Towards a Three-tiered Social Graph in Decentralized Online Social Networks

Felix Beierle
Service-centric Networking
Technische Universität Berlin
Telekom Innovation Laboratories
Berlin, Germany
beierle@tu-berlin.de

Sebastian Göndör
Service-centric Networking
Technische Universität Berlin
Telekom Innovation Laboratories
Berlin, Germany
sebastian.goendoer@tu-berlin.de

Axel Küpper
Service-centric Networking
Technische Universität Berlin
Telekom Innovation Laboratories
Berlin, Germany
axel.kuepper@tu-berlin.de

ABSTRACT

Online Social Networks have become one of the main tools for interpersonal online communication. In the age of the smartphone, mobile user scenarios become more and more important for Online Social Networks. Smartphones enable location-based and context-aware services, but bring the increased risk of privacy violations - at least in centralized OSN architectures. Decentralized Online Social Network architectures are promising as they inherently offer better privacy and less dependence on a single service provider, but they bring new challenges regarding core features of Online Social Networks. In this paper, we introduce a three-tiered view of the social graph and propose a new architecture for decentralized Online Social Networking applications, supporting the three-tiered view and focusing on location-based and context-aware user scenarios.

Categories and Subject Descriptors

C.2.4 [**Computer-Communication Networks**]: Distributed Systems; H.3.5 [**Information Storage and Retrieval**]: Online Information Services

General Terms

Design

Keywords

online social networks; mobile social networks; social graph; location-based services; context-aware services

1. INTRODUCTION

Current *Online Social Network*s (OSN) utilize centralized architectures and store any and all information about their users. The resulting loss of control over one's personal information and the resulting lack of privacy is a major concern for many users. In centralized OSN architectures, lack of privacy is often an issue, because the operator of the OSN might base their business model on the users' private data, while the users might not have intended to share their data to that extent [8, 6, 4]. Furthermore, the usage of proprietary interfaces prohibits communication between different OSN service providers, creating a well calculated lock-in effect for their users. Moreover, with the steep rise in the processing power and enhanced features of smartphones as well as their fast adoption by users, OSNs have moved massively to the mobile world (inspiring the term *Mobile Social Network*s (MSN) for OSNs focused on mobile scenarios [11, 3]). This has opened new opportunities for OSNs and their users, but also brings increased privacy risks. For instance, many location-based and context-aware user scenarios are only possible in MSNs, but the amount of sensitive data collected, stored, and shared typically increases in these scenarios.

In this paper, we address these aspects of OSNs. In particular, we argue that distributing the social graph, which is one of the core pieces of information of any social network, is a vital step towards better privacy. We also argue that, additionally, the capabilities of smartphones can be leveraged to achieve both more flexible use of context information and higher privacy of sensitive data for users of OSNs.

We investigate the relevant core features of OSNs and review existing architectural approaches for implementing online social networks. Based on this, we propose a three-tiered view of the social graph and a new architectural concept, called *T3 architecture*, supporting this three-tiered view. For offering location-based and context-aware services and taking care of better privacy protection at the same time, our approach is to use a global directory for retrieving information on how to contact other users, and to use additional decentralized directories that can offer search and recommendations for specific contexts or locations. Thus, in summary, the three main contributions of this paper are (1) the elaboration of core features of OSNs, (2) the proposal to view the social graph as a three-tiered concept, and (3) an architectural approach for distributed online social network applications supporting the three-tiered social graph and location-based and context-aware user scenarios.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HotPOST'15, June 22, 2015, Hangzhou, China.

Copyright © 2015 ACM 978-1-4503-3517-1/15/06 ...$15.00.

DOI: http://dx.doi.org/10.1145/2757513.2757517.

The rest of this paper is structured as follows: After elaborating the core features of OSNs (Section 2) and investigating existing architectures for OSNs (Section 3), we develop our three-tier view of the social graph in Section 4 and introduce the T3 architecture in Section 5. In Section 6, we conclude and point out future work.

2. FROM ONLINE TO MOBILE SOCIAL NETWORKS

After analyzing user scenarios for OSNs for smartphone users and existing definitions for OSNs and MSNs, we define and discuss the core features of OSNs.

2.1 Mobile OSN User Scenarios

Before detailing the scenarios, we first explain why the smartphone is the essential social networking device: McLuhan imagined each technology as an extension of man [13]: if the car is the extension of feet, glasses an extension of the eyes, then the smartphone is most certainly an extension of all parts that constitute a person's social interactions. We believe that while the smartphone is capable of being such an extension of man, current OSN architectures or applications do not reflect this sentiment. Carrying the smartphone at all times - with its sensors, processing power, and connectivity like Bluetooth and WiFi-Direct - especially local services do not need to rely on cloud services or traditional OSNs.

Imagine Alice, a smartphone user living in a larger city. Going through her daily life, Alice is collecting data on her smartphone, e.g., visited locations or artists she listened to. By collecting this data, Alice's interests are expressed implicitly rather than explicitly as it is done by clicking a "like"-button. Going shopping, she can share selected parts of her collected context and location data to receive personalized recommendations. After shopping, imagine Alice receiving a notification from Bob, who wants to play baseball. She does not play baseball herself, but her smartphone can calculate and suggest friends of hers that like baseball, and are not friends with Bob. Alice can decide to forward recommended missing players.

Studying at a university, Alice can talk to the people that are in the same class. Going to the library to study, she can look for people who are both at the library and in her classes. Going downtown in the evening, when visiting a crowded place, she can chat with strangers that are in her proximity, e.g., to figure out what is at the center of the crowd. At night, Alice goes to a party and quickly realizes that she does not know many people there. She can retrieve information about the people around her and find out mutual contacts. Traveling to another city, Alice should be able to query a public computer at her travel destination for recommendations of local attractions or restaurants. Furthermore, current contact information, like the postal address, could be maintained locally with the users, making an up-to-date version available for contacts.

2.2 Definitions of OSNs and MSNs

Some related work distinguishes between OSN and MSN. In this section, we analyze the different definitions of the two terms.

In their paper about decentralized OSNs, Paul et al. define OSNs as providing a form of *profile management*, *relationship management* between users, *interaction* between users, as well having an *API* (for interaction with third party applications), *search*, *recommendations*, and *social network connector*s (for connections between different OSNs) [14]. In contrast to [14], we argue that an API is not essential for the user of an OSN and is thus not necessarily a defining feature. Social network connectors contradict the lock-in effect that centralized OSN platforms desire; for an approach towards interconnection between OSNs, see [9]. Furthermore, the authors of [14] do not consider the importance of mobile user scenarios as described in Section 2.1.

In [11], Kayastha et al. define *Mobile Social Networks* as a combination of OSNs with location-based services and connections via Bluetooth and Wi-Fi. In [3], Bellavista et al. define MSNs as OSNs that offer the possibility for opportunistic communities based on location information, with the possibility of not needing a connection to a server. On a less technical level, they define an MSN as the combination of virtual communities ("virtual social spaces") of OSNs with "physical social spaces."

Aside from architectural implications regarding the connections between mobile devices, we consider context and location information to be the essential new element that mobile usage and smartphones bring into the world of OSNs. We argue that with the pervasiveness of smartphones, the *mobile* aspect of an OSN is indispensable, even for traditional OSNs like Facebook. In [11], Facebook, typically regarded as an OSN, is defined as an MSN, as it offers location-based services, and is reachable via a mobile website or mobile application. We therefore use the terms OSN and MSN synonymously.

2.3 Core Features

Motivated by the analysis and observations above, we distinguish between the following six core features for OSNs:

Personal Profile. Each user has a personal profile in which he/she can describe oneself, typically with a profile picture.

Social Graph. The social graph consists of the connections between a user of the OSN with other users or products.

Interaction. Messages, updates, posts, or likes by friends are typically displayed to linked users.

Context Information. The user's context, e.g., current location, availability for instant messaging, etc., is often displayed for linked users or used for group creation.

Search. Search for specific users or users with a similar context.

Recommendations. Recommendations of friends or products, based on social graph and context information, is the most common feature with which users can explore the network around them.

There are different OSNs with different focuses. LinkedIn[1], for example, focuses on work relationships and career opportunities, Twitter[2] is about - typically publicly - broadcasting short messages. For each OSN, the six listed

[1] https://www.linkedin.com/

[2] https://twitter.com/

features might be of differing importance, and might have differing complexity.

With the rise of smartphone devices in recent years, Instant Messenger applications seem to fulfill the role of OSNs, with messaging apps like WeChat[3] or WhatsApp offering some OSN features like a short profile. Looking closer, the contact list or address book of each of those applications - or the phone's contact list with PTSN telephone numbers - is actually a part of a social graph: Each user sees just direct connections, he does not see people more than one hop away. A social graph that is stored in a distributed way offers inherently more privacy to the users of the OSN. At the same time, recommendations that are based on knowledge about the graph, pose a new challenge.

Today, each OSN or messaging application is vertically integrated: They have their own identity management, their own social graph, and - most often - proprietary communication protocols. Thus, each OSN acts like a silo, with users being bound to restrictions their OSN of choice comes with.

The modern smartphone is an optimal online social networking device. It contains a camera, sensors that register the user's context, it can connect to the Internet, and it can display websites, videos, or images that other users posted. The generated context information especially deserves protection, as it can contain the most private information like health data or exact user locations at any time. Given the processing power of current smartphones, they can be used to locally process data, e.g., lots of the context data a phone generates. For instance, a phone could track and rank visited places or favorite artists, independent of specific OSN applications, leaving the user in control of his data, letting him decide when to share what data. Additionally, collected context and location data offers the possibility to create dynamic social graphs, connecting users without them having to explicitly befriend each other. This way, people who are in the same city and are interested in a specific style of music genre could automatically be connected and share information about local shows and concerts.

3. ARCHITECTURAL APPROACHES

In this section, we motivate the utilization of decentralized system architectures, before we summarize current research about OSNs and MSNs.

3.1 On Decentralized Architectures

In the beginning of the 1990s, AOL (formerly America Online) provided its costumers access to their catalog of sites and services. With the advent of the World Wide Web and email, users were no longer bound to the web of a single company, but could explore the Internet freely and write emails to any other user using a different ISP. The decentralized concept of the World Wide Web allowed the development of open communication platforms.

A more recent example for a decentralized system is XMPP[4] (Extensible Messaging and Presence Protocol). The messaging works similar to email, with users having an identifier with a server, and servers sending messages to each other. Additionally, the idea of *presence* allows users to register with their server with different devices. This presence can be published to the own address book (*buddies*).

At the same time, the recent past also showed a trend towards new centralized systems, most prominently Facebook[5]. Every few weeks, especially in the EU, Facebook is in the news because of privacy issues and many users are concerned [6, 4].

Overall, one major advantage of decentralized systems is that they offer users the possibility to change service providers. Interconnection between different providers can be guaranteed through standardized interfaces.

3.2 Online and Mobile Social Networks

In [14], decentralized OSNs are surveyed. There is a categorization of architectural approaches that distinguishes between peer-to-peer (P2P), federated OSNs (F-OSN), and hybrid approaches. Going with these definitions, P2P OSNs leverage peer-to-peer-technology without a central server. F-OSNs use client/server infrastructure with multiple interconnected servers maintained by different providers or people. Hybrid approaches combine the two other approaches.

PeerSoN is a P2P OSN that focuses on privacy issues [5]. The project considers offline scenarios by proposing to store messages for offline users in a distributed hash table (DHT). *Prometheus* is another OSN project that utilizes a DHT [12]. A public-key infrastructure is employed, and users can define access control lists for access to their data. By collecting data from existing OSNs, emails, etc., weighted and named edges are constructed in the social graph. Data is replicated with trusted nodes in the DHT. The nodes should always be online, and mobile devices are not considered suitable nodes for the project. Overall, the focus here is both the preservation of privacy as well as an approximation of the *real* social graph.

Diaspora[6] is an example for a federated OSN. Here, users can choose between existing servers or set up their own server. The users benefit from this architectural approach with the ability to change providers, similar to how it is possible with email. Privacy concerns are reduced because no single service provider has all the information about all users. The user can be in control of his own data, but for the system as a whole, search and recommendations are more complicated because of the distributed nature of the data. According to [14], one advantage of F-OSNs in comparison with P2P OSNs is the typically better usability.

The main categories that are further analyzed in [14] are storage, access control, and interaction methods (i.e., communication protocols). According to Paul et al. [14], the main motivation for research in the field of decentralized OSNs is privacy concerns with centralized OSNs. Location-based scenarios and mobile devices are not considered.

Regarding MSNs, in [11], a distinction between three architectural approaches is made: *centralized*, *distributed*, and *hybrid architectures*. A centralized architecture is an architecture, in which a server delivers data to users on mobile devices. One example given is Facebook. For such a centralized architecture, one main advantage is that search and recommendations are better supported through knowledge about all users and the whole social graph. For the users of such a system, there are the disadvantages of the lock-in-effect and privacy concerns.

With "distributed" architectures, the authors describe approaches that do not rely on existing infrastructures, but

[3] `http://www.wechat.com/`
[4] `http://xmpp.org/`
[5] `https://www.facebook.com/`
[6] `https://diasporafoundation.org/`

opportunistic connections between mobile devices via Bluetooth or Wi-Fi [11]. One example is *MobiClique* [15]. Bluetooth is used for the transferal of data between users. Messages are forwarded between peers, instant communication as well as the arrival of a message cannot be guaranteed. While especially the uncertain delivery of a message is a major disadvantage, there are scenarios where such an approach might be useful, e.g., after a natural catastrophe that leaves no intact communication infrastructure behind. A hybrid architecture combines the two models, centralized and distributed, by allowing communication between users on mobile devices utilizing opportunistic connections, while allowing the retrieval of content via central servers.

Kayastha et al. further investigate community detection for location-based scenarios [11]. The presented approaches seem to assume knowledge about the whole social graph. We argue that when utilizing a distributed social graph, the whole social graph is not known to any one entity and that new mechanisms based on such a distributed social graph are needed for community detection and formation.

Bellavista et al. give a classification of MSNs on a different level in [3]. Here, a distinction is made by the parameters by which communities can be formed. The classification system spans pairs of mutually exclusive properties: user-centric (friends of someone) versus place-centric (e.g., fans of a specific museum), location-dependent (people at the location) versus location-independent (people anywhere). Additionally, regarding the storage and analysis of user metadata, Bellavista et al. distinguish between a spatial scope - by defining a global (data from anywhere) versus a local (data stored at the museum only) scope - and a temporal scope (if user metadata is stored for a longer time or only used instantaneously). In Section 4, we propose an new approach to categorize the creation of groups in OSNs.

Regarding the collection of context data on smartphones, some previous work has been done. With *P3 Systems*, Jones and Grandhi describe *people-to-people-to-geographical-places* [10]. The authors describe location-based services and what potential lies in them, especially with respect to online social networking and recommendations. Ideas of using users' proximity or location histories for recommender systems are presented.

In *WhozThat?*, Bluetooth is used for the broadcasting of social IDs, e.g., usernames from existing OSNs [1]. The users' phones can then retrieve information about other users in proximity by requesting public profile information from OSNs. The authors report about an implementation that leverages this approach to realize a music player that automatically plays music according to the tastes of the people in proximity. Security and privacy issues are mentioned, but not discussed further in [1].

In *SocialFusion*, the idea of collection context data is extended to not only data from existing OSNs, but also from mobile phones and nearby sensors [2]. In both *WhozThat?* and *SocialFusion*, OSN aspects, like the interaction between users, is not intended.

4. THREE-TIER VIEW OF THE SOCIAL GRAPH

In contrast to [3], we propose a new, easier, and in our eyes more useful approach to categorize group formation in OSNs. By detailing this approach, we will identify architectural requirements for an OSN architecture. The following three-tier view of the social graph will be motivated by an advanced user scenario:

- Location-based groups: short-term groups which can form spontaneously when people are at the same location.
- Context-based groups: medium-term groups which are formed based on context, e.g., being in the same class at university.
- Friends: long-term relations which are independent from context and location.

Using the terminology from [7], the friend-tier connections represent the *self-reported social network*, while context- and location-tier relations represent the *inferred* or *detected social network*.

To further illustrate these three tiers we will look at another user scenario. Imagine a chain of gyms Alice is registered with and regularly goes to. The machines she works out with transmit data about her workout to her smartphone. Furthermore, her wearable fitness tracker collects data. While she is working out, she listens to music on her smartphone. Her top ten artists are always automatically maintained by her mobile device. From the workout data on her phone, a fitness index is calculated. As all the data is on her phone, Alice is in control of her data. Should she decide to change to another gym chain, she has all her data with her and can carry settings for workout machines etc. with her.

Friend-Tier.

To share her fitness index or her top ten favorite music artists with her friends, Alice first needs to connect to her friends. What is needed here is a global directory in which Alice can look up other users by some form of ID or username. Such a DNS-like directory should further contain information about the presence of the friend, and information about how to contact this friend, e.g., IP address and protocol. Such a global directory can be used independently from actual OSN applications and would be used for looking up how to contact people, not storing social graph information. Connecting to her friends, Alice has formed her part of the distributed social graph in the friend-tier.

Context-Tier.

On the context-tier, Alice can be added to a group with the other people that work out at the same gym chain, independent of the branch location. As the global directory does only store information about how to contact other users, another form of directory is needed for creating a group based on context. The gym chain could offer a Wi-Fi hotspot that offers its customers to share their ID or username.

Additionally to just sharing IDs to form dynamic context-tier groups, a gym could have a TV in the lobby where the users that are currently working out are displayed. Additionally, the top music artists listened to by the people working out could be displayed. The users would stay in control of their data and could decide if and what data they share when entering the gym. Furthermore, the gym could use the provided context data to give personalized recommendations for products that fit a customer's workout style.

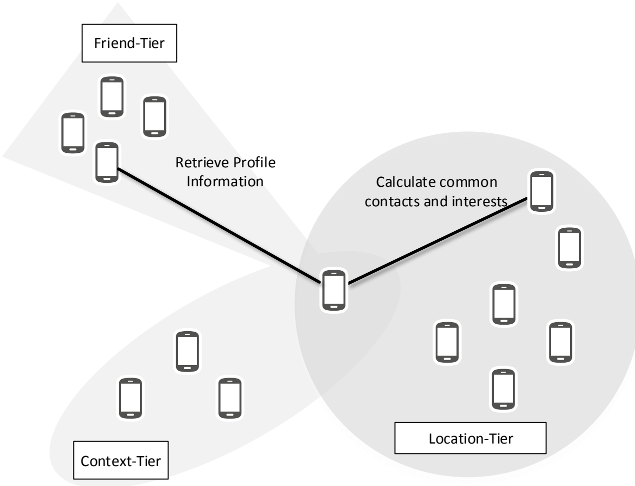

Figure 1: Three-tier Social Graph

Location-Tier.

During the summertime, some gym customers might prefer to go jogging outside in a park instead of working out in a gym. Here, location-tier groups can be formed. Some trusted online directory, e.g., provided by the gym chain or the city, is needed for the collection of users that are currently at the park. When such a directory offers its users the ability to publish more data than just their IDs, e.g., their fitness index, the directory can be used for searching, for example, to find new jogging partners with a similar fitness index.

5. T3 ARCHITECTURE

The three-tier view developed in the previous section induces requirements for a supporting OSN architecture which we will summarize in the following. We introduce an architecture called T3 satisfying these requirements.

To enable users to look up how to contact other users in decentralized OSN architectures, we propose a global directory, independent of the concrete OSN application or used communication protocol. Furthermore, a communication architecture that allows for direct communications between two smartphones over the Internet eliminates the need for a central server.

Additional decentralized directories can be implemented that allow users to publish their context data, enabling search and recommendations, as well as enabling support for the introduced context- and location-tier of the three-tiered view of the social graph. Leveraging the smartphone's capabilities, context data can be collected and processed locally. Standardized data formats are needed for the exchange between different users.

In Figure 1, we give an example of a social graph according to our three-tier approach. The smartphone in the center of the image is in a location-tier group with other users that are at the same location, as indicated by the dark gray circle. The light gray ellipse indicates membership in a context-tier group based on a common context, for example, the people that are enrolled in the same class at a university. The friend tier connections are indicated by the light gray triangle.

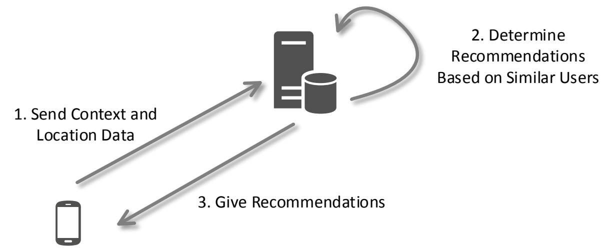

Figure 2: Recommendation Functionality of a Decentralized Directory

The black lines indicate direct interactions. The black line on the right hand side in Figure 1 indicates that the center smartphone is calculating its common contacts and interests with another smartphone in the area. The black line on the left hand side in Figure 1 indicates that from another smartphone, the center smartphone retrieves profile information.

Figure 2 shows the general concept of a decentralized directory. Imagine a travel scenario in which user Alice is querying a computer she encounters at a central train station of the city she is visiting. Alice sends location (history) and context data to the computer and receives personalized recommendations based on the data she sent. Additionally to recommendations, a decentralized directory can be utilized for the search for nearby restaurants or attractions.

In Figure 3, we outline our proposed T3 architecture for OSNs. In the following, we will detail the role the components play by indicating what of the six core features of OSNs (Section 2.3), they realize. The core components are:

Global Directory.

The global directory is used to look up how to contact other users for *interaction*, e.g., calling or messaging. Furthermore, information about where and how to retrieve a user's *personal profile* could be stored in the global directory as well.

Decentralized Directories.

Decentralized directories support *search* and *recommendations* by allowing users registration and by offering the possibility for users to publish context and location data. Compared to a single central directory service, one advantage of the approach of having multiple decentralized directories is trust: users might put higher trust in their gym chain that they visit regularly and where they see and know the people working there. Furthermore, users can choose which data to share with each directory, there is no need for a central directory that has all the data.

Smartphones.

On the users' smartphones, the *social graph* is stored in a distributed manner. Additionally, each user is in control of the locally stored *context information*. The *personal profile* could be hosted on the mobile devices, or on an external server for better availability. Some form of access control will be necessary to ensure proper privacy settings for users. For this, the three tiers could be utilized, or some additional role-based model - with roles like *co-workers* or *classmates* - could be used for access control automatization.

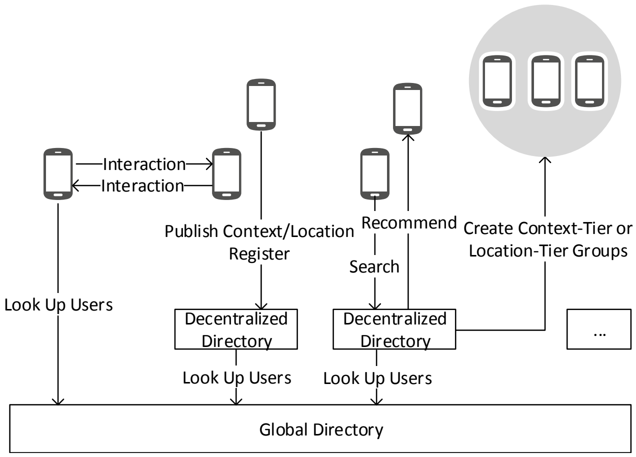

Figure 3: T3 Architecture for OSNs

The presented T3 architecture supports the three-tier concept of the social graph: Users can connect with friends by looking them up in the global directory; distributed directories help form context-tier and location-tier groups. The users' privacy is protected by utilizing a distributed social graph and by giving the users the control over their context data.

6. CONCLUSION AND OUTLOOK

In this paper, we argued that mobile user scenarios are essential for OSNs. We defined the core features of OSNs, introduced a three-tier view of the social graph, and proposed the decentralized OSN architecture T3 supporting this three-tier view, that utilizes a global directory, additional decentralized directories, and leverages the capabilities of current smartphones.

Future work includes the question how user IDs are created for the global directory, what authentication and authorization mechanisms are feasible, especially regarding registration, search, and recommendation with decentralized directories. Additionally, a feasible publish/subscribe mechanism should be investigated to enable users to automatically retrieve status updates from their friends. Furthermore, standardized interfaces for the communication between smartphones, as well as between smartphones and directories, are to be researched.

7. ACKNOWLEDGMENTS

This work has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No 645342, project reTHINK; and from project SONIC (grant number 01IS12056), which is funded as part of the SoftwareCampus initiative by the German Federal Ministry of Education and Research (BMBF) in cooperation with EIT ICT Labs Germany GmbH. The authors would like to thank Begüm İlke Zilci for valuable discussions about user scenarios.

8. REFERENCES

[1] A. Beach, M. Gartrell, S. Akkala, J. Elston, J. Kelley, K. Nishimoto, B. Ray, S. Razgulin, K. Sundaresan, B. Surendar, and others. Whozthat? Evolving an Ecosystem for Context-Aware Mobile Social Networks. *Network, IEEE*, 22(4):50–55, 2008.

[2] A. Beach, M. Gartrell, X. Xing, R. Han, Q. Lv, S. Mishra, and K. Seada. Fusing mobile, sensor, and social data to fully enable context-aware computing. In *Proceedings of the Eleventh Workshop on Mobile Computing Systems & Applications*, pages 60–65. ACM, 2010.

[3] P. Bellavista, R. Montanari, and S. K. Das. Mobile social networking middleware: A survey. *Pervasive and Mobile Computing*, 9(4):437–453, Aug. 2013.

[4] A. Bleicher. The anti-Facebook. *IEEE Spectrum*, 48(6):54–82, 2011.

[5] S. Buchegger, D. Schiöberg, L.-H. Vu, and A. Datta. PeerSoN: P2P Social Networking: Early Experiences and Insights. In *Proceedings of the Second ACM EuroSys Workshop on Social Network Systems*, SNS '09, pages 46–52, New York, NY, USA, 2009. ACM.

[6] A. Datta, S. Buchegger, L.-H. Vu, T. Strufe, and K. Rzadca. Decentralized Online Social Networks. In *Handbook of Social Network Technologies and Applications*, pages 349–378. Springer, 2010.

[7] N. Eagle, A. S. Pentland, and D. Lazer. Inferring friendship network structure by using mobile phone data. *Proceedings of the National Academy of Sciences*, 106(36):15274–15278, 2009.

[8] M. Falch, A. Henten, R. Tadayoni, and I. Windekilde. Business models in social networking. In *CMI Int. Conf. on Social Networking and Communities*, 2009.

[9] S. Göndör and H. Hebbo. SONIC: Towards seamless interaction in heterogeneous distributed OSN ecosystems. In *Proc. Wireless and Mobile Computing, Networking and Communications*, pages 407–412. IEEE, 2014.

[10] Q. Jones and S. Grandhi. P3 systems: putting the place back into social networks. *IEEE Internet Computing*, 9(5):38–46, Sept. 2005.

[11] N. Kayastha, D. Niyato, P. Wang, and E. Hossain. Applications, Architectures, and Protocol Design Issues for Mobile Social Networks: A Survey. *Proceedings of the IEEE*, 99(12):2130–2158, Dec. 2011.

[12] N. Kourtellis, J. Finnis, P. Anderson, J. Blackburn, C. Borcea, and A. Iamnitchi. Prometheus: User-controlled P2P Social Data Management for Socially-aware Applications. In *Middleware 2010*, pages 212–231. Springer, 2010.

[13] M. McLuhan. *Understanding Media: The Extensions of Man.* McGraw-Hill, New York, NY, USA, 1964.

[14] T. Paul, A. Famulari, and T. Strufe. A survey on decentralized Online Social Networks. *Computer Networks*, 75, Part A:437–452, Dec. 2014.

[15] A.-K. Pietiläinen, E. Oliver, J. LeBrun, G. Varghese, and C. Diot. MobiClique: Middleware for Mobile Social Networking. In *Proceedings of the 2nd ACM Workshop on Online Social Networks*, pages 49–54, New York, NY, USA, 2009. ACM.

Additional Information

Bibliographic Data
F. Beierle, S. Göndör, and A. Küpper, “Towards a Three-tiered Social Graph in Decentralized Online Social Networks,” in *Proc. 7th International Workshop on Hot Topics in Planet-Scale mObile Computing and Online Social neTworking (HotPOST)*. ACM, 2015, p. 1–6.

Pre-print from
https://beierle.de

Online at
http://dx.doi.org/10.1145/2757513.2757517

Authors
Felix Beierle

Sebastian Göndör

Axel Küpper

BibTeX

```
@inproceedings{Beierle2015HotPOST,
title = {{Towards a Three-tiered Social Graph in Decentralized Online Social Networks}},
author = {Beierle, Felix and G\"ond\"or, Sebastian and K\"upper, Axel},
booktitle = {{Proc. 7th International Workshop on Hot Topics in Planet-Scale mObile
Computing and Online Social neTworking (HotPOST)}},
publisher = {ACM},
year = {2015},
pages = {1--6},
doi = {10.1145/2757513.2757517}
}
```

Copyright Note
Copyright (c) 2015 ACM
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.