

On the Feasibility of Opportunistic Ad Hoc Music Sharing

Barbara Keller, Philippe von Bergen, Roger Wattenhofer, Samuel Welten
Computer Engineering and Networks Laboratory
Information Technology and Electrical Engineering Department
ETH Zurich, Switzerland
firstname.lastname@tik.ethz.ch

ABSTRACT

We use the mobility traces of the Lausanne data collection campaign to demonstrate the feasibility of opportunistic ad-hoc music sharing over WiFi using smartphones. Instead of specific file requests we simulate the sharing of music using a compact representation of a user's music taste that is based on a measure of music similarity. Our simulations show that one's music collection can be extended with suitable music by sharing music over locally established wireless LAN connections.

Categories and Subject Descriptors

H.3 [Information Storage and Retrieval]: Information Search and Retrieval

General Terms

Experimentation

Keywords

music recommendation, ad hoc networks, file sharing

1. INTRODUCTION

Over a century ago, people wanting to hear music either had to make music themselves or go to the opera. With the invention of sound recording technologies, music was more easily accessible and people were able to build music collections. Since then huge technological improvements have made it simpler to collect music. When the tape recorder became popular in the sixties, it was possible to copy tapes at home. With the emergence of the compact disc in the seventies, it became even more popular to copy music at home as this procedure allows to copy music without suffering loss of quality. In addition, this change made music more mobile; people started carrying around devices to listen to music. With the widespread usage of the mp3 format in the nineties a new phenomenon emerged. Suddenly music

was not bound to a physical medium anymore, allowing for even larger music collections, and sharing with people without the need to meet physically. With the development of mobile phones to so called smartphones, people started to store their music on their phones. Although the way of storing music has eased, the way of getting to know new music drags behind. People still depend on recommendations of friends, journals and blogs or have to browse through music collections of stores or online communities. It is still cumbersome to find out, which music one may like. Even with friends who share one's music taste it is tedious to compare music collections to find music one does not own already. It would be really convenient if smartphones would automatically get appropriate new music once one is in the proximity of a person with a music collection that is potentially matching one's music taste.

Over the last years more and more artists have released songs and albums under a creative commons license. This essentially means that people are not only technically able to share music with everyone, but they are also legally allowed to do so. With the increasing spread of smartphones and creative commons music, the vision of opportunistic music sharing seems to become more and more realistic. But is it really feasible? Will people meet often enough other people that share their taste of music?

Our main contribution is to show that it is possible to extend personal music collections effectively by exchanging music files over local wireless network connections. We use the data collected during the Lausanne data collection campaign to provide a proof of concept. This is achieved by analyzing how music would spread if the participants had used such an automated music exchange program. For this analysis we assume that a smartphone is able to detect another peer in its proximity and is able to build up a direct wireless connection. This wireless connection is first used to exchange a compact representation of music interests, and then to share suitable music files with the other party.

2. RELATED WORK

The idea to exchange information in a opportunistic fashion over locally established networks has been around for several years (e.g. [15]). Even though such networks would be possible with many wireless technologies, most recent approaches focus on wireless LAN. Trifunovic et al. proposed a opportunistic network system for Android smartphones that uses on the one hand existing non-encrypted access points and on the other hand the devices' capability to create an access point on their own [20]. Other recent publications

in this area concentrate on protocols for multi-hop routing [3] or efficient information retrieval [8, 14]. In contrast to those publications we do not explicitly model multi-hop transmissions because we assume the network to be sparsely connected and want to avoid unnecessary transmissions to save energy. Furthermore, these approaches only consider explicit data requests opposed to our similarity-based, un-specific requests.

The feasibility and efficiency of such local opportunistic networks are commonly assessed through real life mobility traces. As smartphones are equipped with more and more sophisticated sensors and offer the possibility to run custom unobtrusive applications, several research groups collected and published this kind of data collections. Popular data sets include the *Reality Mining* project of Eagle et al. [4] and the *CenceMe* system of Miluzzo et al. [18]. The data set from the Lausanne data collection campaign that is used in our study is comparable to these projects in terms of size and collected data [13].

As the music is spreading among the participants of our study, our research is also related to insights on information diffusion. The diffusion of information, ideas and viruses in different networks is a broad area of research [1, 11, 5]. A lot of work focuses on the analysis of the spread of biological viruses. A widespread approach is to use mathematical models to predict and minimize the spread [2, 9]. As digital viruses that use Bluetooth and ad hoc networks as their means of transfer have similar spreading patterns, analogue approaches are used in [19, 21]. We do not rely on an underlying mathematical model for the diffusion of music, but understanding the underlying pattern as it is done in [6] could improve the efficiency of our application. Inverse to our analysis, Madan investigated in [17] the frequency and strength of social interactions using, among other information, the spread of songs between dorm mates. He used the basic assumption that friends generally have a similar taste of music and tend to share songs, to assess the connectivity among the people.

Modeling a person’s music taste and embedding users and music in an Euclidean space for the purpose of recommending music was described by different researchers [12, 7]. Goussevskaia et al. also proposed to use such embeddings for decentralized music sharing. The embedding by Kuhn et al. that is used in this paper has similar properties as the aforementioned embeddings [16]. The following experiments might therefore also be feasible with one of those music similarity measures.

3. METHODOLOGY

To test the feasibility of automatic opportunistic music sharing we were mainly interested in real life motion traces of several persons over a longer period of time. The data set of the Nokia Mobile Data Challenge (NMDC) provided exactly this kind of data: GPS-based location data and WLAN-based location data of smartphones of 38 persons, recorded over the period of one year [10].

The following procedure was used to simulate opportunistic ad hoc music sharing: First, we identified the periods of time when two people were close enough to establish a direct wireless LAN connection (from now on denoted as a *meeting*). Then, we computed for each participant her music taste and used this taste along with the meetings to simulate sharing of music.

3.1 Meetings

The GPS-location log of each of the subjects is a list of timestamps and according GPS-based and WLAN-based location coordinates. Our first goal was to extract meetings. To do so we had to make two basic assumptions:

1. Two persons can establish a direct connection once the distance of the two is below a certain threshold (we used 50 or 30 meters in our simulations).
2. A person is at the latest known location until information about a new location is available or 10 minutes passed.

As the assumption that a person stays at the latest known location is too strong we limited this assumption to 10 minutes to prevent wrong meeting detections for persons who did turn off their phone at some point. We are aware that there are situations, where these assumptions can be violated, but they might hold in the general case.

To find meetings we searched over all possible pairings of persons. For each pair we identified GPS timestamp entries closest to each other and checked if the difference between these timestamps did not exceed 10 minutes. Then we checked if the two participants were at most 30 respectively 50 meters away from each other using the provided GPS data. The duration of such a meeting is determined as the difference between the first and last timestamp fulfilling these requirements. Note that the geographical distance between two persons is calculated assuming a plane surface, without respecting the topology. As we are only interested in short distances up to 50 meters we consider this a reasonable assumption.

3.2 Sharing

To simulate the automatic sharing process we need to compute a representation of the music taste of each participant. This notion of music taste is based on the collection of songs on the phone of a person. Even though the data of the *Lausanne data collection campaign* incorporates lists of the songs stored on these devices, these lists were not useful for our purposes because only a few of the participants stored a reasonable amount of music on their smartphones. As one of our basic assumptions is that such a music sharing application is used by music affine persons, this data did not seem to be reasonable. To overcome this issue we assigned each participant a music collection of a random user of the Android music player *jukefox*. This means we used data about people’s music preferences that was collected in the scope of another project and could be used to simulate realistic music collections on the devices of the participants. Each participant was assigned to a music set between 9 and 827 songs. The mean size of the assigned collections was 264 songs.

3.2.1 Representing Music Taste

Once two persons meet and want to exchange music, their devices have to exchange requests that describe in what kind of music they are interested in. This description should be in a form that allows the other party to decide which of the songs on the device could match the description without consulting a central server. A solution which allows decentralized computation of recommendations is to use an Euclidean music similarity embedding as proposed in [12] or

Algorithm 1 Simulating Ad Hoc Music Sharing

```
Gather music collections and tastes for all users
for all meetings do
  Deduct initialization time of meeting duration
  for all songs in collections do
    if song is suitable then
      if enough time left then
        Add song to receiving collection
        Deduct transfer time
      else
        break
    end if
  end if
end for
end for
```

[7]. Such an embedding allows to assign Euclidean coordinates to every song in a collection: the smaller the distance between two songs, the more similar they are. This property is used to model the music taste of a person as spheres in this space. The same idea was also presented by Goussevskaia et al., who state that an Euclidean music space enables a volume-based representation of the music tastes and therefore easy sharing of music [7]. The embedding that is used in our study is described in [16]. It is computed based on usage data of over 1 million users of the online radio service *last.fm* and user-assigned tags from the same platform.

Based on the songs of the music profile which is assigned to a person we compute her music taste. We apply a k-means algorithm to the positions of all songs in the collection of this user. Each of the class centers is the center of a sphere where the radius is defined as the longest distance between the center and the songs associated to this center. This representation of the music interests allows a user to send these spheres (centers and radii) to another user, who will be able to decide how well the songs of her own collection fit this taste (i.e. if her songs fall into a sphere of the other user and how close to the sphere center they are). Thanks to this representation of music taste the reception of new matching songs does not provoke huge changes of taste. Although new songs would lead to a slight drift of the sphere center we did not adapt a users taste after receiving new songs in our simulations.

3.3 Simulation

Using the described meetings and music collections, we ran several simulations concerning the sharing of music files. The exact simulation algorithm is shown in Algorithm 1. The method *song is suitable* in Algorithm 1 consists of two parts. On the one hand, it checks if this song is already present in the receivers collection and on the other hand, it is a simple threshold decision based on the music taste representation. If the song lies within a sphere and its distance to the sphere center is below a certain threshold relatively to the radius, it is classified as suitable, and as not suitable otherwise.

4. RESULTS

4.1 Meetings

The first results yielded from our simulations were the number and durations of meetings contained in the NMDC




Figure 1: This heat map shows the number of meetings between all pairs of participants during one year. Note that some of participants seem to be highly connected and meet a lot, whereas others meet only a few times.

data set. The previously described assumptions led to a number of meetings between 17.562 and 30.556 over all participants, depending on the parameters. Table 1 summarizes our findings using different parameters for maximal distance and minimally required meeting duration. Note that by requiring a maximal distance of 30 meters and a minimal duration of 90 seconds, there are still over 24.000 meetings. Assuming a discover- and connection-set-up-time of 30 seconds and a data transfer rate of 25 Mbit/s, 90 seconds are enough to transfer an average album. Moreover, in the subsequently described simulations we used a setup time of 30 seconds and a transfer time of 2 seconds per song.

The heat map in Figure 1 visualizes the number of meetings between the individual participants. Some of the participants seems to meet a lot more often than others. Only one subject did not meet with any other participant at all, during the whole data collection period. The distribution of the meeting durations for all meetings for a maximal distance of 30 meters can be found in Figure 2. The mean duration of all meetings is 816 seconds, if the maximal distance is 50 meters and 810 seconds for 30 meters maximal distance. The peculiar increase of meetings of a duration of 600 seconds could be an artifact of the choice of the allowed time gap between two timestamps (see Section 3.1).

4.2 Sharing

To assess the effectiveness of the opportunistic sharing, we defined for each participant a target set of songs. This target




Figure 2: A histogram of the durations of the meetings between all participants for a maximal distance of 30 meters.

| Maximal distance | Minimal duration | Total number of meetings | Mean meetings per user |
|------------------|------------------|--------------------------|------------------------|
| 50 meters | 1 second | 30556 | 804 |
| 50 meters | 90 seconds | 25925 | 682 |
| 50 meters | 300 seconds | 18618 | 490 |
| 30 meters | 1 second | 28804 | 758 |
| 30 meters | 90 seconds | 24390 | 642 |
| 30 meters | 300 seconds | 17562 | 462 |

Table 1: The amount and duration of meetings, depending on the maximally allowed distance and minimal duration.

set consists of all songs of all participants that were in one of the spheres of the respective music taste representation, i.e. for which the method *song is suitable* of Algorithm 1 would return true (see also Section 3.2). This resulted in target sets of songs that fit the participants' music tastes. These target sets had a mean size of 480 songs per user. Even though this number is much bigger than the average collection size of a participant, it is still quite small and specific, considering that the music profiles of all participants together contained 9025 distinct songs.

For each of the participants we evaluated the percentage of her target set she received after simulating all the meetings in chronological order. Figure 3 shows the target set percentages of all participants for a simulation using 30 meters as a maximal meeting distance and respecting the time restrictions that arise from the connection and transmission delays described before. On average the users got 46.8% of the respective target set, which means that the average music collection would have been extended by 206 songs.

One user, for example, started with a collection mostly containing The Beatles, Nine Inch Nails, Jack Johnson, Pink Floyd and Norah Jones. During the simulation she received songs of The Beatles, Nine Inch Nails and Jack Johnson, but she also received songs of John Mayer and Johnny Cash, which she did not possess before at all.

5. CONCLUSION

Even though a real world deployment of a sharing software will have to show if the motion traces of the participants of the Lausanne data collection campaign are representative to model a set of music affine persons, our simulations have




Figure 3: The histogram displays the distribution of the sharing success among the participants. Note that in the mean, people got 46.8% of the recommended music (using a max. distance of 30 meters).

shown that an opportunistic decentralized approach to detecting new music is feasible. Moreover, new wireless technologies might enable an even faster detection of neighboring peers and allow to transfer data at higher rates over larger distances.

6. REFERENCES

- [1] L. Bettencourt, A. Cintron-Arias, D. I. Kaiser, and C. Castillo-Chavez. The power of a good idea: Quantitative modeling of the spread of ideas from epidemiological models. *Physica A: Statistical Mechanics and its Applications*, 364, 2006.
- [2] V. Colizza, A. Barrat, M. Barthelemy, A.-J. Valleron, and A. Vespignani. Modeling the Worldwide Spread of Pandemic Influenza: Baseline Case and Containment Interventions. *PLoS Med*, 4, jan 2007.
- [3] G. Ding and B. Bhargava. Peer-to-peer file-sharing over mobile ad hoc networks. In *Proceedings of the Second IEEE Annual Conference on Pervasive Computing and Communications Workshops, PERCOMW '04*, 2004.
- [4] N. Eagle and A. (Sandy) Pentland. Reality mining: sensing complex social systems. *Personal Ubiquitous Comput.*, 2006.
- [5] S. Eubank, H. Guclu, K. Anil, M. Marathe, A. Srinivasan, Z. Toroczkai, and N. Wang. Modelling disease outbreaks in realistic urban social networks. *Nature*, 2004.
- [6] M. Gonzalez, C. Hidalgo, and A. Barabasi. Understanding individual human mobility patterns. *Nature*, 453, 2008.
- [7] O. Goussevskaia, M. Kuhn, M. Lorenzi, and R. Wattenhofer. From Web to Map: Exploring the World of Music. In *IEEE/WIC/ACM International Conference on Web Intelligence (WI)*, Sydney, Australia, December 2008.
- [8] O. R. Helgason, E. A. Yavuz, S. T. Kouyoumdjieva,

- L. Pajevic, and G. Karlsson. A mobile peer-to-peer system for opportunistic content-centric networking. In *Proceedings of the second ACM SIGCOMM workshop on Networking, systems, and applications on mobile handhelds*, MobiHeld '10, 2010.
- [9] L. Hufnagel, D. Brockmann, and T. Geisel. Forecast and control of epidemics in a globalized world. *Proceedings of the National Academy of Sciences of the United States of America*, 101, 2004.
- [10] Juha K. Laurila, Daniel Gatica-Perez, Imad Aad, Jan Blom, Olivier Bornet, Trinh-Minh-Tri Do, Olivier Dousse, Julien Eberle, and Markus Miettinen. The mobile data challenge: Big data for mobile computing research. In *in Proceedings Mobile Data Challenge by Nokia Workshop, in conjunction with Int. Conf. on Pervasive Computing*, Pervasive, 2012.
- [11] R. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized rumor spreading. In *In IEEE Symposium on Foundations of Computer Science*, 2000.
- [12] M. Khoshneshin and W. N. Street. Collaborative filtering via euclidean embedding. In *Proceedings of the fourth ACM conference on Recommender systems*, RecSys '10, 2010.
- [13] N. Kiukkonen, B. J., O. Dousse, D. Gatica-Perez, and L. J. Towards rich mobile phone datasets: Lausanne data collection campaign. In *Proc. ACM Int. Conf. on Pervasive Services (ICPS), Berlin.*, 2010.
- [14] A. Klemm, C. Lindemann, and O. Waldhorst. A special-purpose peer-to-peer file sharing system for mobile ad hoc networks. In *Vehicular Technology Conference, 2003. VTC 2003-Fall. 2003 IEEE 58th*, 2003.
- [15] G. Kortuem, J. Schneider, D. Preuitt, T. G. C. Thompson, S. Fickas, and Z. Segall. When peer-to-peer comes face-to-face: Collaborative peer-to-peer computing in mobile ad hoc networks. In *Proceedings of the First International Conference on Peer-to-Peer Computing*, P2P '01, 2001.
- [16] M. Kuhn, R. Wattenhofer, and S. Welten. Social audio features for advanced music retrieval interfaces. In *Proceedings of the international conference on Multimedia*, MM '10, 2010.
- [17] A. Madan and A. Pentland. Modeling social diffusion phenomena using reality mining. In *In AAAI Spring Symposium on Human Behavior Modeling*, 2009.
- [18] E. Miluzzo, N. D. Lane, K. Fodor, R. Peterson, H. Lu, M. Musolesi, S. B. Eisenman, X. Zheng, and A. T. Campbell. Sensing meets mobile social networks: the design, implementation and evaluation of the cenceme application. In *Proceedings of the 6th ACM conference on Embedded network sensor systems*, SenSys '08, 2008.
- [19] M. Nekovee. Worm epidemics in wireless ad hoc networks. *New Journal of Physics*, 9(6):189, 2007.
- [20] S. Trifunovic, B. Distl, D. Schatzmann, and F. Legendre. Wifi-opp: ad-hoc-less opportunistic networking. In *Proceedings of the 6th ACM workshop on Challenged networks*, CHANTS '11, 2011.
- [21] P. Wang, M. C. González, R. Menezes, and A.-L. Barabási. New generation of mobile phone viruses and corresponding countermeasures. *CoRR*, 2010.