise to make tracking systems, embedded in the infrastructure, more ubiquitous. This will benefit mobile devices in tracking their relative locations, allowing users to seamlessly interact across devices without manual overhead.

What's Next?

Cross-device interaction is an exciting new area that has seen increased focus from both the research community and industry. Enabling interaction across many devices provides opportunities and potential benefits for application domains such as education, healthcare, and business. However, in practice, interacting across devices is challenging and often not possible. To move forward in cross-device interaction, we need to bridge the gap between the messiness and chaos of real-world ubiquitous computing [7] and the ideal of cross-device interaction and interacting across many different devices. More work is needed to explore in-the-wild use of cross-device ubiquitous computing systems. The seven challenges and opportunities we outlined here open up new and exciting avenues for research.

Endnotes

1. Houben, S., Vermeulen, J., Klokmose, C., Marquardt, N., Schöning, J., and Reiterer, H. Cross-surface: Workshop on interacting with multi-device ecologies in the wild. *Proc. of the 2015 International Conference on Interactive Tabletops & Surfaces*. ACM, 2015, 485–489.
2. Houben, S., Marquardt, N., Vermeulen, J., Schöning, J., Klokmose, C., Reiterer, H., Korsgaard, H., and Schreiner, M. Cross-surface: Challenges and opportunities for bringing your own device in the wild. *Proc. of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 2016, 3366–3372.
3. Houben, S., Vermeulen, J., Klokmose, C., Schöning, J., Marquardt, N., and Reiterer, H. Cross-surface: Challenges and opportunities of spatial and proxemic interaction. *Proc. of the 2016 ACM on Interactive Surfaces and Spaces*. ACM, 2016, 509–512
4. Weiser, M. The computer for the 21st century. *Scientific American* 265, 3 (1991), 94–104.
5. A term borrowed from Gibson's Ecological Psychology perspective describing the dynamic relation and connection between devices and people.
6. Bødker, S and Klokmose, C.N. Dynamics in artifact ecologies. *Proc. of the 7th Nordic Conference on Human-Computer Interaction: Making Sense Through Design*. ACM, 2012, 448–457.
7. Oulasvirta, A. When users do the ubicomp. *Interactions* 15, 2 (2008), 6–9.

Steven Houben is a lecturer (assistant professor) in interactive systems in the School of Computing and Communications at Lancaster University. His research goals are to design, build, and evaluate cross-device interactive systems and physical computing devices for human-data interaction and information handling. s.houben@lancaster.ac.uk

Nicolai Marquardt is a senior lecturer (associate professor) in physical computing at the University College London. At the UCL Interaction Centre, he works on projects in the research areas of ubiquitous computing, interactive surfaces, sensor-based systems, prototyping toolkits, and physical user interfaces. n.marquardt@ucl.ac.uk

Jo Vermeulen is a postdoctoral fellow in the InnoVis group at the Interactions Lab at the University of Calgary. He will be joining the Department of Computer Science at Aarhus University as assistant professor. He is interested in addressing interaction challenges within ubicomp spaces, including providing intelligibility, discoverability, feedback, and feedforward for cross-device interactions. jo@jovermeulen.com

Clemens Klokrose is an associate professor in the development of advanced interactive systems in the Department of Digital Design and Information Studies at the School of Communication and Culture, Aarhus University. His main interest is human-computer interaction that goes beyond personal computing. clemens@cs.au.dk

Johannes Schöning is professor of human-computer interaction science at the University of Bremen. His main research interests lie at the intersection between human-computer interaction (HCI), geographic information science, and ubiquitous interface technologies. schoening@uni-bremen.de

Harald Reiterer is professor at the computer and information science department of the University of Konstanz. His main research interests include fields of human-computer interaction such as interaction design, usability engineering, and information visualization. harald.reiterer@uni-konstanz.de

Christian Holz is a researcher in human-computer interaction at Microsoft Research, Redmond. cholz@microsoft.com

Insights

- We synthesize insights from three conference workshops about the current state of research on interacting across surfaces and devices.
- In particular, we summarize seven key research challenges and opportunities, including: designing for scale and interoperability, enabling easy configuration of device ecologies, and addressing further application domain and use cases.