P2P support for Group-Communication Applications: a Cross-Layer Approach for MANET Environments

Marco Conti[1], Jon Crowcroft[2], Franca Delmastro[1], Andrea Passarella[1]
[1]CNR, IIT Institute
Via G. Moruzzi, 1 – 56124 Pisa, Italy
firstname.lastname@iit.cnr.it
[2]University of Cambridge, The Computer Laboratory
15 JJ Thomson Avenue – Cambridge CB3 0FD, UK
jon.crowcroft@cl.cam.ac.uk

***Abstract*— P2P systems are a natural way of supporting Group-Communication (GC) applications. Furthermore, their core characteristics (e.g., self-organisation, self-recovery, decentralisation) match particularly well MANET features. For these reasons we developed and tested a simple yet significant GC application on top of different P2P substrates. Our aim is to understand which type of P2P system is more suitable to support GC applications in MANET environments. Specifically, in this demo we highlight limitations of legacy P2P systems on MANETs, and advantages of an innovative solution based on cross-layer optimisations.**

I. Introduction

In the framework of the MobileMAN Project [2], we are investigating the viability of developing GC applications on real ad hoc networks. To this end, we developed the Whiteboard application (WB), which implements a distributed whiteboard among MANET users. WB usage is very intuitive (see Figure 1(a)). Users run a WB instance on their own devices, select a topic they want to join, and start drawing on the canvas. Drawings are distributed to all devices subscribed to that topic, and rendered on each canvas. GC applications fit well the underlying features of MANETs since they are distributed, self-organising, and decentralised in nature. In addition, they can represent a good incentive for users to adopt this technology in the daily use. Let us imagine a user entering a shopping center and running the WB on her mobile device. The user will get in touch with other people who share common interests exchanging opinions and suggestions. In addition, users are not supposed to be charged for such a service, since WB exploits the *free-of-charge* 802.11 technology.

The use of a P2P substrate can make the development of GC applications straightforward. However, providing an efficient P2P support in MANETs is a challenging task. Within the MobileMAN Project we have developed a P2P system optimised for MANETs (CrossROAD [1]), and we have compared it with traditional solutions (e.g., Pastry [3]). Our studies show that, unlike legacy P2P systems, CrossROAD is able to support GC applications very efficiently in very dynamic environments such as MANETs.

II. Legacy and Cross-layer P2P solutions

Originally, the WB application had been developed for wired networks where the P2P substrate was a structured overlay network based on a DHT (i.e., Pastry). As a design choice, WB exploits a subject-based multicast protocol to build groups (i.e., identify all user nodes interested in the same topic), and disseminate WB data to the group members. Specifically, each topic is associated to a multicast tree, identified by the topic ID. In the original implementation we used Scribe [4], an application-level multicast protocol built on top of Pastry. As any P2P system implementing a DHT, Pastry identifies each node with a logical ID. Scribe generates the topic IDs in the same address space used by Pastry to define the logical IDs. Then, it assigns the role of tree's root to the node whose logical ID is numerically closest to the topic ID. Other nodes have to register their interest with the root by sending subscription messages specifying as key value the topicID. Each subscription message is routed by the DHT to the root. It is discarded as soon as it reaches a *branching point*, i.e. the first node in the overlay path between the subscription originator and the root that is already in the tree. Finally, a node willing to send data on the tree has to specify the topicID so that they are snt to the root, which relays them to the group members.

In a wired network like Internet it has been shown that a P2P system based on Pastry and Scribe outperforms other similar solutions [4]. Therefore, we wanted to understand if such a system is suitable for multi-hop ad hoc networks, as well. In particular, we compare this solution (referred to as "legacy") with a *Cross-layer* P2P system (named CrossROAD [1]) designed during the MobileMAN Project [2]. The difference between the two systems is illustrated in Figure 1(b). CrossROAD presents to higher layers exactly the same interface and functionalities provided by Pastry, but exploits cross-layer interactions with a proactive routing protocol (OLSR) to build and maintain the DHT. These interactions are handled by the Network Status module (*NeSt*), which provides well-defined interfaces for cross-layer interactions throughout the protocol stack. Specifically, each node running CrossROAD piggybacks advertisements of its presence in the overlay into

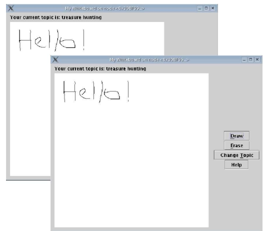

(a)

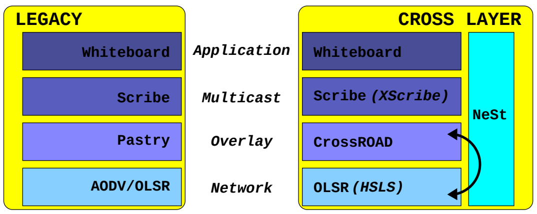

(b)

Fig. 1. The WB application (a) and the compared network stacks (b)

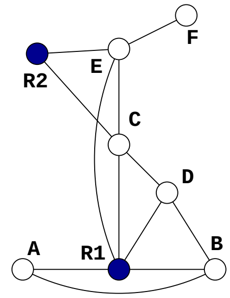

(a) The test-bed

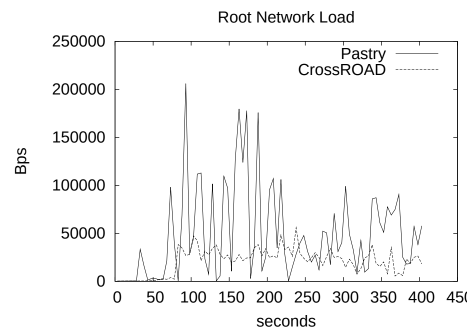

(b) Network Load

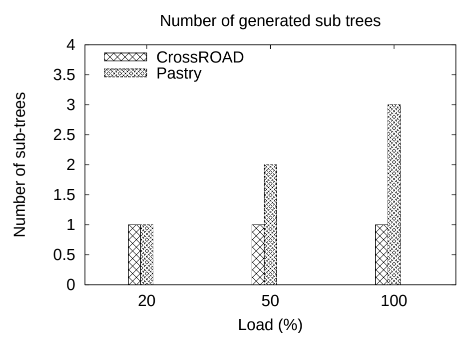

(c) Scribe sub-trees

Fig. 2. Experiments

routing messages periodically sent by OLSR. Thus, each node in the network becomes aware of the other peers in the overlay network. This mechanism drastically reduces the overhead required to build and maintain DHTs in legacy systems such as Pastry. Furthermore, CrossROAD is even more responsive than legacy systems, since it converges as quickly as the routing protocol does in case of mobility or network partitions. In a nutshell, the main idea behind CrossROAD is to build the overlay network exploiting the routing traffic. Thus, i) the system is completely self-organising, and ii) there is no need to generate separate management traffic to build the network and the overlay planes.

A. *Experimental Results*

To compare the legacy and the cross-layer solutions we have developed a small-scale test-bed consisting of 8 nodes (see Figures 2(a)). Six of them run the full stacks of Figure 1(b), while 2 more nodes act just as routers. It is worth pointing out that such small-scale test-beds are the most reasonable environments for realistic ad hoc network developement [5]. We run experiments in which users running WB draw lines on their canvas at a controlled rate. Specifically, a traffic load equal to $x\%$ means that the user draws, on average, $x/100$ lines per second. In all experiments, we noticed a drastic reduction of the network load when CrossROAD is used. To give an idea, Figure 2(b) plots the network load experienced by the root node with a 100% traffic load: CrossROAD cuts the average network load by 60%. For nodes other than root, this reduction is as high as 75%. Finally, Figure 2(c) shows that the reduction of the network load also improves the stability of the Scribe tree. Specifically, the figure shows how many sub-trees are generated during the experiment. Partitions of the tree (i.e., several sub-trees) results from the fact that some node does not receive correct information on the state of the overlay, especially in the bootstrap procedure, creating a new tree rooted at itself. Of course nodes in distinct sub-trees are no more able to communicate considerably increasing the packet loss. We have also measured an higher number of re-subscription when Pastry is used. A node generates a re-subscription when its parent in the tree is not found in the overlay network anymore. Since in Pastry experiments the physical network is overloaded, the overlay network (and the multicast tree) is not very stable. On the other hand, the light load imposed by CrossROAD makes the overlay network and the Scribe tree more stable.

B. *Future Work*

The results we got from our experiments show that using a lightweight, cross-layer, P2P substrate improves the performance of GC applications. The next step we are investigating is how such cross-layer interactions can be exploited to optimise the multicast data distribution. Specifically, Scribe needs to maintain a fairly complex tree structure in order to disseminate the data. In a very dynamic environment such as a MANET, this can be too costly. We are currently designing a structure-less, cross-layer, overlay multicast protocol that be able to overcome these limitations.

III. Demo highlights

During the demo, attendees will be able to run the Whiteboard on laptops, and communicate with each other through it. Making attendees use lively our WB implementation on mobile devices, this demo will show how such simple, “Plug&Play” applications can be enjoyable and valuable to users. It will thus suggest a class of novel applications that can be viably developed on cheap, off-the-shelf 802.11 devices. From a technical standpoint, the demo will compare the two networking approaches we have considered in our experiments (see Figure 1(b)). On one hand WB will run on the legacy stack, exploiting Pastry as P2P support. On the other hand WB will run on the cross-layer P2P solution designed in the framework of the MobileMAN project (CrossROAD). The two instances of WB will run concurrently, and plots will be dynamically generated, showing in a very intuitive way the network overhead of the two solutions. The demo will highlight that the cross-layer solution is able to provide exactly the same service while i) extremely reducing the networking overhead, and ii) guaranteeing the reconfiguration of the system in case of link failures or nodes disconnections with no influence on the application.

References

[1] F. Delmastro, “From Pastry to CrossROAD: Cross-layer Ring Overlay for Ad hoc networks”, in Proc. of Workshop of Mobile Peer-to-Peer 2005, in conjuction with the PerCom 2005 conference, Kauai Island, Hawaii, Mar. 2005.

[2] “Mobile Metropolitan Ad hoc Network (MobileMAN)”, IST-2001-18113 Project, funded by the EC FET-IST Programme, `http://cnd.iit.cnr.it/mobileMAN/`

[3] A. Rowstron and P. Druschel, “Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems”, Middleware 2001, Germany, November 2001.

[4] M. Castro, M. B. Jones, A-M. Kermarrec, A. Rowstron, M. Theimer, H. Wang and A. Wolman, “An Evaluation of Scalable Application-level Multicast Built Using Peer-to-peer overlays”, Infocom 2003, San Francisco, CA, April, 2003.

[5] P. Gunningberg and H. Lundgren and E. Nordström and C. Tschudin, “Lessons from Experimental MANET Research”, Ad Hoc Networks Journal, (Special Issue on “Ad Hoc Networking for Pervasive Systems”), Vol. 3, Number 2, March 2005.