Social-aware Content Sharing in Opportunistic Networks*

Marco Conti, Franca Delmastro, and Andrea Passarella
IIT Institute - CNR
via G. Moruzzi,1 - 56124 Pisa, Italy
email{firstname.lastname@iit.cnr.it}

Abstract

Opportunistic networks represent the emerging pervasive communication paradigm that supports intermittent connectivity scenarios. To develop reliable data exchange and communications both protocols and services can be enhanced exploiting context and content information, defining thus a novel data-centric communication paradigm. In this demo we present a social-aware content sharing service based on the exchange of a combination of the user personal information, interests, social relationships and information related to shared contents. In this way, the system is able to implement cooperative downloading mechanisms allowing users to share files even they will never be directly connected.

Keywords: opportunistic networks, context-awareness, p2p, content-sharing service

1 Introduction

The increasing use of mobile devices as tools to generate and share content among users according to the Web 2.0 model is one of the main motivations to exploit the mobile p2p computing paradigm in pervasive mobile environments. Recently, a lot of studies on p2p over self-organizing networks have been performed providing optimized solutions for mobile users [5, 3]. However, they mainly rely on the assumption that each pair of nodes wishing to communicate are connected through a stable path. Actually, a pervasive network is *intrinsically* disconnected mainly because of nodes' mobility. For this reason we can envision as a realistic scenario clouds of mobile devices, some of which being wirelessly connected in ad hoc mode, some others being connected to the legacy Internet, and others being temporarily disconnected from the network. Opportunistic networks represent the emerging pervasive communication paradigm that supports this scenario as an evolution of mobile ad hoc networks. In this case there is no assumption on nodes' connectivity, and nodes are not required to exchange information about the network topology. Routes are dynamically built while data flows towards the destination, exploiting all the available opportunities to successfully deliver the information. To this aim, opportunistic network communications are not supported by topology-based protocols, but they mainly exploit *context* and *content* information to define new *data-centric* communication paradigms. This information characterizes both data that are exchanged through the network (specified by the developed application), and users carrying mobile devices, defining their habits, interests and specific information that can be useful to optimize network services. In this way, we are moving from the standard concept of computer networks to the definition of networks of *people*. In this direction a novel definition of p2p over opportunistic networks has been proposed recently [4], highlighting the importance of disseminating content and context information in the network to improve the autonomous behaviour of mobile devices and better satisfy mobile users requirements. Within the Haggle[1] and SOCIALNETS[2] EU projects work has been done and it is currently in progress in this area. Specifically, in Haggle a data-centric architecture has been defined (see section 2.1) to support application-driven message forwarding in opportunistic networks, trying to reflect natural ways of interaction between humans. This demo presents an opportunistic contex- and social-aware file sharing service developed on top of Haggle's architecture in which using context and content information allows mobile users to share files even if they will never be directly connected.

2 Social-aware Opportunistic Content Sharing

In order to extend the User-generated content model in pervasive environments, Social-aware Opportunistic Content Sharing service highlights how using context

*This work was partially funded by the IST program of the European Commission under the HAGGLE (027918) and SOCIALNETS (217141) FET Projects.

[1] `http://www.haggleproject.org`
[2] `http://www.social-nets.eu`

978-1-4244-3938-6/09/$25.00 (c)2009 IEEE

information enhances the basic features of standard content-sharing applications in opportunistic networks, exploiting nodes' mobility to evolve the standard p2p connections among content producers and consumers. Focusing on this specific kind of service, our definition of context mainly refers to users' sharing interests (e.g., genre and type of files they want to share), habits, social contacts and mobility patterns. In fact, since users are mobile and carry devices with them, users' social behavior is a valid signal of users' mobility patterns, and it may be really useful to predict future contacts among users, improving forwarding decisions on the network. These assumptions are based on social network models [6] and small-world theories [7], in which users are grouped in communities, and nodes of the same community have strong social links between each other. Some nodes have also social links outside their "home" community, modelling social relationships with users of different groups. Small-world theories have shown that these "external" links act as shortcuts and enable communications across the network with a small number of hops (6 hops in the "classical" small-world models). This behavioral model is exploited by Context-aware Opportunistic FS service to spread context information in the users communities to improve both network protocols and applications features. Specifically, two different notions of context related to a single user are introduced in the presented application: the private context of a user, that consists of personal information[3] (e.g., name, address, job, communities he/she belongs to) that are not strictly related to the developed service (and can thus be used as general context to be shared among several local applications), and the public context, that is mainly defined by the service (e.g., sharing interests in terms of files genre, category, content attributes). All this information are collected in an object called Identity Table (IT), since it identifies both characteristics of the single user, and his/her interactions with the service. The user initializes it directly through a user-friendly GUI on the mobile device (a smartphone in this demo). Following the model presented in [2] for a context-aware forwarding algorithm in opportunistic networks, ITs are exchanged among 1-hop neighbors. Therefore, by exchanging ITs during the neighbor discovery procedure, each node knows the personal and public information of all its neighbors, and it adds them to its *current context*, obtaining thus a snapshot of the surroundings. However, since users usually moves from a community to another, it is also important to locally maintain the *context evolution* over time (i.e., information about context of past encountered users as the evolution of social contacts of the local user). In this way, when a user moves to a new community, the application can exploit the interests of past encountered nodes to identify additional files in which previous neighbors could be interested, and use part of its local resource to pre-fetch them to satisfy possible future requests. This procedure has been called *cooperative downloading* among opportunistic nodes. The possible choices between the files a user can decide to download in favour of previous neighbors mainly depends on the local resource availability and the data utility. Intuitively, the main idea is to select those files that maximize the utility for both the local user (e.g. data belonging to the same category she is interested in) and the users belonging to the different communities she also belongs to. We can thus define the following utility function:

$$U = u^{(l)} + \sum_i \omega_i u_i^{(c)} \tag{1}$$

where $u^{(l)}$ is the utility for the local user, $u_i^{(c)}$ is the utility for the i-th community the user is in contact with, and ω_i is a cooperation index that defines the willingness of the user to cooperate with the i-th community (i.e., to spend own resources to increase data availability for that community). Note that, by using cooperation indexes greater than 0, we can avoid the selfish users behaviour in which they tend to maximize only their own utility. In addition, the autonomous preventive downloads do not involve the files the local user is interested in, since it can directly request them. Another condition for the utility computation is represented by the current context of the local user. Since it is highly probable that users visiting the same community are in communication range, or can directly download the files they are interested in with few-hops connections, we can decide to assign a lighter weight to the current context of the local user, so that the system can privilege the other contexts. A detailed analysis of the *context-aware utility* and related parameters can be found in [1]. In this demo we implement a simplified definition of this function to highlight the main advantages of using context, as the application reliability in case of intermittent connectivity and the reduction of requests dissemination on the network.

As a practical example, let us consider two social communities of users as shown in Figure 1(a). In the example user C belongs to both communities, and in the two communities users share some common interests: on the left side users A and C are interested in mp3 and jpeg files, while B is interested in jpeg and avi files; on the right side both users D and E are interested in mp3 and avi. Assuming that all nodes in each community are able to comunicate with each other, after

[3]Managing private/personal information about users requires to deal with privacy issues (see [4] for details).

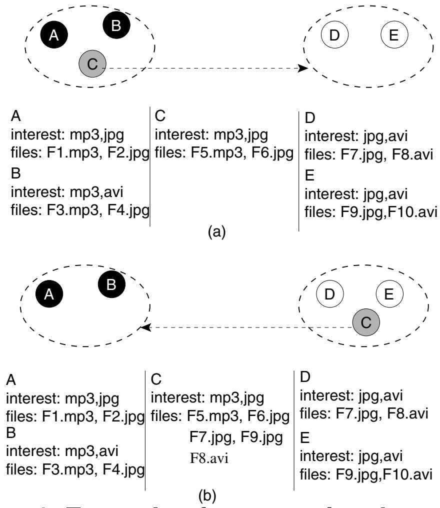

Figure 1. Example of cooperative download.

the neighbor discovery phase each of them knows the interests and the file list of the others, updating thus their current context. As C belongs to both communities, it is highly probable that she will move towards the right-hand-side community and can thus help distributing files to interested users. When she reaches the right-hand-side community she is able to select the files she is interested in and download them on her device. Moreover, looking at the previously stored context information, the application run by C is also able to identify, and possibly download, additional files in which the users of the left-hand-side community are probably interested (see Figure 1(b)).

2.1 Interaction with the Haggle architecture

Figure 2 shows the Haggle node architecture. It is completely layer-less above the data-link, and a set of managers is defined to implement all the features needed to communicate and deliver messages (e.g., security, protocols definition, forwarding algorithms, connectivity, resource and context management). Managers can make use of modules to implement specific functionality related to particular algorithms. All the managers interact with a central entity, the Haggle Kernel, which is a minimal event queue that coordinates actions and communications between managers, and waits for incoming data. All the interactions between managers are performed through a data-centric pub/sub paradigm. Managers publish DataObjects to whom other managers can subscribe. The kernel implements the pub/sub abstraction, and stores the DataObject in a DataStore. Relations between DataObjects can also be defined, enabling to represent rich context information. Haggle collects context information about both encountered users and devices through periodical beaconing. The analysis of beacons and extraction of context information is left to individual managers/modules that subscribe to beacons.

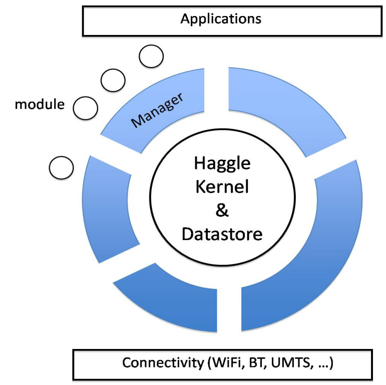

Figure 2. Haggle architecture.

The Social-aware Opportunistic Content Sharing service is partly developed on top of this architecture, and directly interacts with a dedicated module of the Resource Manager, which provides general context management and data dissemination features. Context and content information used by the service are defined through the interaction with the mobile user, and are then passed to the internal module to be spread on the network.

3 Demo highlights and requirements

During this demo attendees will be able to run the service on HTC Touch smartphones, downloading files from neighbors belonging to the same community and moving from a community to another. From the technical point of view, the demo will show the reliability of the system in case of intermittent connectivity and the implementation of cooperative downloading mechanisms based on the elaboration of context information.

References

[1] C. Boldrini, M. Conti, and A. Passarella. Contentplace: social-aware data dissemination in opportunistic networks. In *Proc. of ACM MSWiM '08*, pages 203–210, New York, NY, USA, 2008. ACM.

[2] C. Boldrini, M. Conti, and A. Passarella. Exploiting users social relations to forward data in opportunistic networks: The hibop solution. *Pervasive and Mobile Computing*, 4(5):633–657, 2008.

[3] M. Caesar, M. Castro, E. Nightingale, G. O'Shea, and A. Rowstron. Virtual ring routing: network routing inspired by DHTs. In *Proc. of ACM SIGCOMM*. ACM New York, NY, USA, 2006.

[4] M. Conti, F. Delmastro, and A. Passarella. *Context-Aware P2P Over Opportunistic Networks*. in Mobile Peer-to-Peer Computing for Next Generation Distributed Environments: Advancing Conceptual and Algorithmic Applications, ISBN: 978-1-60566-715-7, 2009.

[5] M. Conti, F. Delmastro, and G. Turi. *Peer-to-Peer Computing in Mobile Ad Hoc Networks*. in The Handbook of Mobile Middleware, Auerbach Publications, 2007.

[6] S. Wasserman and K. Faust. *Social network analysis: Methods and applications*. Cambridge University Press, 1994.

[7] D. Watts. *Small Worlds: The Dynamics of Networks Between Order and Randomness*. Princeton University Press, Princeton, 1999.