Understanding User Behavior in Spotify

Boxun Zhang†, Gunnar Kreitz‡§, Marcus Isaksson‡, Javier Ubillos‡, Guido Urdaneta‡,
Johan A. Pouwelse†, and Dick Epema†
Delft University of Technology†, Spotify‡, KTH – Royal Institute of Technology§

***Abstract*—Spotify is a peer-assisted music streaming service that has gained worldwide popularity in the past few years. Until now, little has been published about user behavior in such services. In this paper, we study the user behavior in Spotify by analyzing a massive dataset collected between 2010 and 2011. Firstly, we investigate the system dynamics including session arrival patterns, playback arrival patterns, and daily variation of session length. Secondly, we analyze individual user behavior on both multiple and single devices. Our analysis reveals the favorite times of day for Spotify users. We also show the correlations between both the length and the downtime of successive user sessions on single devices. In particular, we conduct the first analysis of the device-switching behavior of a massive user base.**

I. Introduction

Spotify – a peer-assisted music streaming service – has gained worldwide popularity in the past few years. Spotify provides millions of users instant access to over 20 million tracks. Our previous studies [1] have introduced the technical architecture of Spotify and analyzed Spotify's P2P network using a one-week dataset. Until now, little has been published about the user behavior in Spotify. In this paper[1], we conduct an empirical study of user behavior in Spotify by analyzing a larger dataset.

User behavior can be fundamentally different in music streaming and video streaming services, because unlike watching videos, listening to music usually does not require constant attention from users. Despite the increasing popularity of music streaming services nowadays, few studies [3] have examined the user behavior in those services. Such knowledge is crucial for improving the system design and operation. In particular, the explosively increasing adoption of smartphones and tablets urges the understanding of the usage pattern of music streaming services on mobile platforms.

Our dataset was collected by Spotify between 2010 and 2011, which covers users in Sweden, UK, and Spain. We study both the system dynamics and individual user behavior in Spotify. For system dynamics, we investigate the arrival patterns of user sessions, playbacks, and daily variations of session length. For individual users, we examine the session switching patterns on multiple devices, temporal patterns of user appearances, and correlations of successive sessions. Our findings are not only key for improving system design and operation of Spotify, but also provide valuable insights to understand user behavior in general music streaming systems.

Our main findings include:

1) We find that the session arrivals, playback arrivals, and session length exhibit strong daily patterns in Spotify.
2) We show that the session arrivals in both 1-hour and 10-minute intervals in Spotify can be modeled as a non-homogenous Poisson process.
3) We observe strong "inertia" of Spotify users to continue successive sessions on the same device.
4) We find that most Spotify users have their favorite times of day to use Spotify.
5) We find the first session length can be used as indicator for both the successive session length and downtime.

[1]For full details of this study, we refer readers to our technical report [2].

II. Dataset

A. Trace collection

For this study, we were granted access to the Hadoop cluster that is used to store and analyze Spotify's log data. From the cluster, we extracted session information of Premium users in July 2010 and March 2011, covering Sweden, UK, and Spain. The Premium dataset not only enables us to study long-term user behavior in Spotify, but also allow us to make comparative studies of user behavior on desktop and mobile devices possible, as a Premium account was required to use the smartphone client.

B. Data Sanitization

One issue we observe from our datasets is that a non-negligible fraction of sessions were logged out due to "idle-timeout". A "idle-timeout" logout happens when the client fails to send its heartbeat messages for 10 minutes, which causes it to be logged out by the Spotify server. This can happen when users experience poor network connections, such as moving in a building with poor WiFi coverage. When a client is logged out, it tries to reconnect immediately until a new connection to the server is established, which in turn will be logged as a new session.

Sessions generated by this reconnection mechanism do not reflect user behavior as they are not initiated by users themselves. Thus, we merge nearby successive sessions that are separated by "idle-timeout" and have downtime less than 30 seconds. We refer readers to our technical reports for the details about the selection of the merging threshold. We find that around 25% and 75% sessions in the original desktop and mobile datasets can be merged, respectively. The much higher percentage of merged sessions in mobile dataset is because mobile users are more likely to experience unstable network conditions, and mobile clients automatically disconnect 3G sessions when WiFi becomes available.

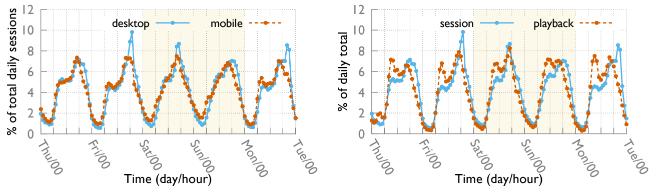

(a) sessions

(b) playbacks

Fig. 1. Arrival patterns for sessions and playbacks.

III. System Dynamics

In this section, we present our analysis of system dynamics in Spotify. Unless otherwise stated, we present only the results of the Sweden 2011 dataset, as many findings are similar among all the countries examined.

A. Session arrival

Strong daily patterns have been observed in both Internet backbone traffic [4] and many P2P systems [5]–[11]. Figure 1(a) shows the number of new sessions in Spotify within 1-hour intervals in a five-day period. The data has been normalized by the total number of daily sessions in the respective category, and the time in this figure is local time. We observe a strong daily pattern and significant variation of hourly arrival rates: the session arrival rate is lowest around 2 am, and increases sharply until 9-10 am, which we define as the *morning peak*. After the morning peak, the arrival rate drops slightly during the lunch break in weekdays. After the lunch break, the session arrival rate increases again and reaches the daily peak around 6-7 pm – the *evening peak*.

An interesting observation here is that the morning peak of mobile sessions in weekdays is often one hour ahead of desktop sessions, which we believe is because the Spotify mobile app is often used while commuting. This hypothesis can also explain the earlier evening peak of mobile sessions. Another observation is the *weekend effect*: during weekends, both the morning peak and lunch break dip of mobile sessions disappear, and the "commuting effect" also disappears. For desktop sessions, we find that the lunch break dip is noticeably less pronounced than in weekdays.

The high variation of session arrival rates clearly indicates that session arrivals in Spotify cannot be modeled as a homogeneous Poisson process. Thus, we hypothesize the session arrivals in Spotify as a non-homogenous Poisson process, which is a Poisson process with its rate parameter λ changing over time. The expected number of events between time a and time b is:

$$\lambda_{a,b} = \int_a^b \lambda(t)\mathrm{d}t \tag{1}$$

To simplify our analysis, we postulate that $\lambda(t)$ is constant in small time intervals so that we can model session arrivals in each interval as a homogenous Poisson process. This approach has been used to model the traffic of various web systems [12]. Session arrivals in a time interval can be modeled as a homogenous Poisson process if the inter-arrivals are exponentially distributed, and the arrivals are independent from each other.

We use the maximum likelihood estimation (MLE) to fit exponential distributions for session inter-arrivals, and the fitting results are tested against the *Kolmorogov-Smirnov* test with a significance level of 5%. We test session arrivals for independence by examining if the inter-arrivals are from a random sequence. For a random sequence, the probability that the autocorrelation at any lag exceeds $1.96/\sqrt{n}$ is below 5%, where n is the number of samples in the sequence [12].

The percentage of exponential inter-arrivals and independent session arrivals is above 95% for both the desktop and mobile datasets of all countries (SE, UK, ES) with 10-minute interval. For 1-hour intervals, the percentage of independent session arrivals of some datasets can drop below 80%. The reason some 1-hour intervals fail the independence test is that during a certain time of a day, 1 hour is long enough for the daily pattern to take effect (e.g., the morning peak) so the lag-1 autocorrelation of inter-arrivals can be significant.

B. Playback arrival

Besides session length, another important metric to measure the user activeness in music streaming services is the number of playbacks in a session. We define the total number of playbacks in sessions that start in each hour as the hourly playback arrival. As we have observed clear daily patterns of session arrivals and session length, we wondered whether the hourly playback arrivals also exhibit similar patterns. Note that a playback does not necessarily generate traffic to the servers, as a large fraction of tracks are cached locally and/or transferred from P2P network [1].

Figure 1(b) shows the hourly playback arrivals together with hourly session arrivals. The hourly playback arrivals are normalized the same way as the hourly session arrivals. For desktop users, we observe both the morning peak and evening peak for playback arrivals. However, we find the daily pattern of playback arrivals differs from that of session arrivals. Take the Monday in Figure 1(b) as an example: the morning peak of session arrivals contribute about 4.5% of total daily sessions but generate nearly 8% of total daily playbacks. This disproportionally high percentage of playbacks indicates that morning sessions are more active in terms of number of playbacks, which can be explained by the long session length in the morning (Figure 2). The evening (or afternoon) peak of playback arrivals is much less significant than the evening peak of session arrivals, which in turn can be explained by the shorter session length in the evening. In contrast, the playback arrivals of mobile users match fairly well with the hourly session arrivals. This means, unlike desktop sessions, mobile sessions are similarly active in terms of number of playbacks through the day. We believe this is due to the different ways of Spotify users using the desktop and mobile clients. For example, it is rare for mobile users to have very long sessions in the morning, while many desktop users tend to have "background music" during that time. The unmatched

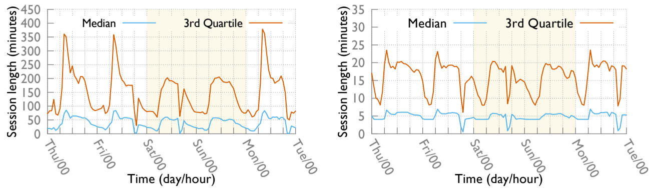

(a) desktop sessions (b) mobile sessions

Fig. 2. Daily patterns of session length (Different vertical scales in figures).

arrivals of session and playback for desktop users indicates that other metrics than the session arrival rates are also necessary to accurately capture user activity or system workload.

C. *Session length*

Besides session arrival patterns, session length is another key property for characterizing important properties such as churn rates in P2P system. A particular question for us is *Does the session length distribution also exhibit daily patterns like session arrivals?* To answer this question, we compute the median and the 3$^{\text{rd}}$ quartile of the length of new sessions starting in each hour, which is shown in Figure 2. An interesting observation is that the length of desktop sessions peaks in the morning and then decreases almost monotonously until late night. The peak of session length matches fairly well with the morning peak of session arrivals, which we believe is because many users launch Spotify to have "background music" at work. Our explanation is confirmed by the weekend effect of session length: the morning peak of session length appears in weekdays but disappears during the weekend. After the morning, the session length of desktop users is very similar in weekdays and weekends.

From Figure 2(b), we find that mobile sessions are much shorter than for desktop sessions, and its morning peak of session length is also much less significant than desktop sessions. We also notice that the median mobile session length exhibits small variation over time. However, we remark that the number of mobile sessions is in fact much larger than the number of desktop sessions, which suggests that the usage pattern is dramatically different between desktop and mobile users.

IV. User behavior

In this section, we explore the behavioral patterns of individual Spotify users on multiple devices and single device, respectively. Similarly to Section III, we present only the results of the Sweden 2011 dataset unless otherwise stated.

A. *Device switch patterns*

Many users have Spotify client installed on multiple devices, but it is not clear how they switch between those devices when using Spotify. In this section, we focus on the device switching behavior of several typical user groups: users with one desktop and one mobile, users with two desktops (e.g., one at work and the other at home) and one mobile, and users with three desktops and one mobile. In addition, we examine the behavior of users with two desktops and two mobiles as a comparison to users with only one mobile.

We study the device switch behavior by measuring the probability of users switching between different devices in successive sessions. First, we rank the desktops and mobiles of a user separately by the usage frequency (number of sessions) of each device. Take a user with p desktops and q mobiles as an example, the desktops and mobiles of that user are ranked respectively as $D_1...D_p$ and $M_1...M_q$ based on usage frequency, so we can have pairs of ranked devices used in successive sessions (e.g., $< D_1, M_1 >$). Then we can obtain the device switch probability by repeating this process for all users in our dataset.

Figure 3 illustrates the device switch probability of our interested user groups. Firstly, we find that Spotify users have strong "inertia" to continue successive sessions on both the most frequently used desktop and mobile devices, but the inertia on mobile devices is considerably bigger than on desktop devices. Secondly, we find that the probability of continuing the next session on the most used desktop and mobile devices decreases insignificantly as the total number of devices increases. Thirdly, we find that the probability of continuing successive sessions on the same device is much lower for less used desktops ($D2, D3$), while it is as high as 0.789 for the least used mobile ($M2$). Last, we notice that the probability of switching from a desktop to the most used mobile is between 0.226 - 0.299. The lesser a desktop is used, the higher the probability of switching to the most used mobile for the successive session.

B. *Favorite times of day*

We now look into the temporal aspect of the multi-device behavior in Spotify. In particular, we are interested in answering the questions *Do users have favorite time of day to use Spotify?* Before diving into details, we first define the terms that are used in the following analysis: *favorite time* is the time period of a day during which the largest fraction of a user's sessions occur. *Concentration ratio* is the ratio between the number of sessions occur in the favorite time and a user's total sessions. *Long session time* is the time period of a day that a user has the longest average session length.

To find the favorite times of Spotify users, we first divide equally the time of a day – 00-24h – into eight parts, and we count the sessions of each user started in each part of a day. Then, the favorite time for each user is the part of the day with the most sessions. Figure 4(a) shows that concentration ratios is quite high for large fraction of users. More than half of desktop users have over 35% of their sessions in their favorite times, and 30% for half of the mobile users. Another finding is that the concentration ratio of mobile users is lower than that of desktop users. We believe it is because mobiles have better accessibility than desktops so users have more opportunities to spread their mobile sessions across different times of a day.

From Figure 5(a), we find that the favorite times of large fractions of both desktop and mobile users spread between 12 - 24 pm. The most popular favorite time for mobile users is

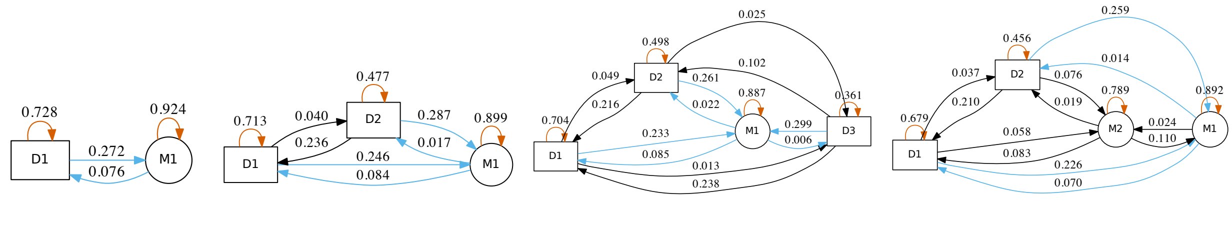

(a) (b) (c) (d)

Fig. 3. Device switch probability for different user groups (red arrows – highest switch probability of each device; blue arrows – probability of switching from and to the most frequently used mobiles.

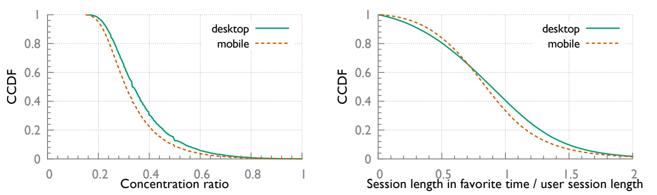

(a) Concentration ratio

(b) Session length in favorite time versus overall session length.

Fig. 4. Concentration ratio and session length of favorite times.

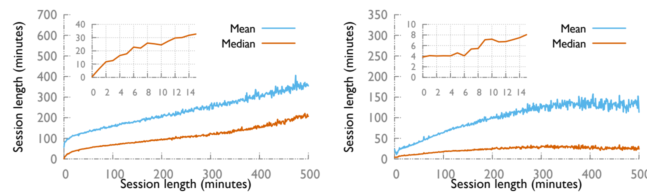

(a) desktop

(b) mobile

Fig. 6. Correlations of length of successive sessions on single device.

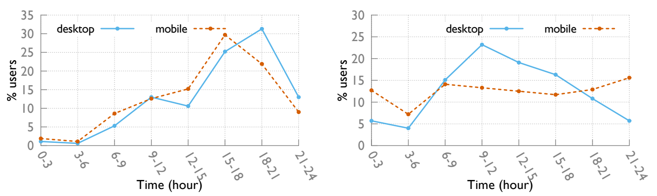

(a) Favorite time distribution.

(b) Long session time distribution.

Fig. 5. Distributions of sessions in favorite time and long session time.

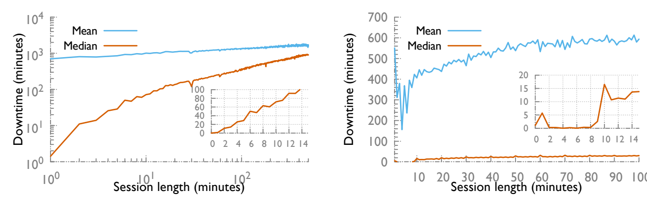

(a) desktop

(b) mobile

Fig. 7. Correlations of current session length and the successive downtime on single device.

between 15-18 pm, which is three hours earlier than the most popular favorite time for desktop users (18-21 pm). We believe this is because many Spotify users tend to use desktops rather than mobiles after they arrive home in the evening.

We notice that the favorite time of a large fraction of users is in the second half of a day, which is the time period when many short sessions occur (Figure 2 in Sec. III). By comparing the average length of sessions in the favorite time and that of all sessions of a user, we find that for many users, the sessions in their favorite times are much shorter than their sessions on average, which is shown in Figure 4(b). After comparing the average session length in different time periods of a day, we find that for 92% of all users, their favorite times are different than their long session times. The average session length in the long session times of around 50% users are twice as long as the average user session length, and four times as long as the average user session length for about 20% users.

In contrast to the favorite time, we find that the long session time of many desktop users are in the morning, while the long session times of a smaller fraction of desktop users are in late night or early morning, which is shown in Figure 5(b). The distribution of long session time for mobile users is much more even across day, with the lowest fractions occur in 3-6 am.

C. Correlations of successive sessions

Now we study the user behavior on single device. We examine the correlation between the first session length, the successive session length, and the successive downtime.

We adopt the approach used in [13] to study the correlation of successive session length: a user with n sessions have $n-1$ consecutive sessions pairs. We group all session pairs of all users by the length of the first session in each pair. Then, for each group with the first session of length x, we compute the median and mean length of second sessions. Figure 6(a) shows an almost linear correlation between the length of consecutive sessions for desktop users, which has also been observed in Gnutella and Kad [13]. This suggests the length of the first session an indicator for the successive session on the same device. However, the large gap between the mean and median values indicates high variations of the length of successive sessions.

To study the correlation between session length and successive downtime, we analyze pairs of the length of first session and the successive downtime. From Figure 7(a), we observe that the median downtime increases as the uptime of session length increases. Here we are particularly interested in the short sessions, because such sessions are the main sources of churn in a system. For sessions shorter than 15 minutes, we observe a linear correlation of session length and the successive downtime, which means many users return to the system shortly after a short session.

Successive mobile sessions on the same device exhibit similar correlations. Figure 6(b) shows the correlation between the successive mobile session length. Since mobile sessions are considerably shorter than desktop sessions, both the median and mean values of successive session length converge quickly. For sessions shorter than 6 minutes, the length of 50% successive sessions are shorter than 4 minutes. Figure 7(b) shows the correlation between the length of current session and the successive downtime. The mean and median downtimes are much shorter than desktop sessions, due to the much shorter inter-arrival times of mobile sessions. For sessions of 2-8 minutes, 50% of the successive downtimes are less than one minute! The peak of mean downtime when the session length is around one minute is probably due to users open the Spotify app accidentally, and then they quickly close the app. These findings confirm the intuition that mobile users can generate much higher churn rates than desktop users.

V. Related Work

Many studies of P2P video streaming systems [3], [5]–[8], [14] have been conducted in the past years. Session arrivals in many P2P video streaming system [5], [7], [8] exhibit strong daily patterns, and similar patterns have been observed in individual video channels [6], [14]. Compared to Spotify, both the morning peak and evening peak in P2P video streaming systems arrive later: the morning peak of session arrivals in these systems comes around lunch time, and the evening peak arrives toward midnight. We believe these differences of the peak times are caused by the types of content provided in Spotify and those video streaming systems. For example, many users listen to music in the morning but few watch movies.

Several studies [6]–[8] find that large fractions of user sessions in individual video channels end within ten minutes, which is significantly shorter than the sessions of desktop users in Spotify. It is suggested [7] that the short sessions in video streaming systems is due to impatience of users and the "intro sampling" behavior. In addition, a measurement study of peer lifetime in P2P video streaming system shows that the median of peer lifetime is less than 20 minutes. A study [3] of RealAudio traffic shows that the median user lifetime in sports and talk shows stream channels is around 50 minutes, similar to the median desktop session length in Spotify. A study [5] of a mobile IPTV system find that the average session length in each channel is around only 3 minutes caused by channel surfing [15], and users tend to stay longer in a channel between 0-6 am, which is very different from the peak time of session length (9-10 am) in Spotify.

Several studies [16]–[18] have pointed out that using multiple devices is increasingly common nowadays and mobile phones are emerging as the primary computing devices for some users. However, all these studies are based on datasets and interview feedbacks collected from small groups of users. To the best of our knowledge, we are the first to measure the device switch probability in a very large user base. Correlations between successive sessions in several P2P file-sharing systems have been studied in [13], and in this study, we take one step further by studying the correlation between session length and downtime of successive sessions.

VI. Conclusion

Spotify has gained worldwide popularity in the past few years, but little has been published about the behavioral patterns of its users, or that in other music streaming systems. In this paper, we study the user behavior in Spotify by analyzing a large dataset collected between 2010 and 2011. We found that in Spotify, not only session arrivals, but also session length and playback arrivals exhibit daily patterns. For individual users, we first studied the behavior of switching between desktop and mobile devices for using Spotify. Second, we found that Spotify users have their favorite times of day to access the service. Third, we observed clear correlations between the session length and downtime of successive user sessions on single devices. Our findings greatly deepen the understanding of user behavior in Spotify, and also provide new insights of user behavior in other music streaming services.

References

[1] G. Kreitz and F. Niemelä, "Spotify – Large Scale, Low Latency, P2P Music-on-Demand Streaming," in *IEEE P2P Computing*, 2010.

[2] B. Zhang, G. Kreitz, M. Isaksson, J. Ubillos, J. A. P. Guido Urdaneta, and D. Epema, "Understanding User Behavior in Spotify," Delft University of Technology, Tech. Rep. PDS-2013-001, 2013.

[3] A. Mena and J. Heidemann, "An Empirical Study of Real Audio Traffic," in *IEEE INFOCOM*, 2000.

[4] K. Thompson, G. Miller, and R. Wilder, "Wide-area Internet Traffic Patterns and Characteristics," *IEEE Network*, 1997.

[5] Y. Li, Y. Zhang, and R. Yuan, "Measurement and Analysis of a Large Scale Commercial Mobile Internet TV System," in *ACM IMC*, 2011.

[6] Y. Huang, T. Z. Fu, D.-M. Chiu, J. C. Lui, and C. Huang, "Challenges, Design and Analysis of a Large-scale P2P-VoD System," in *ACM SIGCOMM*, 2008.

[7] H. Yu, D. Zheng, B. Zhao, and W. Zheng, "Understanding User Behavior in Large-scale Video-on-Demand Systems," in *ACM EuroSys*, 2006.

[8] X. Hei, C. Liang, J. Liang, Y. Liu, and K. Ross, "A Measurement Study of a Large-scale P2P IPTV System," *IEEE Transactions on Multimedia*, 2007.

[9] S. Guha, N. Daswani, and R. Jain, "An Experimental Study of the Skype Peer-to-Peer VoIP System," in *IPTPS*, 2006.

[10] R. Bhagwan, S. Savage, and G. M. Voelker, "Understanding Availability," in *IPTPS*, 2003.

[11] M. Dell'Amico, P. Michiardi, and Y. Roudier, "Back To The Future: On Predicting User Uptime," *CoRR*, 2010.

[12] V. Paxson and S. Floyd, "Wide Area Traffic: the Failure of Poisson Modeling," *IEEE/ACM Transactions on Networking*, 1995.

[13] D. Stutzbach and R. Rejaie, "Understanding Churn in Peer-to-Peer Networks," in *ACM IMC*, 2006.

[14] Z. Liu, C. Wu, B. Li, and S. Zhao, "UUSee: Large-scale Operational On-demand Streaming with Random Network Coding," in *INFOCOM*, 2010.

[15] M. Cha, P. Rodriguez, J. Crowcroft, S. Moon, and X. Amatriain, "Watching Television over an IP Network," in *ACM IMC*, 2008.

[16] A. Karlson, B. Meyers, A. Jacobs, P. Johns, and S. Kane, "Working Overtime: Patterns of Smartphone and PC Usage in the Day of an Information Worker," *Pervasive Computing*, 2009.

[17] D. Dearman and J. Pierce, "It's on My Other Computer!: Computing with Multiple Devices," in *ACM SIGCHI*, 2008.

[18] S. Kane, A. Karlson, B. Meyers, P. Johns, A. Jacobs, and G. Smith, "Exploring Cross-Device Web Use on PCs and Mobile Devices," 2009.