

Collaborating with Users in Proximity for Decentralized Mobile Recommender Systems

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

Tobias Eichinger
Service-centric Networking
Telekom Innovation Laboratories
Technische Universität Berlin
Berlin, Germany
tobias.eichinger@tu-berlin.de

Abstract—Typically, recommender systems from any domain, be it movies, music, restaurants, etc., are organized in a centralized fashion. The service provider holds all the data, biases in the recommender algorithms are not transparent to the user, and the service providers often create lock-in effects making it inconvenient for the user to switch providers. In this paper, we argue that the user’s smartphone already holds a lot of the data that feeds into typical recommender systems for movies, music, or POIs. With the ubiquity of the smartphone and other users in proximity in public places or public transportation, data can be exchanged directly between users in a device-to-device manner. This way, each smartphone can build its own database and calculate its own recommendations. One of the benefits of such a system is that it is not restricted to recommendations for just one user – ad-hoc group recommendations are also possible. While the infrastructure for such a platform already exists – the smartphones already in the palms of the users – there are challenges both with respect to the mobile recommender system platform as well as to its recommender algorithms. In this paper, we present a mobile architecture for the described system – consisting of data collection, data exchange, and recommender system – and highlight its challenges and opportunities.

Index Terms—Social Networking Services; Ubiquitous Computing; Mobile Computing; Smartphones; Context Data; Device-to-Device Communication; Recommender Systems

I. INTRODUCTION

Recommender systems are ubiquitous in several different domains of everyday life. They recommend restaurants to go to (e.g., Yelp, Google Maps), music to listen to (e.g., Spotify, Deezer), or movies to watch (e.g., IMDb, Netflix). The recommender algorithms usually are limited to recommending items that the platform offers (Spotify, Netflix) or at least items that are indexed with the provider (Yelp, Google Maps). There might be certain biases, for example, towards recommending products that create the biggest margin for the service provider. Furthermore, as the service provider is interested in retaining its users, there are certain lock-in effects that try to make the user stay with the current service provider instead of switching to another one. In that sense, from an organizational perspective, the examples given above are centralized – a single service provider has all the data and decides how recommendations are calculated.

On top of that, users have privacy concerns regarding the use of the data they share [1]. Recently, Google was fined 50

million euros due to a violation of the new European privacy laws (GDPR)¹.

The infrastructure that could be a solution for both challenges of lock-in effects and privacy concerns at the same time, is already in the palms of its users. The smartphone can store lots of information about its user and his/her interests, e.g., regarding preferred restaurants, music, or movies. Equipped with capabilities for device-to-device communication, users can exchange data between each other. When considering recommender systems based on content-based filtering or collaborative filtering, data about similar items and similar users is needed. Data about the properties of items can be retrieved through public APIs (e.g., Google Places, Spotify, Open Movie Database (OMDb)). Finding similar users might be simple with smartphones: spending time at the same location might imply similarity – at least to a certain degree. Additionally, from our previous research, other methods of determining similarity between users based on smartphone data are available [2], [3]. Thus, exchanging data between smartphones in proximity in a device-to-device fashion allows to create local databases that allow to filter for similar users. This data can be used for on-device recommender systems that are independent of external service providers.

Typically, each of the mentioned existing recommender systems only offers recommendations for a single user. The ubiquitous system we propose in this paper offers the possibility of spontaneous ad-hoc recommendations for groups of users in proximity.

Combining and expanding approaches from device-to-device computing [4], [5] and decentralized recommender systems [6]–[8], in this paper, we propose a modular architecture for recommender systems for virtually any domain, building on the existing infrastructure of smartphones. The architecture consists of collaborative data collection paired with data exchange via device-to-device communication and local recommender systems running on each device, supported by third-party service providers where appropriate. There are some challenges to overcome when developing such a plat-

¹<https://www.cnils.fr/en/cnils-restricted-committee-imposes-financial-penalty-50-million-euros-against-google-llc>

form. Some data is already readily available on smartphones, for example the most frequently visited locations. Other user preferences/ratings that cannot be assessed automatically might have to be entered manually or retrieved from external service providers, e.g., music listened to or favorite movies. While some short-distance wireless technologies, like NFC, Bluetooth, or WiFi Direct, are available on most modern smartphones, data exchange between users remains a challenge to be implemented for multiplatform apps.

The main contributions of this paper are:

- proposing a general modular architecture for a service-provider-independent mobile platform for recommender systems
- developing a comprehensive approach for collecting data and exchanging data in a device-to-device fashion for multiplatform apps (Android and iOS) to enable domain-independent recommender systems

The remainder of this paper is structured as follows: In Section II, we introduce a general architecture for the proposed mobile recommender system platform. In Section III and IV, we illustrate our approach for collecting and exchanging data. In Section V, we highlight what challenges and opportunities the recommender systems in our proposed platform will face, before pointing out related work in Section VI.

II. PROPOSED ARCHITECTURE

In order for the described architecture to be feasible, there are two main requirements:

- R1 The system has to be multiplatform.
- R2 The system has to be designed in a modular way.

Android and iOS are the remaining relevant mobile operating systems. They have a combined market share of about 100%². The main challenge this brings is related to data exchange, as we describe in Section IV. The system should be developed in a modular way (R2) in order to be able to exchange components easily. Consider short-distance wireless interfaces: In the past, infrared was used. Technological advances now offer larger transmission ranges, shorter connection times, and higher bandwidths via, for example, Bluetooth or WiFi Direct. Similarly, advances in recommender systems and machine learning might offer better recommendations, creating the need to replace the module or offload certain tasks to components available from external service providers.

In Figure 1, we illustrate our proposed general modular architecture. The three main components of the system are *Data Collection*, *Data Exchange*, and *Recommender System*. *Data Collection* is responsible for getting data about the user. *Data Exchange* is responsible for getting data from other users. The *Recommender System* utilizes all available data for recommending items to the user. The mobile OS provides components for sensors (for example for tracking the user’s location for inferring his/her favorite POIs) and wireless interfaces (for exchanging data).

²<https://www.statista.com/statistics/266136/global-market-share-held-by-smartphone-operating-systems/>

As we further analyze in Section III and IV, external service providers might be needed (or be useful) in order to retrieve metadata about items, utilize existing systems, or offload data or computational tasks. Figure 1 shows dashed lines for optional connections to third party service providers. *Data Collection* might use this to retrieve data about the user or to enrich already available data, e.g., find out the genre of the songs the user listened to. More details are given in Section III. The *Recommender System* can optionally be relayed to an external service provider.




Fig. 1: Architecture components of the proposed system.

III. COLLECTING DATA

We identify three different possibilities to retrieve user data:

a) *Tracking data automatically*: Smartphones contain a multitude of sensors that are often used for context-aware applications. Frameworks like the Google Awareness API³ yield the location, weather, etc. for the user. Such data can easily be tracked and so preferences or implicit ratings for, e.g., locations or POIs can be inferred. In previous work, we demonstrated an Android application that tracks a large variety of context data sources [9]–[11]. The data that can be tracked automatically on iOS might differ. In order to create a multiplatform system and ensure that the same data points are available on all systems, additional ways of retrieving the user’s ratings are necessary. Figure 2a shows the sequence diagram of automatic context data tracking.

b) *Retrieving data from existing service providers*: In order to minimize necessary user effort, the second method we suggest is retrieving data from existing service providers. For example, Spotify’s API enables application developers to fetch recently played tracks⁴. Regularly doing this yields a complete music listening history indicating implicit user ratings. Figure 2b shows the sequence diagram for the collection of data from a third-party service provider.

³<https://developers.google.com/awareness/>

⁴<https://developer.spotify.com/documentation/web-api/reference/player/get-recently-played/>

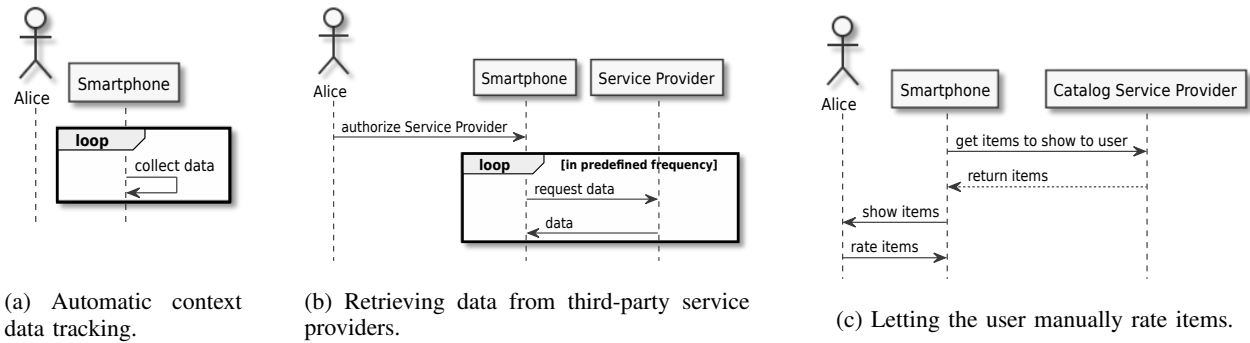


Fig. 2: Data collection mechanisms.

c) *Letting the user manually rate items*: For data that is neither automatically trackable nor available via third parties, the user should be able to enter it manually. By defining an ontology for categories and terms that can be exchanged between users, compatibility between the data from different collection methods can be ensured. Pre-defined categories can be movies, music, or restaurants, where recommender system are often used, but any other category would be possible too. Service providers like The Open Movie Database API⁵, for example, can be used to help employ globally valid identifiers for each item, in this case, for each movie. Figure 2c shows the sequence diagram for manual data collection. *Catalog Service Provider* denotes a service provider that offers structured information about a specific category, like the mentioned Open Movie Database.

IV. EXCHANGING DATA

The idea behind the exchange of data is that when users pass each other, their preference data, i.e., their ratings, are exchanged automatically in a device-to-device manner, building up each user's local database with more data. In order for such a system to work unobtrusively and without user interaction, the exchange of data should be done in the background (*broadcasting*), without establishing explicit connections between smartphones. While device-to-device communication has gained some attention in research, practical application is still lacking, especially when considering broadcasting, and especially when considering R1 (multiplatform app for both Android and iOS). In this section, we give an overview of related work, related applications, and propose a technical solution to be used in our proposed architecture.

During the advent of mobile phones, researchers suggested using device-to-device peer discovery and communication in order to stimulate social interactions. There are several papers between 2005 and 2010 describing exchanging data via Bluetooth, sometimes combining the direct data exchange with retrieving data from a central server [12]–[15]. In both [16] and [17], the authors suggest using WiFi SSIDs and the Bluetooth discovery protocol in order to exchange data between devices. Only very small amounts of data can be

transferred that way but the benefit is that no proper connection has to be established between two devices, thus allowing for broadcasting data to devices in proximity.

With the advent of smartphones, wireless interfaces improved and computing power increased, thus making it worth looking into more recent publications. A lot of work related to device-to-device communication focuses on challenges like offloading, content dissemination, or energy efficiency rather than the application layer [4]. Yet, there are some projects related to applications that describe actual implementations or suggestions for implementations. In [18], the authors use the WiFi ad-hoc mode on an Android device. This mode is not available by default, an extension had to be compiled into the Linux kernel. There is still a lack of support of WiFi ad-hoc since publication of that paper (2016), which shows that only a very limited set of devices would be able to run such an application. In [19], the authors use Bluetooth Low Energy (BLE) in IoT scenarios for ad-hoc communication. The framework they develop allows devices to communicate without predefined roles like client/server. Other recent works suggest using WiFi Direct for exchanging data between smartphones [20], [21], which is not readily supported on iOS devices.

Besides the academic work in the field, it is also worth looking into what related applications are actually available. In the following, we thus look into existing applications and technologies for device-to-device data transfer. Overall, such technologies seem to only operate with devices from the same manufacturer or the same mobile OS. Hand-held gaming devices from Nintendo and Sony are offering data exchange with nearby players in proximity⁶. This is a feature specific to each gaming system and does not work across devices from Nintendo and Sony. Both Apple and Google offer frameworks that enable software developers to interact with nearby devices. Apple's framework is called Multipeer Connectivity⁷ and only works with other Apple devices. Google introduced its framework Google Nearby⁸ with two different APIs: Nearby Connections and Nearby Messages.

⁶<https://www.nintendo.com/3ds/built-in-software/streetpass/how-it-works> and <http://us.playstation.com/psvita/apps/psvita-app-near.html>

⁷<https://developer.apple.com/documentation/multipeerconnectivity>

⁸<https://developers.google.com/nearby/>

⁵<http://www.omdbapi.com/>

Nearby Connections allows for device-to-device data transfer, but only between Android devices. Nearby Messages is only available when the devices are connected to the internet and allows only small payloads, but is available for both Android and iOS. There are multiplatform apps offering device-to-device data transfer solutions, e.g., SHAREit⁹. Such apps often let one user open a WiFi hotspot that is then joined by a second user. Data exchange via Bluetooth between Android and iOS devices is not readily available and typically requires explicit user interaction.

Reviewing related work in academia, software development frameworks, and apps, we summarize that the issue we are facing stems from the interaction of the following three factors:

- (a) *broadcasting*: In an optimal solution, exchanging preference data between passing users in proximity happens in the background without user interaction.
- (b) *multiplatform app*: In order to be able to provide the proposed system for virtually all smartphone users, the system has to be available on both Android and iOS.
- (c) *large message size*: In order to exchange user preferences data needed for a recommender system, larger messages have to be exchanged (in the range of several KB or MB).

Solutions for combinations of two of those aspects are available: (a) + (b): Using Google Nearby Messages, WiFi SSIDs, or Bluetooth Discovery Protocol messages, broadcasting between Android and iOS devices is possible, but only with small payloads. (b) + (c): Apps like SHAREit allow the explicit connection establishment between two devices in order to transfer large amounts of data. Typically, local WiFi hotspots are used. It might be possible to implement a combination of (a) + (c) with OS specific solutions (iOS: Multipeer Connectivity, Android: Google Nearby Connections).

For enabling a combination of all three aspects tough, a workaround is necessary. Building on the existing approaches, we present one workaround to facilitate a multiplatform approach with broadcasting that does not require user interaction and alleviates the issue of size limitations, see Figure 3.

First, Alice authorizes the system to access her account at some Cloud Storage Provider (CSP) like Dropbox, Google Drive, etc. Alternatively, she could use her own cloud storage. In some predefined frequency, Alice’s data is then uploaded to the CSP and shared via a public URL. This URL is then broadcasted via Google Nearby Messages or some other technique that allows multiplatform broadcasting. As only the URL is shared, which can be further shortened via a URL shortener service, the small payload restrictions of multiplatform broadcasting technologies should suffice. Another user, Bob in Figure 3, receives the broadcast with the URL and can download Alice’s publicly shared data. Optimizations like waiting for a WiFi connection can easily be implemented. Note that the only required user interaction by Alice or Bob is the authorization of the Cloud Storage Provider. Deeper investigations have to address the limitations posed by such

a solution, for example considering potential attack vectors created by such an approach.

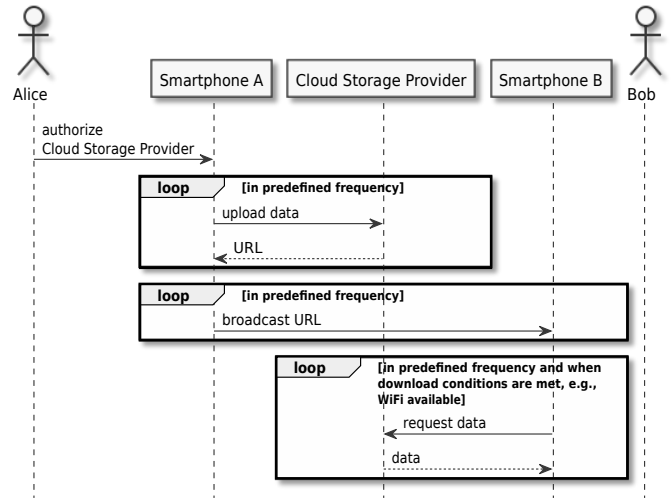


Fig. 3: Data exchange via a third-party cloud storage provider.

Even though technologies like BLE are in widespread use, a deeper look reveals that actually exchanging data between Android and iOS without user interaction is not a trivial task. Even if future developments might reveal that the implementation of multiplatform broadcasting gets even harder, our proposed architecture could still enable ad-hoc group recommendations with explicit connection establishment. A group of users in proximity could explicitly enable the data exchange among the group and locally calculate recommendations or query a third-party service provider for recommendations.

V. RECOMMENDING NEW ITEMS

Decentralized recommender systems traditionally use peer-to-peer networks [6], [7], [22]. Gossip protocols leverage commonly encountered small world properties of overlay topologies in file-sharing networks and allow to find and gather peers with similar preferences quickly and reliably. Once established, communities of interest exchange item ratings among each other. In contrast to such an approach, our ad-hoc fashion of connection and data exchange between smartphones is a network that is essentially fully disconnected. Furthermore, most of the given approaches of decentralized recommender systems deal with personal computers, while we focus on ubiquitous scenarios with mobile devices.

Other related fields are those of ubiquitous recommender systems and context-aware recommender systems (CARS). They consider items in proximity or consider the user’s current context while recommending items, respectively [23]. In contrast to these approaches, our proposed system is general in the sense that any type of item can be recommended, independent of the item’s physical proximity or the user’s context.

Another field, that has gained less attention in industry and academia, is that of group recommender systems [24], [25]. With its ad-hoc nature and immediate preference data exchange, our proposed system is ideally suited to be used for

⁹<https://www.usshareit.com/>

pervasive group recommendation scenarios. Exchanging data between several users in a group setting, a local recommender system can calculate recommendations based on the given data, considering the preferences of each user. When utilizing an external service provider for a recommendation, most likely, before contacting it, the preferences of each group member have to be combined into one group profile as most providers will only recommend items for a single user.

When employing a local recommender system on the smartphone, additional data is needed. For content-based filtering, the properties of items have to be known. Third-party service providers can help with retrieving such needed metadata about items. For user-based collaborative filtering, information about the similarity of users is utilized. While services like Spotify or Netflix have very large databases with millions of users, the local databases in our proposed architecture will be much smaller and thus there is a lower likelihood of finding similar users.

We see two possible solutions for this problem. First, we could let each user disseminate more than just his/her own item preferences/ratings and let him/her also send data from previous encounters – this would also address the cold start problem new users will face. Another approach is to calculate the similarity of users in a different way, independent of the users’ ratings. In psychology, the *propinquity effect* is the well-studied effect that physical proximity is a good predictor of forming interpersonal bonds [26], [27]. Having unique identifiers for each user and counting the number of times and/or the duration of being in proximity would then likely predict a higher bond. Additional methods are available for determining similarity in proximity-based applications. In [2], we developed and evaluated a method for estimating similarity based on the exchange of the users’ context data using probabilistic data structures in device-to-device scenarios. In [3], we developed a privacy-preserving method for determining the similarity of two users based on their text messaging data. Both of those methods can be implemented in our proposed architecture to find similar users, without having the need to have users that rated the same items. Future work will have to show to what extent those similarity metrics and the propinquity effect yield valuable similarity indications for user-based collaborative filtering. Future work could also investigate the feasibility of approaches like federated learning, effectively exchanging trained models or updates to models for recommendations [28].

VI. RELATED WORK

Most related work is from the field of social networking services. Often, the related work does not specifically focus on recommender systems but on other aspects like privacy, utilizing opportunistic networks, or encouraging real-life user interactions.

In our own previous work, we outlined a general concept for a social networking service utilizing context data from smartphones, focusing on how connections between users can be established [29]. In [30], we proposed to utilize the

relationship between smartphone data and personality traits for friend recommendations in device-to-device scenarios.

In [12], the authors let users send their identifiers of a social networking services (via Bluetooth) to enable receivers to visit their publicly available profile. The idea is to encourage social interaction. The authors of [31] developed a framework for ephemeral social networks. The goal is content dissemination in opportunistic networks in scenarios with large crowds like sports events. Westerkamp et al. propose a decentralized social networking service in [32]. The goal is avoiding censorship by utilizing the Ethereum Name Service for identity management and by using self-hosted data storages – or trusted third-party service providers – for each user. With such a solution, the challenges of lock-in effects and lack of privacy are tackled and the user is put in control of his/her data. In our paper, we followed a similar approach, focusing on a different application domain (recommender systems) and on pervasive scenarios.

[5] is a short survey paper that highlights the challenges of device-to-device computing. Regarding the wireless network interfaces, the authors consider cellular networks, WiFi, Bluetooth, and NFC. In our paper, we gave details about the challenges when implementing data exchange between devices.

Regarding decentralized recommender systems, [6] and [8] follows similar approaches compared to our proposed system. In [6], decentrally stored data from the web is used for a recommender system running on the user’s personal computer. Since that paper’s publication (2005), the development of mobile devices enable mobile and ubiquitous scenarios depicted in this paper. In [8], the authors propose that smartphones exchange data in a device-to-device fashion and calculate their own recommendation via collaborative filtering. The focus of that paper is on the recommender algorithm that is evaluated with a music data set. For device-to-device communication, WiFi Direct is proposed.

VII. CONCLUSION AND FUTURE WORK

Current recommender systems often exhibit a lock-in effect for the user and are connected to privacy concerns. Their recommendations are typically for single users rather than groups and might be biased according to the interests of the providing platform. We proposed a decentralized mobile architecture for recommender systems that leverages the preferences/ratings from users that are or have been in proximity. The introduced system runs on the users’ smartphones and utilizes existing external third-party service providers.

The system consists of three main parts, *data collection*, *data exchange*, and *recommender system*. We developed three ways to collect data: tracking context data automatically, retrieving data from third-party service providers, and letting the user manually rate items. We highlighted that while short-range wireless transmission technologies are implemented on all modern smartphones, exchanging larger amounts of data without user interaction on a system available for both Android and iOS remains a challenging task. We proposed a

workaround by broadcasting URLs of cloud storage providers from which the receivers can then download the sender's data. Regarding the recommendation process, when considering user-based collaborative filtering, we looked into ways for determining the similarity between users, including methods independent of the users' ratings.

Future work will be to implement and test the proposed device-to-device data exchange in order to evaluate its reliability, transmission times in real user scenarios, and battery consumption. Regarding the recommender system, future work includes the implementation of a mobile recommender engine. A simulation with a real data set can help evaluate the quality of the recommendations that such a system can provide. Furthermore, potential attack vectors and the overall security of the system should be investigated.




ACKNOWLEDGMENT

We are grateful for the support provided by Robert Staake, Jan Pokorski, Simone Egger, and Yong Wu.

REFERENCES

- [1] M. Falch, A. Henten, R. Tadayoni, and I. Windekilde, "Business models in social networking," in *CMJ Int. Conf. on Social Networking and Communities*, 2009.
- [2] F. Beierle, "Do You Like What I Like? Similarity Estimation in Proximity-Based Mobile Social Networks," in *Proc. 2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications (TrustCom)*. IEEE, 2018, pp. 1040–1047.
- [3] T. Eichinger, F. Beierle, S. U. Khan, R. Middelanis, V. Sekar, and S. Tabibzadeh, "Affinity: A System for Latent User Similarity Comparison on Texting Data," in *Proc. 2019 IEEE International Conference on Communications (ICC)*. IEEE, 2019 (in press).
- [4] M. Ahmed, Y. Li, M. Waqas, M. Sheraz, D. Jin, and Z. Han, "A Survey on Socially Aware Device-to-Device Communications," *IEEE Communications Surveys Tutorials*, vol. 20, no. 3, pp. 2169–2197, 2018.
- [5] W. Zhang, H. Flores, and P. Hui, "Towards collaborative multi-device computing," in *2018 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*. IEEE, 2018, pp. 22–27.
- [6] C.-N. Ziegler, "Semantic Web Recommender Systems," in *Current Trends in Database Technology - EDBT 2004 Workshops*, ser. Lecture Notes in Computer Science, W. Lindner, M. Mesiti, C. Türker, Y. Tzitzikas, and A. I. Vakali, Eds. Springer, 2005, pp. 78–89.
- [7] R. Baraglia, P. Dazzi, M. Mordacchini, and L. Ricci, "A peer-to-peer recommender system for self-emerging user communities based on gossip overlays," *Journal of Computer and System Sciences*, vol. 79, no. 2, pp. 291–308, Mar. 2013.
- [8] L. N. Barbosa, J. Gemmell, M. Horvath, and T. Heimfarth, "Distributed User-Based Collaborative Filtering on an Opportunistic Network," in *2018 IEEE 32nd International Conference on Advanced Information Networking and Applications (AINA)*. IEEE, 2018, pp. 266–273.
- [9] F. Beierle, V. T. Tran, M. Allemand, P. Neff, W. Schlee, T. Probst, R. Pryss, and J. Zimmermann, "TYDR – Track Your Daily Routine. Android App for Tracking Smartphone Sensor and Usage Data," in *2018 IEEE/ACM 5th International Conference on Mobile Software Engineering and Systems (MOBILESoft)*. ACM, 2018, pp. 72–75.
- [10] —, "Context Data Categories and Privacy Model for Mobile Data Collection Apps," *Procedia Computer Science*, vol. 134, pp. 18–25, 2018.
- [11] F. Beierle, V. T. Tran, M. Allemand, P. Neff, W. Schlee, T. Probst, J. Zimmermann, and R. Pryss, "What data are smartphone users willing to share with researchers? Designing and evaluating a privacy model for mobile data collection apps," *Journal of Ambient Intelligence and Humanized Computing*, 2019 (in press).
- [12] A. Beach, M. Gartrell, S. Akkala, J. Elston, J. Kelley, K. Nishimoto, B. Ray, S. Razgulin, K. Sundaresan, B. Surendar, M. Terada, and R. Han, "WhozThat? Evolving an Ecosystem for Context-Aware Mobile Social Networks," *IEEE Network*, vol. 22, no. 4, pp. 50–55, 2008.
- [13] N. Eagle and A. Pentland, "Social Serendipity: Mobilizing social software," *Pervasive Computing, IEEE*, vol. 4, no. 2, pp. 28–34, 2005.
- [14] A.-K. Pietiläinen, E. Oliver, J. LeBrun, G. Varghese, and C. Diot, "MobiClique: Middleware for Mobile Social Networking," in *Proc. 2nd ACM Workshop on Online Social Networks (WOSN)*. ACM, 2009, pp. 49–54.
- [15] Z. Yang, B. Zhang, J. Dai, A. Champion, D. Xuan, and D. Li, "E-SmallTalker: A Distributed Mobile System for Social Networking in Physical Proximity," in *2010 IEEE 30th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 2010, pp. 468–477.
- [16] J. Manweiler, R. Scudellari, and L. P. Cox, "SMILE: Encounter-based trust for mobile social services," in *Proc. 16th ACM Conference on Computer and Communications Security*. ACM, 2009, pp. 246–255.
- [17] J. Teng, B. Zhang, X. Li, X. Bai, and D. Xuan, "E-Shadow: Lubricating Social Interaction Using Mobile Phones," *IEEE Transactions on Computers*, vol. 63, no. 6, pp. 1422–1433, Jun. 2014.
- [18] Z. Lu, G. Cao, and T. L. Porta, "Networking smartphones for disaster recovery," in *2016 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE, 2016, pp. 1–9.
- [19] J. Yang, C. Poellabauer, P. Mitra, J. Rao, and C. Neubecker, "BlueNet: BLE-based ad-hoc communications without predefined roles," in *2017 IEEE SmartWorld, Ubiquitous Intelligence Computing, Advanced Trusted Computing, Scalable Computing Communications, Cloud Big Data Computing, Internet of People and Smart City Innovation (SmartWorld/SCALCOM/UIC/ATC/CBDCOM/IOP/SCI)*. IEEE, 2017, pp. 1–8.
- [20] Z. Mao, J. Ma, Y. Jiang, and B. Yao, "Performance Evaluation of WiFi Direct for Data Dissemination in Mobile Social Networks," in *2017 IEEE Symposium on Computers and Communications (ISCC)*. IEEE, 2017, pp. 1213–1218.
- [21] J. Shu, S. Kosta, R. Zheng, and P. Hui, "Talk2Me: A Framework for Device-to-Device Augmented Reality Social Network," in *2018 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE, 2018, pp. 1–10.
- [22] G. Ruffo and R. Schifanella, "A Peer-to-Peer Recommender System Based on Spontaneous Affinities," *ACM Transactions on Internet Technology*, vol. 9, no. 1, pp. 1–34, Feb. 2009.
- [23] C. Mettouris and G. A. Papadopoulos, "Ubiquitous recommender systems," *Computing*, vol. 96, no. 3, pp. 223–257, Mar. 2014.
- [24] A. Crossen, J. Budzik, and K. J. Hammond, "Flytrap: Intelligent group music recommendation," in *Proc. 7th International Conference on Intelligent User Interfaces*. ACM, 2002, pp. 184–185.
- [25] F. Beierle, K. Grunert, S. Göndör, and A. Küpper, "Privacy-aware Social Music Playlist Generation," in *Proc. 2016 IEEE International Conference on Communications (ICC)*. IEEE, 2016, pp. 5650–5656.
- [26] D. M. Marvin, "Occupational Propinquity as a Factor in Marriage Selection," *Quarterly Publications of the American Statistical Association*, vol. 16, no. 123, pp. 131–150, 1918.
- [27] L. Festinger, S. Schachter, and K. Back, "The spatial ecology of group formation," *Social pressures in informal groups*, pp. 141–161, 1950.
- [28] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-Efficient Learning of Deep Networks from Decentralized Data," in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*. PMLR, 2017, pp. 1273–1282.
- [29] 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, pp. 1–6.
- [30] F. Beierle, K. Grunert, S. Göndör, and V. Schlüter, "Towards Psychometrics-based Friend Recommendations in Social Networking Services," in *2017 IEEE International Conference on AI & Mobile Services (AIMS)*. IEEE, 2017, pp. 105–108.
- [31] A. Boutet, S. Frenot, F. Laforest, P. Launay, N. Le Sommer, Y. Maheo, and D. Reimert, "C3PO: A Network and Application Framework for Spontaneous and Ephemeral Social Networks," in *Web Information Systems Engineering – WISE 2015*, ser. Lecture Notes in Computer Science, J. Wang, W. Cellary, D. Wang, H. Wang, S.-C. Chen, T. Li, and Y. Zhang, Eds. Springer, 2015, pp. 348–358.
- [32] M. Westerkamp, S. Göndör, and A. Küpper, "Tawki: Towards Self-Sovereign Decentralized Communication," in *Proc. IEEE International Conference on Decentralized Applications and Infrastructures (DAPPCON)*. IEEE, 2019 (in press).

Additional Information

Bibliographic Data	F. Beierle and T. Eichinger, "Collaborating with Users in Proximity for Decentralized Mobile Recommender Systems," in <i>Proceedings of the IEEE 16th International Conference on Ubiquitous Intelligence and Computing (UIC)</i> . IEEE, 2019.
Pre-print from	https://beierle.de
Online at	(in press)
Authors	Felix Beierle   Tobias Eichinger 
BibTeX	<pre>@inproceedings{Beierle2019UIC, title = {{Collaborating with Users in Proximity for Decentralized Mobile Recommender Systems}}, author = {Beierle, Felix and Eichinger, Tobias}, booktitle = {{Proceedings of the IEEE 16th International Conference on Ubiquitous Intelligence and Computing (UIC)}}, publisher = {IEEE}, year = {2019} }</pre>
Copyright Note	IEEE Copyright Notice Copyright (c) 2019 IEEE Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.