Nomadic Sharing of Media: Proximity Delivery of Mass Content within P2P Social Networks

Balázs Bakos[1] and Lóránt Farkas[1]

[1] Nokia Research Center, P.O. Box 392
H-1461 Budapest, Hungary

Abstract. In our paper we discuss a novel P2P file sharing system optimized for smart phones, using proximity networking. Group management, peer discovery and proximity multicasting are discussed. The experiment with the proof-of-concept implementation on the Series S60 Symbian platform shows that content sharing in social proximity can happen in a cost- and resource-effective way.

1 Introduction

Today users can consume, create and share digital media content using their smart phones, equipped with peer-to-peer (P2P) clients. P2P content sharing applications have been very successful on the wired Internet. Napster [1], Gnutella [2] and BitTorrent [3] are among the most successful peer-to-peer (P2P) applications, measured to be generating most of the Internet traffic today [5].

Search and network maintenance [6-12], speed and reliability [4,5], free-riding [16] had been the main identified problems in P2P applications. Free-riding is becoming increasingly acute now when P2P clients are increasingly turning mobile [14]. This results in the decrease of the volume of available content, mobiles joining P2P overlay networks only for limited time periods, the overlay network being split into isolated islands, hence the overall performance and reachability of the network degrading.

The community dimension of P2P networks evolved from community-friendly networks such as the first generation Napster, to anonymous ones because of the increasing legal pressure, equivalent to the loss of the community dimension. Recently we experience the reviving of the community dimension with the introduction of darknets, small networks based on trust rather than anonymity. But trust brings the need for these systems to be closed.

In our paper, instead of improving the existing P2P content sharing applications or downscaling them to smart phones, we propose a novel P2P content sharing system optimized for smart phones and proximity links. Section 2 describes the concept. Section 3 focuses on social group management, enhanced peer discovery, multicasting and user interface. Section 4 concludes this paper and outlines future directions to be considered.

Bakos B. and Farkas L. (2006).
Nomadic Sharing of Media: Proximity Delivery of Mass Content within P2P Social Networks.
In *Proceedings of the 3rd International Workshop on Ubiquitous Computing*, pages 81-86
DOI: 10.5220/0002474500810086

2 Nomadic Sharing of Media

The tragedy of the commons comes from the assumption that people consume rather than share. In case of closed groups of mobiles we observe the opposite at least in two different cases. First, people eagerly consume and share, if they think the content is 'cool' or of interest to other peers in the same closed group. Second, people eagerly share content by others via forwarding, if this does not have as an effect a perceivable degradation of the performance of their device.

In nomadic sharing of media content is shared by a user from his phone to a group of people in the proximity. The process of content sharing consists of the following operations:

- People share content with group(s). This means that the content (file(s) or folder(s)) are selected and the group(s) are also selected.
- People meet friends. This is an opportunity for proximity scanning: the user's device scans and if devices belonging to the same group are found, content transfer is initiated.
- Content is transferred to friends' phones. Since multimedia content is typically large, of several megabytes, and proximity events can be very short, it is possible that in a session the content cannot be entirely transferred, which makes it necessary to apply a segmentation mechanism. Namely, content is segmented into standalone chunks of data.
- Friends' phones carry shared content. The shared content can be the entire content or chunks.
- Content (entire or chunks) reaches group members. If all chunks are available, then the content is reassembled and consumed by the group members.

Nomadic share does fulfill the three criteria of darknets from [17] in a particular fashion:

- the distributed content is not necessarily widely distributed, but expandable to such a scenario. Currently there is only a closed group of users participating in the sharing;
- users will not copy but share content which is interesting to themselves and think it would be interesting for others as well (push rather than pull)
- the channels through which users are connected to each other is not high-bandwidth in the strict sense: they are ad-hoc proximity links (of tens of kbyte/s).

3 Application Components

3.1 Social Group Management

One representation of our personal social network in the smart phone is the phonebook or contact list. This is a good starting point for building a more complete database. Since it is phone-centric in its structure (mobile phone number is in general a default field of the phonebook records) it can be efficiently used as a starting point for group creation using cellular short messages.

The contact list of smart phones provides the possibility to defining groups. However, people are in general reluctant towards defining groups because of various reasons: it is time consuming to define the groups, even if the groups are defined they are not static and they can change during the time. Therefore selecting the proper social group management method is a real dilemma for social interaction application designers.

Several solutions can be envisioned to cope with this situation. First, the application designer could rely on the fact that enthusiastic users will create the groups once they have content to share. Ways of doing this represent invitations sent via short messages or invitations sent via proximity. The first scenario is relevant in the typical case when the user sees or hears something and figures out enthusiastically how delighted a group of his friends would be if they had the same experience. Sms is an efficient way of quickly setting up a group and use it for content sharing. The second case is encountered in cases when there is a social event and friends not having seen each other since a long time want to share recent content such as weddings, birth of a child etc. In this case setting up a group can be achieved by proximity messaging.

Second, there might be situations in which users are not very eager to make this step by themselves. For instance, colleagues working in a team frequently meet at some specific location and they call each other typically during daytime. In these cases the smart phone can create groups in an automatic fashion or recommend associations between individuals and groups. This is achieved by logging information related to the user activity: the user's communication patterns; the set of locations visited by the user; the set of visible devices in the proximity/cellular surrounding of the user; the contacts of the user, in the form in which the user stores them in the phonebook or other possible repositories. Based on the collected data the phone can create clusters and assign people to clusters using standard methods of statistics, directly or through recommendations. Fig. 1 shows an exemplary event log containing 4 dimensions: type of event, time, duration, originator, caller/callee.

09/03/2005

Type	Time	Duration	Direction	Remote party
Phonecall	10:53:21	21 min 20 sec	Outgoing	Bob
SMS	11:20	N/A	Outgoing	Alice
Proximity	19:00	> 3 hours	N/A	John
Proximity	8:00	> 8 hours	N/A	Fred

Fig. 1. Event log in four dimensions.

In nomadic share we followed a simple approach: the private group and the social group concept were used. In the case of private groups, friends are selected and added from the phonebook, via proximity scanning or by typing. The remote parties are not explicitly aware of the existence of private groups: they exist as a formal group only in the device of the group creator.

In the case of social groups members are invited and they accept or reject the invitation. The remote parties are explicitly aware of their membership and can act upon, by deciding whether to accept the invitation or not, and by choosing at any time instant to leave the group or remain a member of the group.

3.2 Communications Engine – Peer Discovery and Multicasting in the Proximity

In a typical use case users have some data files that they want to send to other members of the closed social network, whom they occasionally meet. Unless the sender and the target members are in proximity at the moment, the data transfer can take place only when they meet next time. Such meetings are usually not planned but happen spontaneously. The file sharing service is much more user friendly if it does not require frequent user actions, so the user's phone needs to search for the target phones on its own, by periodically running a discovery process in the background.
The discovery process typically uses the same proximity technology (e.g. Bluetooth) as the data transfer. In order to increase the efficiency of the peer discovery, statistical data collected about the location and the time of occurrence of past proximity events, can be used. If the phone is equipped with a GPS, it can directly determine its location. Otherwise the location is derived from the cell information of the cellular network (list of cell ID: network ID doublets). To achieve sufficient accuracy, not only the serving BS, but also the set of all visible BS-s can be monitored by the mobile phone, the cell ID/net ID-s are collected and compared with stored cell ID/net ID lists.
It is possible to store in the phone the most frequently met cell ID's for each potential peer, combined also with timing data. As a consequence, smart phones can predict future proximity events and can concentrate their discovery effort to those locations and occasions in which the targets occur with high probability. Therefore the probability of the proximity communications increases. On the other hand battery power is saved this way, since the discovery process is executed more often only in cases when there are good chances to find other peers in the proximity.
Content transfer in the proximity is performed via Bluetooth piconets, widely available in smart phones. In piconets up to 8 nodes can participate in data transmission, but in a given time instant only two nodes communicate: a sender and a receiver. In order to maximize the number of simultaneous communications taking place we employed parallel piconets consisting of 2 nodes each. Fig. 2 shows such a situation, for the case of 8 nodes in each other's proximity.

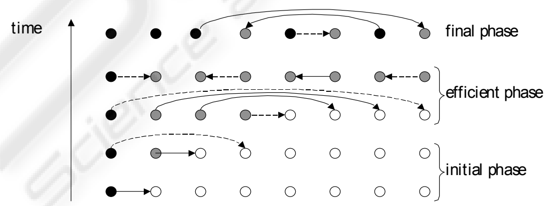

Fig. 2. Efficient multicasting using small parallel piconets.

In Fig. 2. black nodes are sources (primary or secondary), grey nodes are also sources but having received the content only partially. White nodes are recipients for the content.
If the content is split to 2 chunks and marking by continuous arrow the transmission of the first chunk and by dashed arrow the transmission of the second chunk, we can split the transmission in 3 phases: initial, intermediate and final.

In the initial phase only few nodes are sources. In the intermediate phase there are as many sources as recipients (or eventually one source less, the case of odd number of nodes). In this phase all nodes (or eventually all nodes but one) are engaged in data transmission, being either sources or destinations. In other terms piconets of 2 nodes are performing in parallel in the intermediate phase. Therefore, if the Bluetooth interference is negligible, then the throughput in this phase is n/2 times higher than in the one source – one destination case. Then in the final phase we have more potential sources than destinations and the aggregate throughput decreases accordingly.
Fig. 2 does not give an algorithmic scheduling plan, it only shows one possible optimal scheduling for 8 nodes. For scheduling plans the reader is referred to [13].
In the case when nodes are not in each other's proximity, the principle remains the same: the nodes store the data for themselves and also on behalf of other group members, but the forwarding takes place later, when the other group members are seen.
Segmentation makes it possible to cope with the situation of nodes walking out of coverage: the transmission can be resumed later without loss of data and the source can be also another node, different from the one who started the data transmission towards a particular recipient.

3.3 User Interface

Fig. 3 presents a set of snapshots from the application.

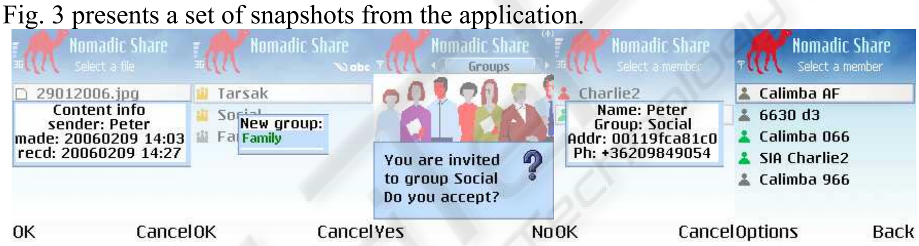

Fig. 3. Application snapshots: content info, group management, presence.

In the implementation phase the concept has been a standalone application with own user interface. Snapshots display various views of this user interface: event notification, group management, and presence are shown.
Notification is given to the user when content is received (in a persistent/non persistent way). The user can check details of the received content, such as date and time of creation, creator, sender.
Group management is available via creation/deletion of groups and adding/invitation of members. The invitation can be accepted or rejected.
Presence info can be visualized: 'green' members have been encountered in the proximity recently. Member details can also be visualized by clicking on the individual members: address, phone number and group membership are displayed.

4 Conclusions

In our work we introduced a novel P2P file sharing application for smart phones to deliver mass media content to a closed group of smart phones over multiple proximity connections. The concept has been implemented on the Nokia 6630 platform. In a user trial session enhancements were suggested to the concept related to the integration with existing Series S60 core applications such as sound and video recorder and contact list. Such an integration would allow users to share content right at creation time and thus minimize the number of steps required to initiate the content sharing.
Future work will focus on the merging of Internet-based P2P protocols and nomadic sharing of media to exploit proximity technologies and context awareness available in mobile devices. This direction may lead to the next generation of P2P protocols and applications suitable both for personal computers and mobile devices.

References

1. The Napster protocol, http://opennap.sourceforge.net/napster.txt
2. The Gnutella protocol specification v0.4, document revision 1.2, http://www9.limewire.com/developer/gnutella_protocol_0.4.pdf
3. Bram Cohen, Incentives Build Robustness in BitTorrent, http://www.bittorrent.com/bittorrentecon.pdf, May 2003
4. Free riding on Gnutella is revisited: The Bell Tolls?, http://csdl2.computer.org/comp/mags/ds/2005/06/o6001.pdf
5. Peer-to-peer file sharing – The impact of file sharing on service provider networks, white paper, December 2000, http://www.sandvine.com
6. G. Csúcs, J.K. Nurminen, B. Bakos, L. Farkas, Peer-to-peer Protocol Evaluation in Topologies Resembling Wireless Networks. An Experiment with Gnutella Query Engine, ICON'2003, Sydney, Australia
7. T. Hu, A. Sereviratne, General Clusters in Peer-to-Peer Networks, ICON'2003
8. K. Sripanidkulchai, B. Maggs, H. Zhang, Efficient Content Location Using Interest-Based Locality in Peer-to-Peer Systems, Infocom'2003
9. H. Meling, A. Montresor, O. Babaoglu, Peer-to-peer Document Sharing using the Ant Paradigm, NIK'2001, Tromso, Norway
10. E. Cohen, S. Shenker, Replication Strategies in Unstructured Peer-to-peer Networks, SigComm'2002, Pittsburgh, Pennsylvania
11. B. F. Cooper, H. Garcia-Molina, Modeling and Measuring Scalable Peer-to-peer Search Networks, SIGCOMM'2002
12. B. F. Cooper, H. Garcia-Molina, Ad hoc, Self-supervising Peer-to-peer Search Networks, Technical Report, Stanford University, 2003, www-db.stanford.edu/~cooperb/pubs
13. L. Farkas, B. Bakos, P. Spányi, A practical approach to multicasting in Bluetooth piconets, accepted for publication in Proc. WCNC'2006, Las Vegas, April, 2006
14. http://symella.aut.bme.hu
15. J. Hummel, U. Lechner, Social profiles of virtual communities. In: Proc. of HICSS'2002, Hawaii, 2002
16. J. Roettgers, Social networks: the future of p2p file sharing, http://freebitflows.t0.or.at/f/about/roettgers
17. P. Biddle et al., The darknet and the future of content distribution, ACM Workshop on digital rights management, Nov. 18, 2002, Washington DC, USA