Advances and Challenges in Ad-hoc Mobile Spatial Tracking for Seamless Interaction across Commodity Devices

Haojian Jin
Yahoo Labs
Sunnyvale, CA
haojian@yahoo-inc.com

Christian Holz
Microsoft Research
Redmond, WA
cholz@microsoft.com

Second International Workshop on Interacting with Multi-Device ecologies "in the wild"

Abstract
In this workshop paper, we assess the progress and lessons learned in developing *ad-hoc* cross-device tracking for mobile *spatial* interaction and point out the challenges existing systems still face. We identify the next steps that are necessary to bring truly fluent cross-device interaction to commodity devices. We finish with a discussion of the emerging opportunities in today's and future mobile devices as well as infrastructure systems that will facilitate accurate ad-hoc tracking on commodity devices.

Author Keywords
Ad-hoc mobile tracking, commodity devices, BLE.

ACM Classification Keywords
H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous;

Introduction
To enable users to interact seamlessly across devices, involved devices need to be aware of surrounding devices. Tracking surrounding devices thereby needs to be ad-hoc, such that users can continue an activity on one mobile device *seamlessly* on another. A big challenge of such tracking is that users interact with devices in their *natural* space, yet current devices lack a no-

tion of this ‘natural’ physical space. To support seamless interaction, devices thus need to understand this 3D space, which is particularly important when multiple users and devices are involved.

Unfortunately, current commodity mobile devices detect only the *presence* of surrounding devices to enable cross-device interaction (e.g., iOS 8 Continuity, Android Wear notifications), typically by analyzing the BLE signal strength to infer the proximity of other devices. This works particularly well if devices are held close.

However, when interacting across several present devices, apps often have to resort to showing a list of devices from which a user must pick. Ideally, though, users would interact naturally—much like they interact with physical objects in their environment.

Recent research projects have made progress in bringing ad-hoc tracking to mobile devices. They often integrate one or more signal types to enable UI apps to detect the *3D location* of surrounding devices. In this workshop paper, we first review the progress in mobile tracking and assess assumptions and shortcomings. We then outline the challenges and directions for future mobile commodity tracking systems.

Advances in ad-hoc tracking

Since most interaction across devices takes place on *mobile* devices, tracking systems need to work ad-hoc. This insight has brought researchers’ attention to the built-in sensors in today’s commodity devices.

Camera-based tracking on mobile devices
Much related work has used the video feed from the camera for ad-hoc tracking. For example, TouchProjector observes the screen contents of other stationary devices’ screens to infer their relative positions [1], but requires a constant visual connection to other screens . Orienteer requires both mobile devices to observe a shared view for registration, such as users’ shoes [4].

While the camera is a strong sensor for vision-based tracking, existing vision technologies are mostly developed for stationary rather than mobile cameras. The mobile nature and suboptimal camera position for mobile vision requires different solutions. Moreover, finger occlusion and table placement will further limit the robustness of ad-hoc camera tracking.

Inertial Motion Units (IMU) for mobile tracking
IMU sensors are the most responsive built-in sensors, which make them suitable for gesture recognition [9] and rough layout inference [6]. However, most mobile devices are not equipped with high-quality IMUs, causing sensor drift and rendering dead reckoning unusable for ad-hoc 3D tracking. To avoid drift, Tracko [9] integrates *temporary* dead reckoning and local coordinate transformation to prevent errors from sensor drift.

Audio processing for ad-hoc mobile tracking
Recent commodity mobile processors can process audio in real-time. BeepBeep [14] is one of the first projects that uses this to determine round-trip *distances* between devices. Tracko [9] establishes 3D tracking based on multiple audio units, producing high accuracy 3D tracking across devices, but is subject to noisy environments and restricted to limited interaction ranges.

Mobile radio to enable device tracking
Radio sensors (e.g., Wifi, GSM, BLE) are common on on today’s mobile devices. Many systems use the signal strength as an indicator to estimate the distance to remote devices. But since these radio sensors are not

designed for ranging purpose, predictions are inaccurate up to several meters (e.g., Tracko [8]). However, using better radio sensors could substantially increase the accuracy, such as QSRCT radio nodes [12].

Towards truly ad-hoc tracking in the wild

While researchers have examined ad-hoc tracking for a long time, it is the new technologies and advances in *commodity* hardware that bring about opportunities that let us think about mobile tracking in new ways—ways that were not possible previously.

Higher-quality built-in sensors

Future devices will pack higher-quality sensors, such as multiple speakers for communication. The popularity of mobile music has brought stereo speakers to many recent phones; the latest iPad Pro even has four speakers. Ad-hoc mobile tracking will benefit from these developments, such as enabling watches or other small devices with only a microphone to be tracked in 3D.

New sensors and low-cost hardware accessories

One potential way to improve commodity ad-hoc sensing is by integrating additional novel sensors into commodity devices—sensors that are expensive now, but will become low-cost through mass production. For example, Google's Tango understands its environment by using special-purpose depth cameras in a tablet.

Cameras on current mobile devices are under-utilized for tracking due to their position and limited field-of-view, which limits the tracking area. One step forward is SurroundSee [16], which is a mobile omni-directional camera that enables peripheral vision around the device to augment daily mobile tasks. We expect that more low-cost hardware, such as camera filters and phone cases [2] will improve tracking significantly.

Internet of Things and Wearables

A plethora of small devices is currently emerging, each dedicated to accomplishing a small task. While such Internet-of-Things devices typically connect using Bluetooth low energy, they often require knowledge of *where* they have been deployed. Wearables face a similar challenge: activity trackers would substantially benefit from an *spatial* awareness of where the user chooses to wear or carry them. For example, WristQue [13] adjusts the local heating and cooling system depending on its (static) location inside a building. We think that emerging systems for ad-hoc tracking, such as Tracko [9], will bring rich capabilities to smart devices that adapt to changed locations in smart environments..

Ubiquitous infrastructure

The recent development of indoor positioning systems makes tracking systems in the infrastructure ubiquitous, which also benefits ad-hoc tracking on commodity devices. WiFi-SLAM [5] determines the physical location of a mobile device based on wireless signal strengths from access points in the environment. Chung et al. [3] and Epsilon [11] explore the space of magnetic filed positioning and light-based positioning. Tracking infrastructure also helps further advancing the relative positioning on current systems to absolute positioning.

Convenient cross-device user authentication

Identification tokens, such as smartphones or wearables increasingly aid users in authentication with local and remote systems. Ad-hoc spatial device tracking will allow current implementations to increase their security by seamlessly ensuring that such identification tokens are close-by [7] and have the potential to fundamentally change scenarios in which multiple users interact with another device simultaneously, tracking and correctly associating all input with a particular user [8].

Development tools

The development in ad-hoc cross-devices is more complex than single device interactions as the developers need to debug on multiple devices at the same time. The efforts to develop ad-hoc cross-device interaction increase exponentially as the number of involved devices increases. We are looking forward to more software development tools like HuddleLamp [15].

Error-prone correction user interface

Tracking accuracy may be not perfect all the time. Thus, the ideal user interface should be able to handle errors implicitly. For example, Corona [10] uses implicit behavior to correct predictions. Depending on the accuracy of the tracking systems, there should be adaptive interfaces for different contexts. Balancing the tracking accuracy and the interactions it can enable can be a promising future research direction.

Conclusion

The recent technology innovations on hardware and software make ad-hoc tracking possible on mobile devices. Developments in this area are still early stage yet promising and indicate that they will impact users' future interactions profoundly. We analyzed the challenges of existing approaches and offered potential paths that we think will allow current systems to make big leaps forward. We believe that in 10 years with next-generation sensors and processors, ad-hoc tracking will be fully mobile and part of commodity devices and operating systems. Future systems will seamlessly integrate 3D ad-hoc tracking much like GPS today.

References

1. Boring, S. et al. Touch projector: mobile interaction through video. Proc. CHI '10, 2287–2296.
2. Butler, A. et al. SideSight: multi-"touch" interaction around small devices. *Proc. UIST '08*, 201–204.
3. Chung, J. et al. Indoor location sensing using geo-magnetism. *Proc. MobiSys '11*.
4. Dearman, D. et al. Determining the orientation of proximate mobile devices using their back facing camera. *CHI '12*.
5. Ferris, B. et al. WiFi-SLAM Using Gaussian Process Latent Variable Models. *Proc. IJCAI '07*.
6. Goel, M. et al. SurfaceLink: using inertial and acoustic sensing to enable multi-device interaction on a surface. *CHI '14*.
7. Holz, C., Bentley, F. On-Demand Biometrics: Fast Cross-Device Authentication. *Proc. CHI '16*.
8. Holz, C., Knaust, M. Biometric Touch Sensing: Seamlessly augmenting each touch with continuous authentication. *UIST '15*.
9. Jin, H. et al. Tracko: Ad-hoc Mobile 3D Tracking Using Bluetooth Low Energy and Inaudible Signals for Cross-Device Interaction. *UIST '15*.
10. Jin, H. et al. Corona: Positioning Adjacent Device with Asymmetric Bluetooth Low Energy RSSI Distributions. *UIST '15*.
11. Li, L. et al. Epsilon: A visible light based positioning system. Proc. NSDI 14, 331-343.
12. Marquardt, N. et al. Crossdevice interaction via micro-mobility and f-formations. Proc. UIST '12.
13. Mayton, B.D. et al. WristQue: A personal sensor wristband. Proc. BSN '13, 1–6.
14. Peng, C. et al. BeepBeep: a high accuracy acoustic ranging system using cots mobile devices. *SenSys '07*.
15. Rädle et al. HuddleLamp: Spatially-Aware Mobile Displays for Ad-hoc Around-the-Table Collaboration. *ITS '14*.
16. Yang, X.D. et al. Surround-see: enabling peripheral vision on smartphones during active use. *UIST '13*.