B. Finley^{a,*}, T. Soikkeli^a

^aDepartment of Communications and Networking, Aalto University, Otakaari 5, Espoo, Finland

Abstract

The increasing number of users with multiple mobile devices underscores the importance of understanding how users interact, often simultaneously, with these multiple devices. However, most device based monitoring studies have focused only on a single device type. In contrast, we study the multidevice usage of a US-based panel through device based monitoring on panelist's smartphone and tablet devices. We present a broad range of results from characterizing individual multidevice sessions to estimating device usage substitution. For example, we find that for panelists, 50% of all device interaction time can be considered multidevice usage.

Keywords: Mobile Device, Multidevice Usage, Mobile Session, Device Substitution

1. Introduction

Today a growing number of users own multiple mobile devices including, for example, a smartphone, tablet, and smartwatch. This proliferation of devices can cause significant changes in usage behavior both in terms of individual sessions and general usage. For example in general, users can shift certain activities from an existing device (i.e. smartphone) to a new device (i.e. tablet), or the new device can prompt totally novel usage. While inside single sessions, users can use multiple devices either in combination or simply simultaneously, and these complex sessions can be described as so called multidevice usage sessions.

Understanding these types of changes is increasingly important for a variety of stakeholders in the mobile ecosystem. For example, mobile application developers are interested in improving the quality of activities that start on one device and continue on another device. A concrete implementation is already seen in web browsers that can sync open web browser tabs across devices¹. Whereas mobile advertisers are interested in understanding how users divide their time between mobile device types to help guide mobile advertising strategies. For example, advertisers can understand the relative ad inventories (number of available impressions) between specific device types of users by considering usage times and typical ad refresh rates.

However, previous studies on mobile device usage have not thoroughly examined this multidevice emergence for a variety of reasons. Many mobile device usage studies primarily tracked only a single device type (typically a smartphone) and often only captured a single mobile platform (typically Android) [1, 2, 3, 4]. Several studies have captured multidevice usage but have used interview, survey, or diary collection methods [5, 6, 7, 8, 9] that are susceptible to recall bias [10]. Other studies have captured multidevice usage but only from a single service provider view point such as multidevice search [11, 12]. Furthermore few studies have explicitly investigated device type substitution (shifting usage to a new device or novel usage of new device) and again these studies typically rely on survey or diary methods [13].

In contrast, we examine multidevice (smartphone and tablet) usage of a US based user panel from February 2015. Our data collection is device based (panelists installed custom device based monitoring

^{*}Corresponding author

Email addresses: benjamin.finley@aalto.fi (B. Finley), tapio.soikkeli@aalto.fi (T. Soikkeli)¹https://support.apple.com/en-us/HT202530

Preprint submitted to Pervasive and Mobile Computing

applications) and multi-platform (Android and iOS). Furthermore, we study device type substitution by utilizing a supplementary (smartphone only) user panel also from February 2015. To our knowledge, this is the first study of a multidevice user panel with device based monitoring.

We make the following concrete contributions in this work.

- We present a simple method for grouping app sessions into multidevice usage sessions based on Allen's interval algebra.
- We describe the statistical properties of these multidevice sessions on both an overall and per user level.
- We present and utilize a novel method for studying the diversity of temporal patterns in multidevice sessions through a matrix representation.
- We examine if users use their smartphones differently during multidevice sessions.
- We examine the usage substitution effects between smartphones and tablets by utilizing a complementary panel of users without a tablet device.

The remaining sections of the paper are organized as follows. Section 2 discusses related work and introduces Allen's interval relations, while Section 3 introduces the panel datasets. Section 4 details the different session definitions and the aggregation of individual app sessions into usage sessions. Section 5 describes and discusses the primary results. Section 6 discusses some general implications and the replicability and generalizability of the results and finally Section 7 presents the conclusions and future work.

2. Background

2.1. Related Work

Related work can be primarily divided into single device usage studies that use a similar device based monitoring method, multidevice studies that also characterize multidevice usage, and multidevice interaction enablers and frameworks.

2.1.1. Single Device Usage Studies with Device Based Monitoring

Eagle and Pentland [14] pioneered device based monitoring with mobile phones in their Reality Mining study, though they did not explicitly focus on studying mobile device usage. Böhmer et al. [1] analyzed mobile application usage in one of the first large scale mobile usage studies. They used a similar concept of a mobile app session described by an app moving to the foreground or background of the device. Furthermore, they extended the concept to describe an app chain (sequential app sessions with a maximum timeout value between sessions of 30s). These app chains are similar to our concept of single device usage sessions. Soikkeli et al. [4] had a similar approach to [1] with 1s and 30s maximum timeout values for app chains, or usage sessions in their terminology. Falaki et al. [2] also based their analysis on foreground application sessions with an extension to interaction sessions. However, an interaction session was defined as continuous screen-on time period (essentially a timeout value of 0s). Oliver [3] also studied this type of interaction session.

All the mentioned studies included only smartphone usage, in some cases because tablets were not yet widespread. Also, the data were usually from only a single platform, such as Android, Windows Mobile, Symbian, or BlackBerry. In some cases, iOS devices at the time had limitations regarding background measurements and the release of applications in the app store [1].

2.1.2. Multidevice Usage Studies

Multidevice studies can be further subclassified by the utilized data collection method or methods.

Interviews and Direct Observation. Many early multidevice studies employed interviews or direct user observation. These methods are less scalable and accurate than device based monitoring but can provide contextual information such as why a user switches between devices during specific activities.

Dearman and Pierce [5] examined multidevice usage of a panel of industry and academic users through semi-structured interviews. They highlighted, for example, that users use multiple devices for a variety of reasons including differences in device form factor and portability. Though tablet devices were not yet widely used at the time of the study. Similarly, Oulasvirta and Sumari [15] and Santosa and Wigdor [16] also performed semi-structured interviews and user observation studies of multidevice users and workflows in the workplace context. Santosa and Wigdor [16] especially highlighted the inadequacies of current methods for handing parallelism and data fragmentation in multidevice use cases.

Finally, Matthews et al. [17] performed semi-structured interviews of industry smartphone users and examined the interplay of smartphone and PC usage. They found that, in contrast to basic cellphones, smartphones prompt both novel usage and substituted usage from PCs as functionality between smartphones and PCs partially overlaps. Furthermore, they report that smartphone attributes such as low initial access time gave smartphones an advantage over PCs for certain tasks in certain contexts (such as checking email). Thus their results are similar to the device substitution we explore between smartphone and tablet devices (wherein, for example, significant mobile game usage is substituted to larger screen tablets).

User Completed Electronic Diaries. Other multidevice studies employed user completed electronic diaries in which a user tracks their own multidevice usage. This method is more scalable than interview or observational methods but still less accurate than device based monitoring as the approach relies on users, that are typically susceptible to recall bias [10], to complete the diaries. However, similar to the interview method, user completed diaries can collect contextual information such as why a user performs a specific task.

Muller et al. [9] analyzed smartphone and tablet usage of a user panel through user completed electronic usage diaries. Similar to our study, their panel includes users that tracked both their smartphone and tablet device usage. They also report on the types and frequencies of secondary activities performed during tablet and smartphone use. Similarly, Google [6] used usage diaries and surveys to study the multidevice usage of a large US based panel. The research found that over 90% of panelists reported starting a task on one device and continuing the task on another device. However relatively few panelists (between 3-8% depending on activity) reported starting an activity on a smartphone and later continuing on a tablet or the reverse sequence.

Combinations of Methods and Other Methods. Finally, a few studies have used a combination of these methods or examined multidevice usage from a service provider rather than user centric viewpoint.

Jokela et al. [8] investigated multidevice usage of a small Finland based panel through both electronic diary and interview methods. They included a broad range of devices including tablets and, for example, home media centers in their multidevice usage definition. The investigation found the users reported significant sequential and parallel related but little parallel non-related multidevice usage, potentially because users do not easily realize non-related parallel usage. Contrastingly, device based monitoring will always detect parallel usage but classification into related or non-related categories requires additional information and analysis. Similarly, Microsoft [7] also used electronic diary and interview methods to study multidevice usage. The research studied a mix of smartphones, tablets, e-readers, gaming consoles and laptops and classified so called multi-screening behavior across these devices as sequential, simultaneous or separate.

Montanez et. al [11] and Wang et al. [12] studied cross device search in the context of a major US web search provider. Wang et al. found, for example, that about 15% of device switching between searching involves contiguous tasks, in other words sequential multidevice sessions.

Hintze et al. [18] analyzed both smartphone and tablet usage through device based monitoring via the large and well known Device Analyzer dataset. Device Analyzer [19] is a large scale dataset based on a popular Android monitoring app. The dataset contains both smartphone and tablet devices however correlating ownership of smartphone and tablet devices to a single individual is not possible thus limiting multidevice session analysis. Furthermore individual application names are not available, thus further limiting analysis.

precedes	meets	overlaps	finishedBy	encloses	starts	equals	startedBy	enclosedBy	finishes	overlappedBy	metBy	precededBy
A	AB	AB	AB	AB	AB	A B	AB	A	A B	B	B	B

Figure 1: Allen's Time Interval Relations

2.1.3. Multidevice Enablers and Frameworks

A logically related and currently very active area of study is the building of enablers and frameworks for more convenient multidevice usage.

Taivalsaari et al. [20] discuss liquid software as a general manifesto of guidelines that multidevice interaction should follow. Hamilton and Wigdor [21] present Conductor; a framework for the construction of multidevice applications, whereas Desrulle and Gielen [22] discuss the high level architectures of such systems in the context of web applications. Finally, Schreiner et al. [23] present a versatile web based multidevice app framework that is agnostic to, for example, the underlying device to device communication method. Relationally, our research looks to help find potentially useful (app) targets for such enablers and to help describe the current state of multidevice usage so the value of these enablers can be estimated.

2.2. Allen's Time Interval Relations

Allen's thirteen basic time interval relations are the building blocks of Allen's Interval Algebra [24]. The relations are *distinct* (a pair of intervals are described by one and only one relation), *exhaustive* (any pair of intervals can be described by one of the relations) and *qualitative* (no numeric time spans are required). Figure 1 shows all of the thirteen basic relations that two finite time intervals a and b can have. In the figure, time is expressed from left to right. So, for example, in the first cell interval a starts and ends before interval b starts. In the second cell interval b starts exactly as interval a ends. Intuitively, six of the relations can be described as converses of six other relations. For example, the converse of a meets b is b metBy a.

In this paper we often use the notion of a session. In the context of, for example, smartphone usage an app session can be defined as a period of application usage with a clear start and end time. This type of session is equivalent to Allen's notion of a finite time interval and thus we can use Allen's relations in defining and examining multidevice usage. The relations formalize the temporal relationships between (usage) sessions of one or more devices.

This work is the first multidevice paper, as far as we know, to exploit such a formalization and thus to bring this foundation into the multidevice usage area.

3. Datasets

The two primary datasets are subsets of a large United States based user panel organized by Verto Analytics² in February 2015. Panelists were recruited online and were given an initial recruitment survey that asked about the devices they own. Panelists were instructed to install custom monitoring applications to all of their applicable devices (in other words, the devices they reported owning including smartphone, tablet, and PC). The monitoring applications logged events such as an app moving to the foreground or background of the device. Only panelists that installed the monitoring applications to all their applicable devices were considered for the panel. All panelists were paid for participation. All analyzed user data was anonymized with no personally identifying information.

The first dataset (hereafter dataset 1) consists of one month (February 2015) of smartphone and tablet application usage data from the 65 panelists with both a smartphone and tablet device. Only active panelists were considered valid and included in the subset. We define a panelist as active if for both of their devices the length between their first usage of the month and last usage of the month is at least 23 days. The subset

²http://vertoanalytics.com/

includes panelists with devices from both major mobile platforms: Apple iOS and Google Android. Specifically, 46 panelists had both an Android smartphone and tablet, 10 panelists had both an iOS smartphone and tablet, 5 panelists had an iOS smartphone and Android tablet, and 4 panelists had Android smartphone and iOS tablet.

The second dataset (hereafter dataset 2) consists of one month (February 2015) of smartphone application usage data from 496 panelists with a smartphone but not a tablet device. In other words, these panelists reported on their panel recruitment survey that they do not own a tablet device. Again only active panelists were considered valid and included in the subset and the subset includes both major mobile platforms. 402 panelists had Android smartphones and 94 panelists had iOS smartphones.

In terms of representativeness, the large panel was recruited with the purpose of obtaining a nationally representative panel and thus is relatively diverse. However all opt-in panels inherently use non-probability sampling and thus representativeness is a concern³. We refer to Hays et al. [25] for further discussion. Furthermore, as mentioned, the recruitment process utilized a recruitment survey that screened potential panelists to improve the demographic and technographic match between the accepted panelists and the population (known as a quota-sampling approach).

For reference we provide a summary of panelist demographic and technographic data for each dataset along with demographic data for US smartphone users in general in Table 1. The clearest demographic discrepancies are that the datasets over represent females and users with lower household incomes. Thus these factors should be considered when generalizing the results. We discuss our overall view of generalizability in panel based studies in Section 6.

In terms of the types of monitored device types, personal computer usage data was also collected from panel users but considered out of scope for this paper. Though analysis of PC usage in a multidevice context would also be valuable, a significant issue is that the simple foreground app session definition is not directly transferable to the personal computer context. For example, PC users often have multiple viewable app windows open simultaneously. Furthermore, the limited mobility of laptops and desktops means that these devices often remain continuously connected to an AC socket, thus users can often leave their screens on even when not interacting with the device (for example by disabling sleep mode). Therefore, more advanced user activity detection would be needed.

4. Usage session

4.1. Session Definitions

In order to study mobile usage sessions we first define several important concepts we will use throughout the analysis. The definitions are partly based on Allen's time interval relations described in section 2.2. First, we define an *app session* as a time interval starting with an application moving to the foreground of the device and ending with the application moving out of the foreground (either replaced by a different app or screen off). This app session definition has been used in previous literature [1, 2, 4]. Next, we define a *single device usage session* as a collection of app sessions on a single device with a maximum timeout value between two sequential app sessions. This definition has also been used previously [1, 4], but in comparison to previous work, the Allen's relations approach we adapt to usage session construction (see section 4.3) offers a more solid foundation for characterizing usage sessions and their relations, and more importantly can be easily extended to multidevice usage.

Thus we describe our definitions in terms of Allen's relations. Since by definition single device app sessions are sequential, only **precedes** and **meet** of Allen's relations (and their respective converse relations) are applicable. If two app sessions **meet**, the sessions will always belong to the same usage session. If an app session **precedes** or is **precededBy** another app session within the maximum timeout, then the app session will also belong to the same usage session. Henceforth, we denote this relation as **precedes** (**precededBy**) within time-window (abbreviated as **precedes** (**precededBy**) within **TW**). Thus

 $^{^{3}}$ We note that in contrast to panels were panelists take surveys, the device based monitoring collection method is not generally susceptible to, for example, false/fake answers, careless answers, or repeatedly giving the same answer.

	Dataset 1	Dataset 2	US Smartphone Users ^a
Mean Age (Years) ^b	38.53(11.91)	37.06(12.63)	41.30 (15.08)
Gender (% Male)	23.08	26.61	50.08
Education (% w/ Some College or Less)	56.92	63.10	53.04
Marital Status (% Married)	44.62	41.33	48.96
Household Income ($\% < 50 \text{K USD}$)	47.69	65.93	40.72
Mean Household Size	2.97(1.54)	2.95(1.50)	3.05(1.61)
Mean Children in Household	0.94(1.06)	0.93(1.21)	0.74(1.27)
Race (% White)	72.31	70.97	71.53
Mean Smartphone Display Size (inches)	4.61(0.71)	4.58(0.60)	-
Mean Smartphone Display Density (ppi) ^c	358.00(22.79)	343.86(84.67)	-
Mean Smartphone Total CPU Frequency (GHz) ^d	5.57(3.33)	5.50(3.20)	-
Mean Time Since Smartphone Release $(years)^{e'}$	2.06~(1.11)	1.92(1.04)	-
Mean Tablet Display Size (inches)	8.59(1.27)	-	-
Mean Tablet Display Density (ppi) ^c	203.89(57.86)	-	-
Mean Tablet Total CPU Frequency (GHz) ^d	3.58(2.04)	-	-
Mean Time Since Tablet Release (years) ^e	2.05(0.86)	-	-

^a US smartphone user demographic data is from Pew Research survey (June-July 2015, subpop with smartphone n=1327) [26]. The survey utilizes weighting to population parameters of census data to create nationally representative results (refer to [27]). We note that Verto Analytics also performs its own national surveys, we utilize the Pew Research survey only for brevity.

^b All mean values also include standard deviations

^c In pixels per inch

^d Total frequency of CPU as number of cores multiplied by frequency per core

^e Time from approximate release month of device model until Feb 2015

through a combination of these four relations multiple *app sessions* can be grouped into *single device usage sessions*.

For multiple devices (of a single user in our case) all of Allen's relations are applicable. As mentioned **precedes** and **meet** (with their respective converse relations) describe sequential usage, while the remaining relations describe (at least partly) simultaneous usage of the devices. For constructing a *multidevice usage session* we follow a two-step process. First, we apply the single device usage session definition separately for each device. Second, we examine the relations of the usage sessions of the multiple devices. If the single device usage sessions are (at least partly) simultaneous, **meet** (are **metBy**) or **precede** (are **precededBy**) **within TW** then they belong to the same *multidevice usage session*. Figure 2 illustrates this process with smartphone and tablet sessions (two devices).

Finally, we define several descriptive session types that we will use throughout the paper. Specifically, each single device usage session (and by extension each app session contained within a session) can be classified along two distinct binary dimensions: device type (*smartphone* or *tablet*) and containment in a multidevice usage session (if contained in a multidevice session then *mixed* or if not then *pure*).

4.2. Maximum Timeout Value

In both the single and multidevice cases, the maximum timeout value (used in session construction) is an important parameter that affects the eventual interpretation of the results. Ideally the timeout value should be a moderate value that accurately reflects how users psychologically divide their daily device usage. However, as mentioned in section 2.1, various authors have used different values with no best practice prevailing, therefore we briefly examine the effects of a range of different timeout values.

Figure 3a details the mean number of different types of usage sessions per user as a function of timeout value. Interestingly, the mean number of multidevice sessions is relatively stable for timeout values of 1-1000s, indicating that in this range multidevice sessions are primarily subsuming time-adjacent pure smartphone and tablet usage sessions. Whereas after 1000s, the significant decrease in multidevice sessions suggests that time-adjacent multidevice sessions themselves are merged. Figure 3b illustrates the mean app sessions per usage session per user as a function of timeout value. The increasing number of app sessions per usage session further supports the previous indications.

In terms of selecting a practical timeout value for analysis purposes, we select a mid-range compromise value of 60s that preserves a large number of pure smartphone and tablet usage sessions and does not



Figure 2: Multidevice usage session construction - Step 1: Apply usage session definition to each device's app sessions (merge A,B,C to H and E,F to J; D to I and G to K are degenerate single app session cases), Step 2: Collapse simultaneous, meeting and preceding within time-window usage sessions into multidevice usage sessions (merge H,J to L and I,K to M).



Figure 3: a) Mean Usage Sessions Per User by Timeout Window, b) Mean App Sessions Per Usage Session (Per User) by Timeout Window

indicate significant merging of multidevice sessions. Furthermore, we use the same value for creating both single device usage sessions and multidevice usage sessions. Though an argument could be made that switching devices in sequential multidevice sessions might often take longer than 60 seconds and thus a longer timeout value could be used. However, we leave this analysis for future work.

4.3. Usage Session Construction

In order to analyze smartphone, tablet and multidevice usage on a session level we first construct the usage sessions. We follow the session definitions of section 4.1. Table 2 details statistics related to the usage session construction process of dataset 1.

In the single device case one app session can have from zero to two relations: no relation to a preceding or following app session, a single relation to a preceding or a following app session, or relations to both preceding and following app sessions. For Table 2 we count all these relations. Table A.12 in Appendix A details similar statistics for dataset 2.

Table 2 also breaks down the types of Allen's relations observed between the smartphone and tablet usage sessions contained in multidevice usage sessions. The relations are written with the smartphone session preceding the relation and the tablet session following. In other words, we can read the statistics as follows: smartphone session [relation] tablet session. In the majority of cases (over 61 %) a smartphone session is **enclosedBy** a tablet session. The second most frequent case (with 11 %) is that a smartphone session **encloses** a tablet session. A majority of the rest of the cases fall relatively evenly between **overlaps** and **precedes within TW** and their converses.

Smartphone usage	Tablet usage	Multidevice usage
330523	63239	97620
78798	12708	18136
-	-	4517
8.16 %	4.65 %	6.16 %
42.07 %	45.44 %	0.10 %
-	-	8.50 %
-	-	0.04 %
-	-	10.70 %
-	-	0.12 %
-	-	0.00 %
-	-	0.11 %
-	-	61.40 %
-	-	0.11 %
-	-	7.40 %
41.60 %	45.26 %	0.11 %
8.16 %	4.65 %	5.27 %
	Smartphone usage 330523 78798 - 8.16 % 42.07 % - - - - - - - 41.60 % 8.16 %	Smartphone usage Tablet usage 330523 63239 78798 12708 - - 8.16 % 4.65 % 42.07 % 45.44 % - -

Table 2:	Usage	session	construction	statistics
10010 2.	Ubugu	BCBBIOII	construction	5000150105

5. Results

In this section, we present and discuss the main results of our analysis and where appropriate compare these results to prior work.

5.1. Basic Descriptive Statistics

5.1.1. Dataset Level Statistics

In this section we study the dataset level basic characteristics of the single device and multidevice usage sessions. In other words, the statistics describe the aggregate of all sessions from all users of dataset 1. Table 3 details descriptive statistics related to usage session lengths and the number of app sessions per usage session (for dataset 2 see Table A.13 in Appendix A). We calculate the statistics separately for all smartphone usage sessions (combined pure and mixed), all tablet usage sessions (combined pure and mixed), pure smartphone usage sessions, pure tablet usage sessions and multidevice usage sessions.

On general level, we find that tablet usage sessions are significantly longer than smartphone usage sessions but that the number of app sessions in a single device usage session is similar for both smartphones and tablets. Whereas, multidevice usage sessions are about 10 times longer and have about 5 times as many app sessions as compared to the single device usage sessions.

These reported mean and median values are indicative of the differences between sessions types but do not illustrate the full distributions very well. Figures 4b and 4c show empirical distribution functions of usage session lengths and number of app sessions per usage session, respectively, for the different types of usage sessions. In terms of session lengths, the distributions show clear peaks for smartphone and especially for tablet sessions. A small peak appears at 10s and is associated with certain Samsung devices' touch user interface (known as TouchWiz). A larger more noticeable peak appears at around 15-17s and is associated with iOS applications not properly identified by the monitoring application. The difference in the share of these particular iOS sessions between pure and mixed tablet sessions is the main reason for the relatively large difference in session length values between these types. For brevity, we leave the further examination of these app specific anomalies for future research.

In terms of number of app sessions, only 8% of all multidevice usage sessions have exactly two app sessions (the minimum for multidevice usage sessions by definition) while half of multidevice usage sessions have more than 10 app sessions. Thus app and interaction designers likely need to consider complex use cases of multiple devices well beyond the degenerate case of two simple sequential or simultaneous app sessions (as these degenerate cases are rare in practice). We further illustrate the wide range of multidevice app usage patterns in Section 5.2.

Table 3 also shows how the usage is divided among the session types based on different definitions of usage including number of app sessions, number of usage sessions, and interaction time. Most interestingly when the division is based on interaction time, multidevice and single device usage divide almost equally

Table 3: Basic statistics on the whole dataset level^a

	$\begin{array}{l} {\rm Smartphone \ - \ all} \\ {\rm N} = 78798 \end{array}$	Tablet - all $N = 12708$	Smartphone - pure $N = 63930$	Tablet - pure N = 6871	$\begin{array}{c} \text{Multidevice} \\ \text{N} = 4517 \end{array}$
Usage session length (s)					
mean	312	1461	245	506	4241
median	61	197	56	114	1147
std. dev.	1509	8154	829	1074	14284
App sessions per usage					
session					
mean	4.13	4.95	3.96	3.97	21.61
median	2.00	2.00	2.00	2.00	10.00
std. dev.	6.46	7.92	5.90	5.98	45.47
Percentage of all usage					
based on app sessions	83.5 %	16.5 %	67.0 %	7.2 %	25.8 %
based on usage sessions	85.9 %	14.1 %	84.9 %	9.1 %	6.0 %
based on interaction time	55.8 %	44.2 %	40.0 %	9.2~%	50.9~%

^a N is the number of usage sessions. Smartphone (all) and Tablet (all) indicates all smartphone and tablet usage or app sessions, whereas Smartphone (pure) and Tablet (pure) indicates only pure smartphone and tablet usage or app sessions.



Figure 4: Cumulative distributions (the inter-session time axis is shifted by 60s for readability)

with 50.9% multidevice usage and 49.1% single device usage. In other words, over half of all interaction time is considered multidevice usage. This reinforces the significant prevalence of multidevice usage for multidevice users and the potentially large impact of optimizations for multidevice usage.

If we examine the device types carefully, we find that a tablet usage session is much more likely to be part of a multidevice usage session, when compared to smartphone usage sessions. In fact, over 90% of tablet interaction time and 93% of tablet usage sessions are associated with multidevice usage. Comparatively, Google [6] survey results, for example, indicate only 35% of tablet sessions also include smartphone usage. Similarly, Muller et al. [9] reports device multitasking during tablet usage in only a small fraction (3.4%) of user reported usage incidents. Though these studies are not fully comparable, the large differences in results still reinforce the importance of utilizing device monitoring methods in addition to self reporting methods.

Finally, in terms of temporal distributions, Figure 5 shows hourly distributions of the different types of usage sessions. Smartphone usage is distributed relatively uniformly across the office hours with a small shift towards the evening. Whereas tablet and multidevice usage is focused in the late evening. These distributions support the notion of smartphones being the ubiquitous all-day work horses of mobile device usage, while tablets are related to leisure time and the home context. Furthermore, the distributions also support and validate previous work based on self reported diaries such as [9] that illustrate similar distributions.

5.1.2. Per User Level Statistics

In the previous section we examined the usage sessions on the entire dataset level, in other words all sessions of all users were aggregated. Table 4 shows similar statistics, again for dataset 1, but based on a per



Figure 5: Hourly Distributions

user level (for statistics for dataset 2 see Table A.14 in Appendix A). More specifically, the session lengths and number of app sessions per usage session are based on per user medians and the usage sessions per day and interaction time per day are themselves per user statistics.

The per user results reveal considerable diversity across users also in the multidevice usage case. Session length medians show an order of magnitude difference between users and in the case of app sessions per usage session from 5 to 10 fold differences. Thus we extend to the multidevice context previous studies, for example [18], that report significant diversity in usage across users.

In terms of per user interaction time, we find that typically no single device dominates but users spend considerable time with both their smartphone and tablet. This relatively equitable time distribution suggests that mobile ad inventories per user will also be roughly equitable depending on ad refresh rates. Thus informing multidevice mobile advertising strategies.

Furthermore, the overall interaction time with both smartphones and tablets is considerably higher than previous studies⁴ (see Table 1 in [18]). These differences are potentially due to increased pervasiveness of these devices over time, differences in the types of users between studies, or a combination of these factors. For reference, the dataset from [18] has users from over 175 countries whereas our dataset has only US users.

5.2. Patterns within Multidevice Usage Sessions

In this section, we study the temporal patterns within multidevice usage sessions. In other words, the configurations of the device usage sessions that form each multidevice usage session.

First, to quantify the temporal similarity between two multidevice usage sessions we need a similarity metric. We use an interval sequence similarity measure known as Interval-Based Sequence Matching (IBSM) [28]. The measure is based on representing sequences of finite intervals as matrices. The similarity between these matrices is then the Frobenius norm of the difference between the matrices. In our case, each matrix represents a multidevice usage session, each matrix row represents a specific device (in our case always two rows for smartphone and tablet, with smartphone always the first row and tablet the second) and each matrix element represents the binary usage state (active or inactive) at a discrete time (a second). In other words, a multidevice usage session with a smartphone session from t = 1 to t = 4 and a tablet session from t = 3 to t = 5 would be represented as:

[1	1	1	1	0]
0	0	1	1	1

In cases where the matrices to be compared are of different lengths then each row of the shorter matrix is re-sized through linear interpolation.

 $^{^{4}}$ For smartphones [18] reports a median daily usage of 82 mins, compared to our median daily usage of 261 mins for dataset 2. For tablets [18] reports a median daily usage of 67 mins, compared to our median daily usage of 115 mins for dataset 1.

	Smartphone - all	Tablet - all	Smartphone - pure	Tablet - pure	Multidevice
Usage session length (s)					
mean	85	463	78	211	2041
median	57	215	54	121	1063
std. dev.	109	623	99	259	2983
App sessions per usage					
session					
mean	2.28	2.73	2.25	2.15	13.03
median	2.00	2.00	2.00	2.00	11.00
std. dev.	1.04	2.00	1.00	1.42	8.95
Usage sessions per day					
mean	44.5	7.15	36.05	3.85	2.59
median	43.13	5.01	34.3	2.82	1.79
std. dev.	22.17	6.78	19.96	4.6	2.38
Interaction time per day					
(min)					
mean	234.21	174.95	146.69	32.57	184.22
median	213.86	115.01	114.78	19.23	121.80
std. dev.	167.89	188.12	97.95	44.37	206.05

Table 4: Basic statistics on per user level^a

^a Usage session length and App sessions per usage session statistics are calculated from the per user medians of these metrics. Smartphone (all) and Tablet (all) indicates all smartphone and tablet usage sessions, whereas Smartphone (pure) and Tablet (pure) indicates only pure smartphone and tablet usage sessions.

The IBSM measure is used instead of Allen's relations because while Allen's relational algebra is a fundamental tool for understanding interval relations, the categorical relations are quite strict. In other words, only a small difference in position of a session relative to another session can result in a totally different relation, even though the two configurations might be considered in a sense similar.

With the similarity measure, we divide the multidevice usage sessions into groups based on the similarity of each session's matrix to well defined artificial prototype matrices. Specifically, we use all 256 configurations of the size 2×4 (row \times column) binary matrix as prototypes and each prototype represents a group (thus we have 256 distinct groups). Importantly since we are representing relatively large matrices, on average 2041 columns, by small matrices, always 4 columns, these prototype groups are rough highly aggregated approximations. We refer to each group by the integer number represented by the binary of the prototype matrix when the second row is concatenated after the first.

The choice of 4-column prototype matrices (rather than for example larger or small prototype matrices) relates to a trade-off over the level of detail that proves useful in understanding common patterns. Specifically, large prototype matrices (for example the same size as the actual matrices) would just represent every actual matrix by it's own prototype, whereas, very small prototype matrices (1 column) would just obfuscate most of the interesting differences between the matrices. Thus the 4-column size represents a good compromise size that maintains a useful level of detail.

Table 5 details the ten most frequent groups of all multidevice usage sessions. Interestingly, we find that the top two groups represent primarily usage of a single device over the length of the session along with sparse usage of a second device. The relatively high frequency of group 15 is expected given that tablet usage sessions are generally longer and smartphone usage sessions are shorter and burstier. However surprisingly group 240 (the bit-wise complement of group 15) is also quite frequent.

With regard to device ordering, we find that group 135 (with a starting smartphone session) is found more as often than the bitwise complement group 30 (with an ending smartphone session). We find similar relationships for other groups and in general we find that, in terms of all multidevice sessions, smartphone sessions tend more often to precede tablet sessions than vice versa. Google [6] also identifies smartphones as the most common starting point for multidevice usage (when considering all sequential multidevice sessions).

In terms of diversity, we find that at least 25 groups (10% of all possible groups) have frequencies of 1% or greater and 138 groups (54% of all groups) have at least one session thus suggesting significant diversity in the configuration of multidevice usage sessions. This significant diversity again reinforces the implication that app and interaction designer should consider a large variety of complex multidevice interactions and

All Mult	idevice Usage Sea	sions	Multidevice Usage Sessions Per User			
Group Number	Group Matrix	Frequency	Group Number	Group Matrix	Frequency	
15	$\begin{bmatrix} 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix}$	24.07%	15	$\begin{bmatrix} 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix}$	25.70%	
240	$\begin{bmatrix} 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}$	10.90%	240	$\begin{bmatrix} 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}$	12.95%	
135	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 \end{bmatrix}$	4.45%	135	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 \end{bmatrix}$	3.62%	
143	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix}$	3.72%	30	$\begin{bmatrix} 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 0 \end{bmatrix}$	3.37%	
30	$\begin{bmatrix} 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 0 \end{bmatrix}$	2.90%	143	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix}$	3.32%	
79	$\begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix}$	2.50%	79	$\begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix}$	2.60%	
195	$\begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix}$	2.24%	31	$\begin{bmatrix} 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix}$	2.36%	
225	$\begin{bmatrix} 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$	2.17%	195	$\begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix}$	2.11%	
31	$\begin{bmatrix} 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix}$	2.17%	225	$\begin{bmatrix} 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$	1.87%	
47	$\begin{bmatrix} 0 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix}$	2.04%	60	$\begin{bmatrix} 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \end{bmatrix}$	1.86%	

Table 5: Most Frequent Prototype Groups of Multidevice Usage Sessions (All Sessions and Per User)

usage in future designs.

To account for differences in the number of multidevice sessions per user, we also calculate the relative frequency of each group for each user and then average these percentages. Table 5 also details the ten most frequent groups by this mean frequency per user. Nine of the top ten groups are the same as in overall analysis. However, we find much smaller differences in device ordering than in the overall analysis. For example, groups 135 and 30 have almost equal mean frequencies of 3.62% and 3.37% respectively. Thus suggesting no preference in device ordering on a user level.

Finally we examine the types of apps used within the most frequent group (group 15) and the second most frequent group (group 240) compared to all other groups. Specifically we calculate the differences in normalized usage times by app category between group 15 and 240 (hereafter G15 and G240) sessions and the complements of group 15 and group 240 (hereafter C15 and C240) sessions for both smartphones and tablets.

For G15, as expected, we find the sparse usage of smartphones to be related to app categories with frequent notifications such as news or social networking (83% and 41% more usage in G15) rather than app categories such as video or entertainment (280% and 87% less usage). Interestingly tough, we find that the usage of tablets in G15 to be related more with gaming (33% more usage) rather than video (77% less usage). Thus game developers should be especially aware that tablet users might often need to interrupt games for frequent smartphone notifications.

For G240, we find the sparse usage of tablets to be related to app categories such as productivity (email and web browsing) (60% more usage in G240) rather than app categories such as games or video (160% and 580% less usage). Whereas, we find that the usage of smartphones in G240 to be related to categories such

as photos and gallery, maps and navigation, or video (85%, 70%, and 62% more usage) rather than app categories such as camera, gaming, or social networking (252%, 76%, and 36% less usage).

Thus we find that although G15 and G240 can be illustrated as inverses, the types of apps commonly used in these sessions are different. In other words, the app categories that are commonly used on smartphones in G15 are not necessarily commonly used on Tablets in G240 and vice versa. Therefore, understanding the dominant device of a multidevice session is likely to be useful in, for example, inferring the types of apps used in that session. This type of inference is useful in cases where app sessions themselves might be inferred from other sources and thus no direct app type information is available.

5.3. Smartphone Usage with and without Simultaneous Tablet Usage

Beyond the interplay within single multidevice usage sessions, we study the diversity of a user's smartphone usage behavior across different types of usage sessions. Specifically we study whether multidevice users use their smartphones differently when they are also using a tablet device in the same usage session⁵. We can formalize this question as comparing smartphone usage behavior between smartphone pure usage sessions vs. smartphone mixed usage sessions.

In practice, our comparison consists of testing for statistically significant differences in normalized usage times⁶ of both app categories and individual apps between the session types. In addition, since differences in app usage times are affected by the time of day and different session types are more prevalent during certain times, we limit the analysis to the typical evening period between 17:00 and 24:00.⁷

The examined app categories are custom categories that unify the category differences between application stores of different platforms (iOS and Android). In the analysis we only test app categories that at least 50% of users used. This threshold keeps the sample sizes from becoming too small, thus rendering the tests unreliable. These custom categories are listed in Table 6. The individual apps are a collection of popular apps (again that at least 50% of users used). The apps are listed in Table 7. Refer to Table A.11 in Appendix A for general usage statistics for the app categories and individual apps.

To test for significant differences, we use a percentile bootstrap method for paired data that compares difference scores⁸ of the 20% trimmed mean (section 5.9.12 in [29]). The percentile bootstrap method is selected because the method is robust under non-normality and heterogeneity of variances. Furthermore, we use the 20% trimmed mean as the location measure of the tests because the distributions have significant outliers. Thus the null hypothesis is that the 20% trimmed mean of the distribution of the difference scores of the groups is equal to zero. All tests are performed in R with the bootdpci command from [29]. Furthermore, for significant results (p < 0.05) we report a robust effect size measure, ξ , based on a generalization of explanatory power [30]. We interpret effect sizes $\xi > 0.15, 0.35, 0.50$ as small, medium and large effects respectively, as suggested in section 5.3.4 in [29] and we report the absolute value of effect sizes. All effect sizes were found in R with the yuenv2 command from [29].

In terms of application categories, column C1 in Table 6 details the tests results. We find significant differences in the normalized usage time of several app categories. In all cases we find higher usage time in the smartphone pure sessions. In the cases of Camera, Photo and Gallery, and Video the reduced usage in mixed sessions can be explained by users preferring the larger tablet screen, which is typically optimized for viewing high resolution images. All three categories also have large effect sizes. In terms of individual app usage, column C2 in Table 7 details the tests results. We find significant differences in the usage time of Web Browsers, Email, and Messengers. Again, we find higher usage among the smartphone pure sessions.

Overall, the results suggest that users do use their smartphone differently when they are also using a tablet device in the same usage session. However the usage differences vary significantly for both application categories and individual apps. These differences suggest that platform or app developers could optimize

⁵We note that multidevice users' tablet usage behavior in different situations is also interesting but not as easily studied because certain session types, for example tablet pure usage sessions, are relatively rare in practice. Thus we omit such analysis. ⁶The times must be normalized because the total usage times for these two sessions types are not equal.

 $^{^7\}mathrm{We}$ also performed the analysis without this constraint and found similar results.

 $^{^{8}}$ This difference score approach, as opposed to a marginal distributions approach, tells about the typical differences for a randomly selected user. Refer to section 5.9.9 in [29]

C1 P-values ^a
0.776 PS>MS
$0.036^* \text{ PS>MS} (0.66)$
b
0.122 PS>MS
-
0.670 PS>MS
0.000*** PS>MS (0.71)
$0.004^{**} \text{ PS>MS} (0.38)$
0.000*** PS>MS (0.81)
-
0.432 PS>MS
0.152 PS>MS
0.066 PS>MS
-
$0.026^* \text{ PS>MS} (0.66)$
-

Table 6: Smartphone Usage across Session Types - App Category Tests Results

^a For significant results the significance level, relationship between trimmed means, and robust effect size are also reported. ^b Indicates that not enough users of the category to reach threshold of at least 50% usage.

Table 7: Smartphone Usage across Session Types - Individual App Tests Results

App Name	C2 P-values ^a
Other	0.000*** MS>PS (0.16)
Web Browsers	$0.040^* \text{ PS>MS} (0.19)$
YouTube	_b
Email Apps	$0.002^{**} \text{ PS>MS} (0.41)$
Facebook	0.088 PS>MS
Messenger Apps	$0.008^{**} \text{ PS>MS} (0.12)$

^a For significant results the significance level, relationship between trimmed means, and robust effect size are also reported.

^b Indicates that not enough users of the app to reach threshold of at least 50% usage.

the user experience in cases where, for example, the smartphone can detect or infer whether a tablet is being used simultaneously. Such optimizations would be an logical extension of already existing contextual optimizations (such as detecting significant location like home or work) to the case of using multiple devices. For example, when viewing gallery photos on a smartphone with a nearby tablet device, the smartphone can suggest that the tablet be automatically woken up with the same photo in the gallery loaded so the larger screen tablet can be utilized for viewing.

5.4. Smartphone and Tablet Usage Substitution

In this section we study the potential substitution effects between smartphones and tablets. More specifically, we study whether tablet usage is primarily usage shifted from smartphones (thus substituting

Table 8: Smartphone and Tablet Substitution - Total App Time Tests Results

	C3 P-values ^a	C4 P-values ^a
Total App Time	0.001*** MD>NMD (0.33)	0.015* NMD>MD (0.22)

^a For significant results the significance level, relationship between trimmed means, and robust effect size are also reported.

the tablet for the smartphone), truly novel usage, or a combination of both. Furthermore, we study which specific activities are part of any shifted usage or novel usage.

Undoubtedly, the most sound approach to study such substitution effects would be to compare smartphone usage before and after the acquisition of a new tablet. However given a one month panel, relatively few users acquire a tablet during the observation period. Thus, we instead use an indirect approach to measuring these effects. Specifically, we compare the usage behavior of multidevice users (dataset 1, hereafter MD users) and non-multidevice users (dataset 2, hereafter NMD users). Thus under the assumption that these users are relatively similar, we can identify the tablet usage of MD users as purely novel usage, usage where a tablet was substituted for a smartphone, or a combination of both. We perform this identification by studying two distinct issues. First we study whether MD users have significantly more overall usage (across all devices) compared to NMD users (thus indicating novel usage). Second, we study whether MD users have significantly less smartphone usage than NMD users (thus indicating substitution). Furthermore, we make these comparisons first with regard to all usage, and then for individual app categories and individual apps. We use the same threshold for including app categories and individual apps as in Section 5.3.

In practical terms, performing both of these comparisons is again reducible to testing for significant differences between usage distributions. To perform these tests, we use a percentile bootstrap method that compares the 20% trimmed means (section 5.4 in [29]). Again we use the 20% trimmed mean as the location measure because of significant outliers. Thus the null hypothesis is equality of the 20% trimmed means between the groups. All tests are performed in R with the trimpb2 command from [29]. And again, for significant results (p < 0.05) we also report a robust effect size measure.

First, we compare the overall usage distributions (all usage across all devices) of the groups, Column C3 in Table 8 details the test result. We find that group MD has significantly higher trimmed mean usage time, 244.64 compared to 172.81 minutes per day. Second, we compare the smartphone usage distributions of the groups (all usage from smartphones), column C4 in Table 8 details the test result. We find group NMD has significantly higher trimmed mean usage time, 172.81 compared to 138.00 minutes per day. Thus these results suggest an overall combination of substitution and novel usage. Specifically, the results suggest that of the mean 71.83 minutes per day of tablet usage about 48% is substitution and 52% is novel usage. Finally, in terms of magnitude, the robust effect sizes suggest weak to moderate effects. This combination of substitution and novel usage have been previously observed in the case of smartphone and PC users where usage information was collected through semi-structured interviews [17].

Next, we perform the same comparisons but with regard to app usage within app categories rather than overall usage. Column C5 in Table 9 details the tests results for app categories across all devices. Whereas column C6 in Table 9 details the tests results for app categories from smartphones.

Interestingly, we find relatively few categories with significant novel usage or substitution, only the productivity and video categories indicate potential novel usage while only the games category indicates substitution. The substitution of a tablet for smartphone while playing games is unsurprising given that many games require touch based gestures and interaction that is more easily performed on the larger tablet display. Interestingly though, the productivity and video categories together only account for about 38% of the total novel usage and games only accounts for about 20% of total substituted usage. Thus other categories are likely also affected though not to a statistically significant level. Also we note that in all results from this comparison the effect sizes are relatively small.

Finally, we perform the comparisons with regard to app usage of a collection of individual apps. Column C7 in Table 10 details the tests results for individual apps across all devices and column C8 in Table 10 details the tests results for individual apps from smartphones. The results indicate web browsers and email

Category Name	C5 P-values ^a	C6 P-values ^a
Other	0.013* MD>NMD (0.26)	0.443 NMD>MD
Shopping	0.490 MD > NMD	0.337 NMD > MD
Health	0.054 MD > NMD	0.050 MD > NMD
Productivity	0.016* MD>NMD (0.20)	0.083 NMD > MD
Maps and Navigation	0.449 MD > NMD	0.835 NMD > MD
Entertainment and Sports	0.115 MD > NMD	0.202 NMD > MD
Camera	0.190 MD > NMD	0.331 MD>NMD
Weather	0.348 MD > NMD	0.634 MD > NMD
Photos and Gallery	0.164 MD > NMD	0.315 MD>NMD
E-Books	0.061 MD > NMD	0.520 NMD > MD
Games	0.216 MD > NMD	0.004*** NMD>MD (0.29)
Finance	0.103 MD > NMD	0.880 MD>NMD
Social Networking	0.334 MD>NMD	0.850 MD > NMD
Music and Audio	0.359 NMD > MD	0.057 NMD > MD
Video	$0.011^* \text{ MD} > \text{NMD} (0.28)$	0.201 NMD > MD
News	0.173 MD > NMD	0.877 NMD > MD

Table 9: Smartphone and Tablet Substitution - App Category Tests Results

^a For significant results the significance level, relationship between trimmed means, and robust effect size are also reported.

Table 10: Smartphone and Tablet Substitution - Individual App Tests Results

App Name	C7 P-values ^a	C8 P-values ^a
Other	0.005** MD>NMD (0.27)	0.002** NMD>MD (0.27)
Web Browsers	0.003** MD>NMD (0.25)	0.057 NMD>MD
YouTube	0.140 MD>NMD	0.148 NMD > MD
Email Apps	0.029* MD>NMD (0.19)	0.781 MD > NMD
Facebook	0.133 MD>NMD	0.772 NMD > MD
Messenger Apps	0.969 NMD > MD	0.913 NMD > MD

^a For significant results the significance level, relationship between trimmed means, and robust effect size are also reported.

apps as novel usage which is expected given that they are both classified as productivity apps. However, the most surprising result is that the YouTube app is not significant despite the significance of the larger video category. This might be in part due to the typically shorter and lower resolution clips on YouTube that do not benefit as much from a larger tablet display, as compared to other applications in the video category such as Netflix, Flixter, and Google Play Movies.

Importantly, our analysis of substitution effects relies on the assumption that the users of the two groups are relatively similar. This assumption is certainly not strictly true since certain types of users (basic, advanced, etc.) are both more likely to acquire a tablet device and to, for example, use certain applications more often. However, the differences in user types are unlikely to be large enough to cause all of the observed effects. We find that only household income is statistically significantly different (p < 0.05) in the demographics and smartphone technographics described in Table 1 between the groups. Furthermore, we look to address these concerns in future work (refer to Section 7).

6. Discussion

We briefly reiterate a few implications of the analysis results for two different mobile ecosystem players and we discuss our views on replicability and generalizability. For developers and designers of mobile apps, our results have implications in terms of illustrating both the ubiquity of and diversity within multidevice mobile interaction. For example, our results indicate that, for our panelists, over 50% of mobile device usage time can be considered multidevice usage. This ubiquity implies that developers and designers should *already* be thinking and designing in terms of multidevice mobile sessions. Towards this end, developers and designers can formalize their thinking through the application of already available multidevice interaction frameworks such as the 4C framework [31].

Whereas the diversity within multidevice interaction suggests that multidevice solutions should adapt to a particular user's usage patterns (as less adaptive solutions would not be able to handle this diversity). Relatedly this diversity will likely grow as the number of device types increases and new device types continue to fill small under-served niches (the smartwatch [32] being a clear example).

Additionally for mobile developers and designers, we also provide potential targets for improved mobile interaction design such as certain smartphone usage sessions where simultaneous usage of a tablet device can be inferred (as we illustrated based on smartphone app usage). For example, a user might start watching a video on their smartphone and the smartphone can suggest that the video be played directly on the tablet if simultaneous tablet usage is inferred. Such contextual inference can be viewed as helping users exploit the strengths and weaknesses of different device types without the need for the user to consciously consider, for example, device switches (as is the current case with systems such as Apple airplay).

Finally, in relation to related work, given our device based data approach, our work can be viewed as complimentary to expert survey based studies of design and development challenges for multidevice interaction such as [33] and [34]. For example, [33] specifically identified "unverified assumptions about multidevice use" as a challenge in multidevice experience design. They detail that one expert believed that users seldom access multiple mobile platforms (for example Android smartphone and iOS tablet). However we found that 14% of panelists had different platforms on their smartphones compared to their tablets. Though as we mention later, we have to be cautious with generalizations.

Whereas for advertisers, our results have implications in terms of understanding in-app advertising in the smartphone and tablet multidevice context. For example advertisers can understand the relative ad inventories between specific device types of users. Additionally we show that, given fixed ad refresh rates (in short timescales) and the effect of impression time on ad recognition and recall [35], the differences in smartphone and tablet session characteristics (length and frequency) suggest simple metrics such as number of ad impressions might need to be complemented with more complex metrics such as multidevice ad impression time.

Finally, we also discuss our views on replicability and generalizability. Replicability and generalizability are important concerns in panel based studies in general. Furthermore, given the rapidly changing nature of mobile devices, these concerns are especially relevant for mobile device based panels and we attempt to alleviate some of these concerns by having a relatively diverse multi-platform panel. However, overall we agree with Church et al. [36] that these studies are mainly about the panel populations themselves and the value is primarily in allowing researchers and the community to compare and contrast different experiences with different user populations to build a comprehensive understanding of the variety of mobile users and their unique behaviors. In this sense our multidevice usage study provides an initial benchmark that we hope spurs further investigations.

7. Conclusions and Future Work

In this paper, we provided a first look at multidevice usage by examining the smartphone and tablet usage of a US-based user panel from February 2015. Our contributions include both a descriptive analysis of the panel's multidevice usage and novel methods for examining multidevice usage in general.

We present a basic method for aggregating app sessions into multidevice usage sessions through well known Allen's interval relations.

Then we provide extensive descriptive statistics of both the multidevice and single device usage sessions. The statistics indicate, for example, that about 50% of total device usage time is multidevice usage, 40% smartphone only usage, and 10% tablet only usage.

We describe and apply a novel method to study diversity in temporal patterns in multidevice sessions. The results indicate that about 35% of multidevice sessions are dominated by a single device with only sparse usage of the second device. However, we also find significant diversity with 25 different temporal patterns having at least 1% share of multidevice sessions.

We compare the smartphone usage of users inside and outside multidevice sessions. Paired statistical tests indicate that in general users do use their smartphones differently during multidevice sessions compared to single device sessions. However the magnitude of the differences depends on the specific applications used.

Finally, we study usage substitution between smartphone and tablet devices by comparing the usage of the multidevice panel and a supplementary panel of users without tablet devices. The results indicate that tablet usage is likely a combination of usage shifted from a smartphone (thus substituting the tablet for the smartphone) and truly novel usage. Quantitatively, the results suggest that tablet usage is about equally divided between these cases, 48% and 52% respectively. Furthermore, we find that productivity apps such as web browsing and email and video apps account for a significant portion of novel usage, while game apps account for a significant portion of shifted usage.

In terms of general future research, a natural direction would be to expand the multidevice analysis to include personal computer usage. Personal computer usage data was collected from panel users but considered out of scope for this paper. As mentioned in Section 3, a difficulty in such an expansion is that the simple foreground app session definition is not directly transferable to the personal computer context. Though we plan to examine this difficulty in future work.

For future research in device substitution effects, the demographic information might allow a propensity score matching [37] approach to user group comparison. Such an approach would strengthen resulting conclusions. Furthermore since the panels are continuously ongoing, eventually the gold standard comparison of usage before and after device acquisition might be feasible.

Acknowledgment

The authors would like to thank Verto Analytics for providing the data, and Heikki Hämmäinen and Kalevi Kilkki for providing feedback during the process.

References

- M. Böhmer, B. Hecht, J. Schöning, A. Krüger, G. Bauer, Falling asleep with angry birds, facebook and kindle: A large scale study on mobile application usage, in: Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services, MobileHCI '11, 2011, pp. 47–56.
- [2] H. Falaki, R. Mahajan, S. Kandula, D. Lymberopoulos, R. Govindan, D. Estrin, Diversity in smartphone usage, in: Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services, MobiSys '10, 2010, pp. 179–194.
- [3] E. Oliver, The challenges in large-scale smartphone user studies, in: Proceedings of the 2Nd ACM International Workshop on Hot Topics in Planet-scale Measurement, HotPlanet '10, 2010, pp. 5:1–5:5.
- [4] T. Soikkeli, J. Karikoski, H. Hammainen, Characterizing smartphone usage: Diversity and end user context, International Journal of Handheld Computing Research 4 (1) (2013) 15–36. doi:10.4018/jhcr.2013010102.
- [5] D. Dearman, J. S. Pierce, It's on my other computer!: Computing with multiple devices, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '08, 2008, pp. 767–776.
- [6] Google, The new multi-screen world: Understanding cross-platform consumer behavior (2012).
- URL https://think.withgoogle.com/databoard/media/pdfs/the-new-multi-screen-world-study_research-studies. pdf
- Microsoft, Cross-screen engagement (2013).
 URL http://advertising.microsoft.com/es-xl/WWDocs/User/display/cl/researchreport/1932/global/Cross_ ScreenWhitepaper.pdf
- [8] T. Jokela, J. Ojala, T. Olsson, A diary study on combining multiple information devices in everyday activities and tasks, in: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI '15, 2015, pp. 3903–3912.
- [9] H. Müller, J. L. Gove, J. S. Webb, A. Cheang, Understanding and comparing smartphone and tablet use: Insights from a large-scale diary study, in: Proceedings of the Annual Meeting of the Australian Special Interest Group for Computer Human Interaction, OzCHI '15, 2015, pp. 427–436.
- [10] H. H. N. de Reuver, Mark; Bowman, H. Verkasalo, Smartphone measurement: Do people use mobile applications as they say they do?, in: International Conference on Mobile Business, 2012.

- [11] G. D. Montanez, R. W. White, X. Huang, Cross-device search, in: Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management, CIKM '14, 2014, pp. 1669–1678.
- [12] Y. Wang, X. Huang, R. W. White, Characterizing and supporting cross-device search tasks, in: Proceedings of the Sixth ACM International Conference on Web Search and Data Mining, WSDM '13, 2013, pp. 707–716.
- [13] I. Shmorgun, D. Lamas, M. Saks, A sample of technology substitution, in: Proceedings of the International Conference on Multimedia, Interaction, Design and Innovation, MIDI '13, 2013, pp. 15:1–15:7.
- [14] N. Eagle, A. Pentland, Reality mining: sensing complex social systems, Personal and Ubiquitous Computing 10 (4) (2006) 255–268.
- [15] A. Oulasvirta, L. Sumari, Mobile kits and laptop trays: Managing multiple devices in mobile information work, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '07, 2007, pp. 1127–1136.
- [16] S. Santosa, D. Wigdor, A field study of multi-device workflows in distributed workspaces, in: Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing, UbiComp '13, 2013, pp. 63–72.
- [17] P. J. T. J. Matthews, T., No smartphone is an island: The impact of places, situations, and other devices on smartphone use, iBM Research Report RJ10452.

URL http://domino.research.ibm.com/library/cyberdig.nsf/1e4115aea78b6e7c85256b360066f0d4/f5fd878b5b062aca85257635004ec3f5

- [18] D. Hintze, R. D. Findling, S. Scholz, R. Mayrhofer, Mobile device usage characteristics: The effect of context and form factor on locked and unlocked usage, in: Proceedings of the 12th International Conference on Advances in Mobile Computing and Multimedia, MoMM '14, 2014, pp. 105–114.
- [19] D. T. Wagner, A. Rice, A. R. Beresford, Device analyzer: Understanding smartphone usage, in: Mobile and Ubiquitous Systems: Computing, Networking, and Services, Springer, 2014, pp. 195–208.
- [20] A. Taivalsaari, T. Mikkonen, K. Systa, Liquid software manifesto: The era of multiple device ownership and its implications for software architecture, in: Proceedings of the 38th IEEE Annual Conference on Computer Software and Applications, COMPSAC'14, 2014, pp. 338–343.
- [21] P. Hamilton, D. J. Wigdor, Conductor: Enabling and understanding cross-device interaction, in: Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems, CHI '14, 2014, pp. 2773–2782.
- [22] H. Desruelle, F. Gielen, Distributed sensor-driven web applications through multi-device usage patterns, in: Y. H. Jeong, M. S. Obaidat, Y. N. Yen, J. J. Park (Eds.), Advances in Computer Science and its Applications: CSA 2013, Springer Berlin Heidelberg, Berlin, Heidelberg, 2014, pp. 13–18.
- [23] M. Schreiner, R. Rädle, H.-C. Jetter, H. Reiterer, Connichiwa: A framework for cross-device web applications, in: Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems, CHI EA '15, 2015, pp. 2163–2168.
- [24] J. F. Allen, Maintaining knowledge about temporal intervals, Commun. ACM 26 (11) (1983) 832-843.
- [25] R. D. Hays, H. Liu, A. Kapteyn, Use of internet panels to conduct surveys, Behavior Research Methods 47 (3) (2015) 685–690.
- [26] Pew Internet and American Life Project, June 10-july 12, 2015 gaming, jobs and broadband, http://www.pewinternet. org/datasets/june-10-july-12-2015-gaming-jobs-and-broadband/.
- [27] Pew Research Center, Our survey methodology in detail, http://www.pewresearch.org/methodology/ u-s-survey-research/our-survey-methodology-in-detail/#data-weighting.
- [28] A. Kotsifakos, P. Papapetrou, V. Athitsos, Ibsm: Interval-based sequence matching, in: Proceedings of the 2013 SIAM International Conference on Data Mining, 2013, pp. 596–604.
- [29] R. Wilcox, Introduction to Robust Estimation and Hypothesis Testing, 3rd Edition, Elsevier, 2012.
- [30] R. R. Wilcox, T. S. Tian, Measuring effect size: a robust heteroscedastic approach for two or more groups, Journal of Applied Statistics 38 (7) (2011) 1359–1368.
- [31] H. Sørensen, D. Raptis, J. Kjeldskov, M. B. Skov, The 4c framework: Principles of interaction in digital ecosystems, in: Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing, UbiComp '14, ACM, 2014, pp. 87–97.
- [32] S. Pizza, B. Brown, D. McMillan, A. Lampinen, Smartwatch in vivo, in: Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, CHI '16, ACM, 2016, pp. 5456–5469.
- [33] T. Dong, E. F. Churchill, J. Nichols, Understanding the challenges of designing and developing multi-device experiences, in: Proceedings of the 2016 ACM Conference on Designing Interactive Systems, DIS '16, ACM, 2016, pp. 62–72.
- [34] J. Grubert, M. Kranz, A. Quigley, Challenges in mobile multi-device ecosystems, mUX: The Journal of Mobile User Experience 5 (1) (2016) 1–22.
- [35] D. G. Goldstein, R. P. McAfee, S. Suri, The effects of exposure time on memory of display advertisements, in: Proceedings of the 12th ACM Conference on Electronic Commerce, EC '11, 2011, pp. 49–58.
- [36] K. Church, D. Ferreira, N. Banovic, K. Lyons, Understanding the challenges of mobile phone usage data, in: Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services, MobileHCI '15, 2015, pp. 504–514.
- [37] P. Rosenbaum, D. Rubin, The central role of the propensity score in observational studies for causal effects, Biometrika 70 (1) (1983) 41–55.

Appendix A.

	Dataset 1		Dataset 2	
Category Name	Smartphone	Tablet	Smartphone	
Other	37.97 %	42.84~%	42.70 %	
Social Networking	31.62 %	$5.77 \ \%$	20.61~%	
Games	8.09 %	27.81~%	13.22 %	
Video	6.06~%	12.27~%	3.93~%	
Productivity	4.21 %	2.09~%	3.20 %	
Shopping	3.72 %	1.30~%	3.57 %	
Photos and Gallery	1.55 %	0.18~%	0.70~%	
E-Books	1.47 %	0.62~%	1.81 %	
Entertainment and sports	1.25 %	3.03~%	6.07~%	
Music and Audio	1.08 %	0.40~%	1.48~%	
Maps and Navigation	0.67~%	0.11~%	0.60~%	
Camera	0.55~%	0.08~%	0.41~%	
Health	0.54~%	0.03~%	0.40~%	
News	0.48~%	3.16~%	0.49~%	
Finance	0.38~%	0.16~%	0.47~%	
Weather	0.36~%	0.15~%	0.33~%	
App Name				
Other	49.69 %	62.04~%	59.63 %	
Facebook	16.43~%	4.60~%	10.57~%	
Messenger Apps	15.19~%	1.17~%	10.04~%	
Web Browsers	12.59 %	18.33~%	13.51~%	
Email Apps	4.16~%	2.03~%	3.19~%	
YouTube	1.94~%	11.83~%	3.06~%	

Table A.11: Shares of app category and app usage time^a

^a The values are calculated over the whole datasets

Table A.12: Usage session construction statistics for dataset $2\,$

	Smartphone usage
Number of app sessions	2917208
Number of single device usage sessions	629471
Share of Allen's relations	
precedesWithinTW	8.23 %
meets	41.95 %
metBy	41.60 %
precededByWithinTW	8.23 %

Table A.13: Basic statistics on the whole dataset level for dataset 2

	Usage session length (s)	App sessions per usage session
mean	382	4.61
std. dev.	2156	8.44

Table A.14: Basic statistics on per user level for dataset $2^{\rm a}$

	usuge session	per day	per day (min)
171	2.57	47.12	298.87
66	3.00	42.06	261.09
839	1.33	27.76	205.64
	171 66 839	171 2.57 66 3.00 839 1.33	171 2.57 47.12 66 3.00 42.06 839 1.33 27.76