Device Names in the Wild: Investigating Privacy Risks of Zero Configuration Networking

Bastian Könings, Christoph Bachmaier, Florian Schaub, and Michael Weber
Institute of Media Informatics
Ulm University, Germany
{ bastian.koenings | christoph.bachmaier | florian.schaub | michael.weber }@uni-ulm.de

***Abstract*—Zero configuration networking aims to support users in seamlessly connecting devices and services. However, in public networks associated service announcements pose substantial privacy risks. A major issue is the inclusion of identifying information in device names, often automatically set or suggested by devices upon initial configuration. Focusing on mDNS, we assess this issue by studying its actual extent, awareness about the problem, and potential consequences for privacy. We collected a one-week dataset of mDNS announcements in a semi-public Wi-Fi network at a university. Of 2,957 unique device names, 59% contained real names of users, with 17.6% containing first and last name. An online survey (n=137) revealed that 29% of the participants did not know the current device name of their smartphone, but that the vast majority considered periodic announcement of their full names worrisome. We further discuss specific potential privacy threats and attack scenarios stemming from mDNS device names.**

***Index Terms*—device name, mDNS, privacy, Wi-Fi, Zeroconf**

I. Introduction

In recent years, mobile devices gained more and more popularity. The increase of smartphone shipments of 38% from 2011 to 2012[1] shows that those devices are especially of high interest to consumers. Mobile access to Web and Internet services is a common aspect of smartphone use. However, even with on-going deployments of 4G infrastructure (e.g., LTE), availability of fast mobile Internet is still limited. Thus, Wi-Fi hotspots providing high bandwidth and free Internet access are very popular. Thailand's government, for example, plans to invest 27 million USD to provide free Wi-Fi access to 80% of the country[2]. In January 2013, JiWire's hotspot database[3] listed 824,276 registered Wi-Fi hotspots in 145 different countries. While the availability of free Wi-Fi offers convenience, Wi-Fi functionality of such mobile devices is often tailored for use in personal or small-scale Wi-Fi networks, which results in potential privacy threats when used in public networks.

In this paper, we investigate the privacy risks stemming from periodically sending out device names in protocols for zero configuration networking. A device name is used in those protocols in order to ease discovery and connection setup of nearby devices and services in local networks, e.g., to connect a smartphone to a printer or a TV. In case of Apple's Bonjour protocol, the device name is periodically transmitted in multicast messages. To understand the privacy risks posed by such device names, we analyzed and categorized the device names of 2,957 different devices observed during one week in our campus Wi-Fi network. We found that in 17.6% of all cases the device name contained a user's full name. We further report on the results of an online study conducted to gather insights about naming practices and awareness of users about associated privacy risks. We found that 29% of the participants did not know the current device name of their smartphones and 32% of the participants were not aware that this name was transmitted in local networks as part of those protocols.

After providing a short overview of zero configuration networking in Section II, we present the categorization of device names together with a detailed discussion of the analyzed multicast messages in Section III. The results of our online survey are provided in Section IV. Informed by these results, we discuss potential privacy threats that arise from inclusion of identifying information in device names in Section V. An overview of related work is given in Section VI. Section VII concludes the paper.

II. Zero Configuration Networking

The goal of zero configuration networking (*Zeroconf*) is to avoid manual configuration by providing a decentralized solution for discovering services of nearby devices and for announcing own services in a local network. Zeroconf was proposed by the IETF Zero Configuration Networking Working Group,[4] which specified three main conceptual requirements in order to reach this goal: IP address assignment without a DHCP server, host name resolution without a DNS server, and local service discovery without any rendezvous server. The first requirement was addressed by the standard for self-assigned link-local addressing (RFC 3927 [1]). However, no standard exists for the second and third requirement, which led to the development of diverse solutions from different parties. With *Bonjour*,[5] Apple introduced one of the most adopted Zerconf implementations, which proposes Multicast DNS (mDNS)[6] and DNS-based Service Discovery (DNS-SD)[7] as solutions for these requirements.

[1] http://bit.ly/VTWG5E (ABI Research, July 2012)

[2] http://bit.ly/ZckibR (MuniWireless, December 2012)

[3] http://v4.jiwire.com/search-hotspot-locations.htm (January 2013)

[4] http://datatracker.ietf.org/wg/zeroconf/charter/

[5] http://www.apple.com/support/bonjour/

[6] http://www.multicastdns.org/

[7] http://dns-sd.org/

In Workshop on Privacy and Security for Moving Objects (PriSMO 2013), IEEE MDM'13
Author version. Definitive version at ieeexplore.org

TABLE I
CATEGORIZATION AND DISTRIBUTION OF MDNS DEVICE NAMES FROM OUR DATASET CONTAINING 2,957 UNIQUE DEVICES.

Category	Description	Example	Number of Devices	Percentage
A	first and last name with model name	John Doe's MacBook Pro	420	14.2%
B	first and last name	John Doe	100	3.4%
C	last name with model name	Doe's MacBook Pro	47	1.6%
D	last name	Doe	19	0.6%
E	first name with model name	John's MacBook Pro	753	25.5%
F	first name	John	399	13.5%
G	nickname/alias with model name	Gandalf's MacBook Pro	271	9.1%
H	nickname/alias	Gandalf	719	24.3%
I	model name	MacBook Pro	218	7.4%
K	miscellaneous/random	iBR7tvf9Bg	11	0.4%

Multicast DNS uses conventional DNS record types ending in `.local` and packet formats, which are used by hosts in a local link, i.e., the network segment the host is connected to. Queries are sent via UDP multicast to all hosts on port 5353. Whenever a host enters a new local link, it starts a probing and announcing procedure. The probing procedure ensures that a host's chosen resource records are unique in the current link and not already taken by other hosts. Thus, a host sends a mDNS query asking for those records and resolves any potential conflicts, before announcing its own registered resource records via multicast. Host names are resolved to IP addresses via DNS type A records, e.g.:

```
Some-Computer.local A 169.254.200.50
```

DNS-SD defines certain record types to be used in service discovery. PTR records enumerate service instances, which can be reached at the host name and port number of the corresponding SRV records. TXT records provide additional information about a service instance. A specific service can either be a hardware service (e.g., a host's printer) or software service (e.g., a music player or document share).

While other Zeroconf variants with similar features exist (e.g., NetBIOS, LLMNR, or UPnP's SSDP), Bonjour with its protocols mDNS and DNS-SD is an especially interesting target for privacy analysis due to widespread adoption. In large public and semi-public Wi-Fi networks of universities, airports, or shopping malls, a host's multicast messages can reach a large number of other hosts. According to Hong et al. [2], mDNS traffic consumes about 13% of total bandwidth in such wireless networks. Host names are used in mDNS records to help users and hosts identify other hosts in the local network. However, the periodic announcement of host names can have varying implications on user privacy in such public and semi-public networks, depending on how the host name is composed. A host name is typically composed of the device name, which in turn can range from pseudonyms that do not directly reveal any personal information about the user to device names that disclose the type of device, personal interests, as well as nicknames and full names of users (e.g., John-Doe's-iPhone). In the next section, we investigate the extent of this privacy issue in a real network setting.

III. DEVICE NAMES IN THE WILD

In order to investigate the potential impact of device names on user privacy, we captured all mDNS responses within a certain subnet of our campus Wi-Fi network over the period of one week. According to prior agreement with our data protection officer, only necessary data was extracted from mDNS messages and partially anonymized before storage, i.e., only the hashed MAC address and the device name were stored. The collected dataset includes 2,957 unique devices and their device names.

A. Device Name Categorization

After initial analysis of our dataset we derived ten device name categories as listed in Table I, ordered by descending privacy sensitivity. Due to high diversity of device names, we had to manually classify them into the categories. If names could not clearly be classified into one category (e.g., if a name could be both a first name or nickname), the less sensitive category was chosen. Table I shows the distribution of the different device names and gives examples for each category.

B. Results

We found that in 59% of all cases the device name contains either the user's first name, last name, or both (categories A to F). A user's full name was found in 17.6% of all mDNS messages (categories A and B). In other words, almost two out of three device names contain at least a part of the user's name, if not even the complete name.

A model name was found in 58% of all messages. Looking at the combined categories with and without model names (A/B, C/D, E/F, and G/H), we found that with decreasing sensitivity level the number of device names including a model name also decreases. While in categories A/B about 81% contain a model name, categories C/D contain 71%, E/F 65%, and only 27% of device names in categories G/H contain a model name.

The high percentage in categories A/B could lead to the assumption that most of the corresponding device names were created by Apple's default device naming practice, which suggests inclusion of a user's first or even full name[8] and

[8] The default device name of a newly configured Apple device depends further on iTunes and user account settings of the host computer.

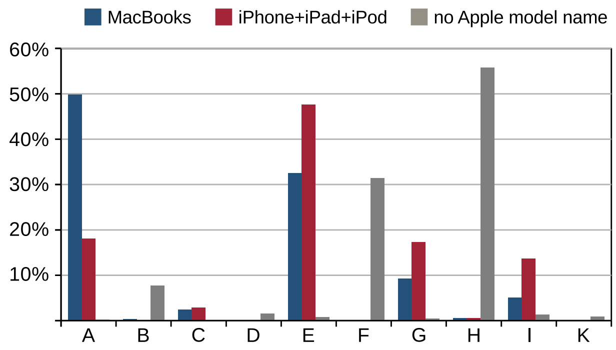

Fig. 1. Distribution of device names for Apple MacBooks, mobile Apple devices (iPhone, iPad, iPod), and devices without Apple model name.

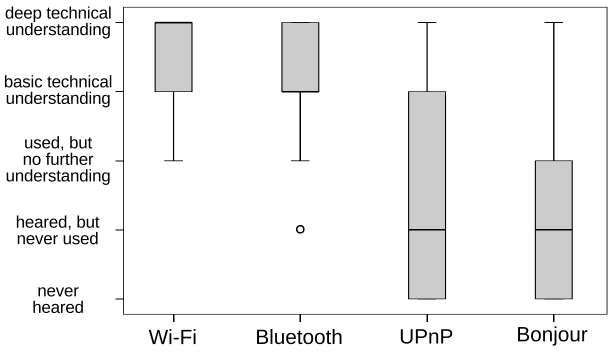

Fig. 2. Participants' familiarity with Wi-Fi, Bluetooth, UPnP, and Bonjour; ranging from "never heard before" to "deep technical understanding".

device model name. The low percentage in categories G/H suggests that users mostly do not include a model name when deliberately choosing a nickname as device name.

Of all 1,685 devices revealing their model name, we found 59% to be iPhones, 20% MacBooks (Pro/Air), 11% iPads, 8% iPods, and 2% others (e.g., iMac). The distribution of device names for MacBooks and mobile Apple devices (which we considered to be iPhones, iPads, and iPods), shows that MacBooks more often revealed the full name (50%) than mobile devices (18%), which more often revealed only the first names (48%) compared to MacBooks with 32% (see Fig. 1). Devices that we could not identify as Apple devices by their model name, in most cases used a nickname or alias (56%), which corresponds to the former finding, that with lower sensitivity level also less often the model name is provided.

IV. User Awareness

We conducted an online survey in order to better understand how users select their device names and if they are aware of privacy risks stemming from service announcements including device names. A total of 137 individuals aged 19-55 (Mdn=25) participated in our survey (32 female, 105 male). The majority of participants (65%) work or study in the ICT sector. Notebooks were owned by 130 participants; smartphones by 105. Three owned neither a notebook nor smartphone.

A. Expertise and Privacy Proclivity

Participants were asked to answer a series of questions regarding their knowledge of Wi-Fi, Bluetooth, UPnP and Bonjour using a five-point Likert scale. Depending on their level of familiarity with each of these four technologies, we categorized participants as *novices* (23% of the participants), *average users* (51%) or *experts* (26%). Figure 2 shows the very different levels of familiarity with Wi-Fi and Bluetooth in contrast to UPnP and Bonjour. All participants had heard of Wi-Fi and had used it at some point, whereas 40% stated that they had never heard of Bonjour and 31% stated that they had heard about Bonjour but had never used it. However, most of the 54 participants (83%) who never heard of Bonjour did not own an Apple device, which suggests that most Apple users are aware of Bonjour. Out of the 43 participants who stated to have heard about Bonjour but never used it, 14 (33%) did own an Apple device. As Bonjour is activated by default on Apple devices, it is fair to assume that these persons as well as the Apple users who never heard of Bonjour had used it at some point in time without being aware of it.

Furthermore, we assessed their general privacy proclivity with five questions, of which three were adopted from Westin's Privacy Index studies [3] and the other two were chosen to reflect the topic of our survey. If not otherwise stated, participants had to use a four-point Likert scale to answer whether they *strongly agreed*, *agreed*, *disagreed*, or *strongly disagreed* with the presented statements. As suggested by Westin, we used the answers to these questions to categorize participants as *privacy unconcerned*, *privacy pragmatists*, and *privacy fundamentalists*. *Privacy unconcerned* (5% of the participants) do not worry about their privacy and do not mind revealing personal information. *Privacy pragmatists* (36%) are reluctant to give out personal information but are willing to do so if the benefit warrants it. *Privacy fundamentalists* (59%) protect their privacy without compromise and are very concerned about how others treat their personal information.

B. Results and Discussion

Using Spearman's rank correlation, we analyzed survey replies in relation to expertise and privacy proclivity of participants. Interestingly, we did not find a correlation between technology understanding and privacy proclivity.

Concerning the awareness about device names being periodically announced in local networks for service discovery, we found no correlation between awareness and privacy proclivity. However, awareness of this problem is significantly higher with increasing level of expertise (r=.425, p<.001, n=61). This aligns with expectations that participants with deeper understanding of Wi-Fi, Bluetooth, UPnP and Bonjour know about the announcement of device names, while privacy fundamentalists may have a strong desire to protect privacy but may lack the knowledge for doing so. Overall, 32% of the participants were not aware of the existence and nature of device name announcements.

Independently of their actual knowledge about zero configuration networking, we asked participants to rate how alarmed

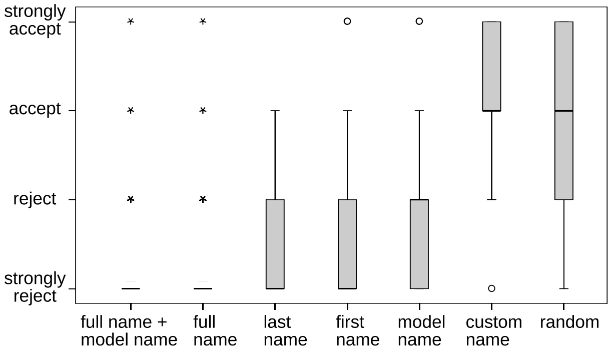

Fig. 3. Participants' ratings of initial default device names for smartphones ranging from "strongly reject" to "strongly accept" as default name.

TABLE II
DISTRIBUTION OF DEVICE NAME CATEGORIES CHOOSEN BY PARTICIPANTS FOR SMARTPHONES AND NOTEBOOKS.

Category	Smartphone	Notebook
A (full name with model name)	1%	0%
B (full name)	2%	3%
C (last name with model name)	0%	0%
D (last name)	1%	2%
E (first name with model name)	0%	1%
F (first name)	9%	10%
G (nickname/alias with model name)	0%	0%
H (nickname/alias)	32%	56%
I (model name)	23%	7%
K (miscellaneous/random)	3%	4%
Don't know	29%	17%

they would be by their full name being announced in local networks in relation to different source devices. Both, for smartphones (92%) and notebooks (89%), a large majority of participants would be highly alarmed by such announcements regardless of their privacy proclivity.

Participants were further presented with examples of initial device names and asked whether or not they would (1) *strongly reject*, (2) *reject*, (3) *accept*, or (4) *strongly accept* them as default names. These examples were (see Fig. 3) the user's full name and model name (Mdn=1, 3% accept), the user's full name (Mdn=1, 3% accept), the user's last name (Mdn=1, 2% accept), the user's first name (Mdn=1, 14% accept), the device model name (Mdn=2, 25% accept), prompting for a custom name (Mdn=3, 85%), and a random string (Mdn=3, 60% accept). Thus, being asked to specify a device name is preferred over any default device name, followed by a random string. However, we found a positive correlation between increased level of expertise and the desire to manually specify device names (r=.179, p=.018, n=137). The results also show that using the first name as device name is considered less sensitive than using last or full name.

Asked about their current device names, we found that with increasing level of expertise, participants were more likely to choose more privacy friendly device names for mobile devices (r=.359, p<.001, n=118), while there was no correlation between privacy proclivity and device naming practice. Table II shows the distribution of device names used by participants on their smartphones and notebooks. The dominant category H (nickname/alias) matches the identified preference for customized device names as well as the results of the mDNS analysis (24% of all devices in category H). In contrast, category I (model name) was deemed not to be a suitable device name, yet, Table II shows that about 23% of all smartphone users had the model name set as their device name. The results of the mDNS analysis with only 7% in category I aligns better to the participant's opinion. Finally, one might suggest that the participants not knowing their device names (29% for smartphones, 17% for notebooks) still have set the default device name, which in case of Apple products is often the full name and model name.

V. Privacy Threats & Attack Scenarios

Our survey results suggest that most users (about 90%) perceive the announcement of a user's full name in wireless networks as a privacy risk. However, other studies have shown that the name is not always perceived as an information with high privacy sensitivity [4] and that users are often unaware of associated privacy risks [4]–[6]. In order to clarify the potential privacy risks that stem from device name announcements, we subsequently discuss a number of specific threats and attack scenarios enabled by Zeroconf announcements.

A. Identification & Social Relation Inference

In smaller public Wi-Fis with a limited number of users, the announcement of a device name could directly identify a person. As stated by Aura et al. [7], "using a laptop computer is akin to wearing a name badge that reveals the person's identity." For instance, Alice is sitting with her iPad in a cafe with free Wi-Fi and is the only person using an iPad. Her device name is configured as *Alice-Doe's-iPad*. If Bob wants to know her name, he just needs to connect to the same Wi-Fi network, start a Zeroconf browser[9] and search for iPads. As Alice's iPad is the only device of this type, Bob can directly identify Alice, gather background information about her social profile with online searches, and approach her as an old acquaintance for whatever purpose ("Hi Alice! Do you remember me from MDM 2013?"). If Bob wants to know who Alice is meeting regularly, he could infer her social relations by monitoring the network over a period of time to identify correlations of device names in data samples including Alice at different points in time.

In larger networks containing multiple devices of a specific type, an adversary can leverage further context information to reduce the number of potential device name matches. For instance, if more persons are using iPads but Alice is the only woman, Bob can search for device names that contain the correct model type (iPad) and female full or first names. Another possibility in wireless networks is to configure a Wi-Fi sniffer in monitor mode and analyze the signal strength of received messages in order to limit the number of mDNS

[9] Zerconf browsers are available for several platforms, e.g., as apps for Android: https://play.google.com/store/apps/details?id=com.melloware.zeroconf

messages to those from nearby devices. The use of directed antennas would further increase the probability of receiving only messages from the targeted person.

In situations where no wireless network is available, an adversary could perform an evil twin attack by creating an access point that spoofs SSIDs of common hotspots[10] and public Wi-Fi providers in order to get devices that were previously connected to a network with the same SSID to automatically join the adversary's bait network.

B. Location Tracking & Behavior Profiling

Location privacy and the involved issue of user tracking is a prominent privacy issue that has attracted considerable attention in recent years [8]. One common approach to track users is to follow the unique MAC address of their Wi-Fi devices [9], [10]. While the MAC address on its own does not directly allow identification of a user, the device name in mDNS messages sent by the same MAC address facilitates matching of a captured track to a particular person.

Once device and identity have been linked, individual behavior profiles can be created. Consider the following example. Alice works at the university and always takes the bus. Bob has already linked Alice's MAC address to her identity. Bob wants to know when Alice arrives at university and places a mobile sniffer at the university's bus stop. As mobile devices are frequently sending Wi-Fi probe requests in order to discover known wireless networks [10], Bob can easily receive Alice's MAC address once she arrives at the bus stop, without being connected to any network. If Bob has enough resources to track Alice over a longer time period and at different locations, he can easily build a detailed behavior profile containing information about how long she stays at work, when she leaves home, is at the cafe, or goes to the gym.

C. Theft Targeting

The fact that Zeroconf protocols are not only announcing details about software services but also about existing hardware services (e.g., printers, scanners, cameras) allows attackers to easily build inventory lists of users. Especially, when the device name itself contains the type of device, as was the case in 58% of all analyzed mDNS messages in our dataset. Following the previous example, we assume that Alice has a printer at home and prints documents with her MacBook Air. This printer would be announced as *AirPrint Canon iP3600 series @ Alice-Doe's-MacBook-Air._ipp._tcp.local.* If Bob receives this announcement in the cafe, he knows not only that Alice has a MacBook Air but also a Canon iP3600 printer at home. With his already deployed tracking system, Bob knows the best time to burglarize Alice's home. In combination with the former discussed setup of a fake access point, this method could also be used by thieves to explore what active devices are inside nearby cars. A similar approach had already led to several car break-ins when mobile phones announced their presence via Bluetooth.[11]

D. Counter Measures

Different counter measures are possible to mitigate the discussed privacy risks. An obvious solution to avoid identification and inference of social relations would be to change device names to pseudonyms or non-descriptive identifiers. However, as device names should allow users to easily identify devices a descriptive name is preferable. A better solution would be to selectively disable announcing procedures in public networks, as also proposed by Aura et al. [7]. Users should be able to explicitly choose in which networks service discovery should be enabled. This solution would also mitigate the problem of theft targeting. However, in situations where the user wants to announce this information and does not want to change the device name to a non-descriptive identifier, different solution is required. Pang et al. [11] propose *Tryst* an architecture that ensures confidentiality of announced information by extending existing service discovery protocols with encryption. The drawback of such solutions is that they always require additional configuration overhead like secure pairing, key generation, or key exchange. Tracking of unique MAC addresses can be mitigated by dynamic address changes [9] or encryption mechanisms on the link layer as proposed by Greenstein et al. [12], which also entails additional configuration overhead.

VI. Related Work

The leakage of personal information when using mobile devices in public Wi-Fi hotspots is a well known problem. Aura et al. [7] developed a monitoring tool for detecting identifier leaks of laptops, stemming from signaling protocols, packet headers, and communication metadata. They find that all layers of the protocol stack leak various plaintext identifiers of users, their computers and their affiliations. However, as the authors did not consider mDNS, they conclude that a user's real name is not leaked as frequently as other identifiers.

Kowitz and Cranor [5] study how visualization of leaked information can influence users' privacy expectations when using wireless networks. They developed a monitoring tool that visualizes leakage of sensitive strings on a large display in a common workplace environment. However, in a two-week trial they could not find any significant change in participants' measured comfort level for communication mediums (e.g., whether participants prefer a phone call instead of writing a text message over an unencrypted wireless medium), although some participants noted a change in their privacy expectations.

Consolvo et al. [4] propose the *Wi-Fi Privacy Ticker*, which notifies users about sensitive information being transmitted over Wi-Fi during web browsing. The Ticker allows users to set custom watch lists of terms they consider to be privacy sensitive and rate the terms with a sensitivity level of low, medium, and high. Nearly all participants added their email addresses as low sensitive, and passwords as high sensitive to the watch list. The last name was added by 11 users and rated as low or medium sensitive. In a three-week field study with 17 participants, they found that improved awareness of situations in which information was exposed over Wi-Fi led to a change

[10] https://wigle.net/gps/gps/main/ssidstats

[11] http://bit.ly/WRI8C5 (Cambridge News, March 2007)

in behavior. Participants stated to be more concerned about exposure on Wi-Fi and to think more about what information they may expose to others. As the study did only investigate information leakage by web browsers, it would be interesting how the privacy ticker would influence users' expectations of privacy when applied to mDNS messages.

In an exploratory study, Klasnja et al. [6] find that users from the general public often have a detailed practical understanding of Wi-Fi, which matches our study results (see Fig. 2). But despite having this practical understanding, Klasnja et al. find that most users lack understanding of implied privacy risks and are unaware of the visibility of their communications in unencrypted Wi-Fi networks. Even if the results of our study showed that only 32% of the participants were unaware of device name announcements, this is still a relative high number considering that most participants were working in ICT-related fields. However, Klasnja et al. state further, that clarifying privacy risks did have a positive effect on users' willingness to take privacy supporting actions, e.g., using Wi-Fi less often, being more careful about which networks they connect to, or not using their full names as usernames.

Kindberg and Jones [13] studied users' naming practices of Bluetooth devices. They analyzed Bluetooth names of 1,703 devices captured at three locations in the city of Bristol, UK. They categorize names into *default* (the model name) and *user-defined* names. A user's full name was considered to be a user-defined name. They find that in the city center 58% of Bluetooth devices had user-defined names, whereas at the university campus and a pub in the city center about 80% had user-defined names. However, the results do not differentiate further between full names and other user-defined names.

Cheng et al. [14] captured and analyzed the network traffic at 20 airport hotspots in four different countries. Their captured dataset covers a time period of 15-60 minutes per hotspot and reveals that two thirds of travelers leak privacy sensitive data by DNS queries, web browsing, or querying search engines. They find that name leakage is the most prominent factor, caused in more than 90% by mDNS messages. Similar to our work, they discovered over 1,800 unique device names in mDNS messages, of which about 600 messages contained real names. However, it is not discussed whether device names contained full names, first names, or model names.

VII. Conclusion

Zero configuration networking facilitates user-friendly setup of device connections and service usage in home networks. However, current implementations pose several privacy risks, especially when mobile devices are operated in public wireless networks. In this paper, we investigated the privacy risks of Apple's Bonjour protocol leveraging Multicast DNS (mDNS) and DNS-based Service Discovery (DNS-SD). The results from our analysis of a one-week captured dataset with 2,957 unique devices revealed that 59% of device names in mDNS messages included a user's real name (first, last, full), a full name in 17.6%, and the model name in 58% of all devices names. MacBooks more often revealed the full name (50%) than iPhones, iPads and iPods (18%) what could be caused by different default naming practices on those devices. Furthermore, the results of our online survey (n=137) showed that participants often did not know the current device name of their smartphones (29%) and were not aware of the fact that device names are announced in local networks (32%). As only 13 participants indicated to use their name as part of their device names, it is possible that especially users not knowing their device names involuntarily included parts of their real names due to automatically set default device names. Interestingly, most participants did not want to have parts of their name set as default device names and prefer custom chosen device names. We discussed several privacy threats that stem from the announcement of real names in mDNS messages and believe that most severe threats could be mitigated if not the real name would be used. However, even if mitigation is as simple as changing the device name, the lack of awareness of privacy risks may prevent users to do so. Thus, novel approaches are needed which help users to gain better awareness of such risks. Further, the default naming practices of devices names should be revised and users should be able to limit service discovery to a selected set of networks.

References

[1] S. Cheshire, B. Aboba, and E. Guttman, *Dynamic configuration of IPv4 link-local addresses*, IETF RFC 3927, 2005.

[2] S. G. Hong, S. Srinivasan, and H. Schulzrinne, "Measurements of multicast service discovery in a campus wireless network," in *Proc. GLOBECOM '09*. IEEE, 2009.

[3] H. Taylor, "Most people are "privacy pragmatists" who, while concerned about privacy, will sometimes trade it off for other benefits," Harris Interactive Survey, Rochester, New York, 2003.

[4] S. Consolvo, J. Jung, B. Greenstein, P. Powledge, G. Maganis, and D. Avrahami, "The wi-fi privacy ticker: Improving awareness & control of personal information exposure on Wi-Fi," in *Proc. UbiComp '10*. ACM, 2010.

[5] B. Kowitz and L. Cranor, "Peripheral privacy notifications for wireless networks," in *Proc. Workshop on Privacy in the Electronic Society (WPES '05)*. ACM, 2005.

[6] P. Klasnja, S. Consolvo, J. Jung, B. M. Greenstein, L. LeGrand, P. Powledge, and D. Wetherall, "When I am on Wi-Fi, I am fearless": Privacy concerns & practices in everyday Wi-Fi use," in *Proc. CHI '09*. ACM, 2009.

[7] T. Aura, J. Lindqvist, M. Roe, and A. Mohammed, "Chattering laptops," in *Proc. PETS '08*. Springer, 2008.

[8] J. Krumm, "A survey of computational location privacy," *Personal and Ubiquitous Computing*, vol. 13, no. 6, pp. 391–399, 2008.

[9] M. Gruteser and D. Grunwald, "Enhancing location privacy in wireless LAN through disposable interface identifiers: A quantitative analysis," *Mobile Networks and Applications*, vol. 10, no. 3, pp. 315–325, 2005.

[10] B. Greenstein, R. Gummadi, J. Pang, M. Y. Chen, T. Kohno, S. Seshan, and D. Wetherall, "Can Ferris Bueller still have his day off? protecting privacy in an era of wireless devices," in *Proc. Workshop on Hot Topics in Operating Systems (HotOS XI)*. USENIX Association, 2007.

[11] J. Pang, B. Greenstein, D. McCoy, S. Seshan, and D. Wetherall, "Tryst: The case for confidential service discovery," in *Proc. Workshop on Hot Topics in Networks (HotNets-VI)*. ACM, 2007.

[12] B. Greenstein, D. McCoy, J. Pang, T. Kohno, S. Seshan, and D. Wetherall, "Improving wireless privacy with an identifier-free link layer protocol," in *Proc. MobiSys '08*. ACM, 2008.

[13] T. Kindberg and T. Jones, "Merolyn the phone: A study of bluetooth naming practices," in *Proc. UbiComp '07*. Springer, 2007.

[14] N. Cheng, X. Wang, P. Mohapatra, and S. Aruna, "Characterizing privacy leakage of public WiFi networks for users on travel," in *Proc. INFOCOM '13*. IEEE, 2013.