C-VeT, The UCLA Campus Vehicular Testbed: Preliminary Measurements and Results

Eugenio Giordano*†,Abhishek Ghosh*, Giovanni Pau*, Mario Gerla*
*Computer Science Department - University of California, Los Angeles, CA 90095, US
e-mail: {giordano|aby|gpau|gerla}@cs.ucla.edu
†DEIS, WiLab - Università di Bologna, Bologna, 40136, Italia

***Abstract*—Vehicular ad-hoc networks (VANET) are becoming more and more interesting nowadays. In fact the technology advances allow high speed communication among vehicles, prompting a new set of possible applications such as Peer to Peer (P2P) content sharing, on-line games etc. It is then critical to understand the interplay of several unrelated technologies such as GPS, Wi-Fi, P2P etc. With this poster we present our experience in the creation of the UCLA Campus Vehicular Testbed (C-VeT) and some preliminary results on popular internet applications such as BitTorrent and Video Streaming.**

I. Introduction

The UCLA Campus Vehicular Testbed (C-Vet) is an innovative project that aims to provide a platform supporting car-to-car and car-to-infrastructure experiments with various data traffic types and mobility patterns. The unique virtualized router implementation allows us to simultaneously test multiple protocols and applications in the identical conditions, sharing common resources. In a two to three year span C-VeT will evolve to order of 50 cars each equipped with a mini PC, two Wi-Fi radio interfaces, a digital long range wireless modem, a GPS receiver and several sensors for special purpose applications. Each vehicle router can be remotely monitored and, if necessary reconfigured. The router will run multiple virtual machines allowing, for example, the simultaneous testing of several routing schemes. C-VeT will include Access Points (APs) providing access to/from the Internet. C-Vet is now under development; not all the planned features are currently available. This poster describes the current state of development, reports on the protocols installed so far and offers a preview of the experiments enabled by C-Vet.

II. Testbed Description

Currently the testbed consists of 6 nodes, each of them is a Dell Latitude D410 laptop. Every node is equipped with a GPS receiver, a 802.11 b/g MIMO wireless card, and a webcam. Each laptop runs Windows XP Professional. The multi-hop connectivity is achieved using the OLSR Daemon application [1]. Only one of the nodes is equipped with an EVDO wireless card that connects the whole ad-hoc network to the internet acting as a gateway. On each node runs an application that collects the node's position, processes the routing table extracting the one-hop neighbors and periodically forwards all the data to the control station. The control station is a fixed machine connected to the internet that receives all the positions of the mobile nodes and their routing information and at every instant displays the state of the network. Additionally the video coming from the webcam is saved by each node and optionally it can be streamed to any of the mobile nodes. With this configuration we performed several tests around the UCLA campus. During these tests we focused on the time a node needs to reconnect after it gets disconnected from the network. Hence we obtained a tuning for the OLSR Daemon that best performs in this vehicular environment. In table I we show such configuration.

TABLE I
OLSR parameters

Parameter	Value
Hello Interval	1.0 sec
Hello Hold	5.0 sec
TC interval	1.0 sec
TC Hold	5.0 sec

III. Experimental Section

In this section we present the results of our Peer to Peer and Video Streaming experiments.

A. Peer to Peer

To investigate the performance of P2P applications in a mobile, vehicular environment, we carried out our first set of experiments around the Engineer IV building at UCLA. We used the free BitTorrent [6] client Azureus [2]. Azureus is an open source java application that includes a distributed tracker, and essential feature to distribute files among nodes without internet tracker support. This is crucial in our ad-hoc network, since the access to the internet is through a shared gateway of limited data rate. In our first setup we uploaded a 50 MB file on 2 source nodes placed at the two extreme corners of the building. Using Azereus we attempted to download this file to 2 static nodes placed at the remaining two corners of the building. and to 2 mobile nodes (laptops placed in cars) driving clockwise around the building. The second experiment uses a configuration similar to the first. A 100 MB file is uploaded on just one source placed at one

This work was supported by Microsoft Research through its PhD Scholarship Programme.

corner of the building. Using Azureus, this file was uploaded to 3 static nodes placed at the remaining three corners of the building. At the same time the same file was uploaded to a single mobile node which again took clockwise revolutions around the building. In figure1 we show the average download speed we achieved in the 2 different experiments. In the first scenario the average download speed is much higher for the fixed nodes. In fact the fixed nodes are always one hop away from both the sources, instead the mobile nodes are always one hop away from one source and 2 hops away from the other. As shown in figure 1 the average download speed is much higher for the moving node in the second case. This happens because one of the fixed nodes is 2 hops away from the source and the other two fixed nodes act as relays hence decreasing their download speed. On the other hand, the mobile node is 1 hop away from the source for half of the time and most likely never acts as a relay, because for the fixed nodes a path to the source in always present (through the other fixed nodes) and will be used to route the packets.

B. Video Streaming

We used 4 static nodes placed at four extreme corners of Engineering IV building. At the same time, using a webcam mounted on a car and VLC we streamed the live video to exactly one of the static nodes. The webcam generated the video stream at a resolution of 176X144 pixels at 15 frames per second. Thus the stream was generating an average of 128 Kbps (since the codec used was DIV3 the bitrate was not constant due to dynamic compression). The video was streamed using UDP so the lost frames were not retransmitted. The VLC server was set with a cache of 200ms. On the receiving node we were both saving and displaying the video. In the real time video the missing frames were much more than 10% but since we were saving the raw data received from the source we were able to reconstruct and re-encode the video received. In table II we show the percentage loss for the video after the reconstruction. As shown in table II the percentage of loss for both frames and blocks is approximately 10%. Such a loss still grants the possibility of actually displaying the video. In real time then the video was not displayable due to fact that the frames were not always delivered in order (due to UDP), this issue could be solved using a more reliable transport protocol or implementing a longer buffer for the frames. Another important detail is the time the frame losses occurred. In fact they occurred when the mobile node was switching from one relay to another. This means that the refresh of the route is not fast enough to be transparent for the video stream.

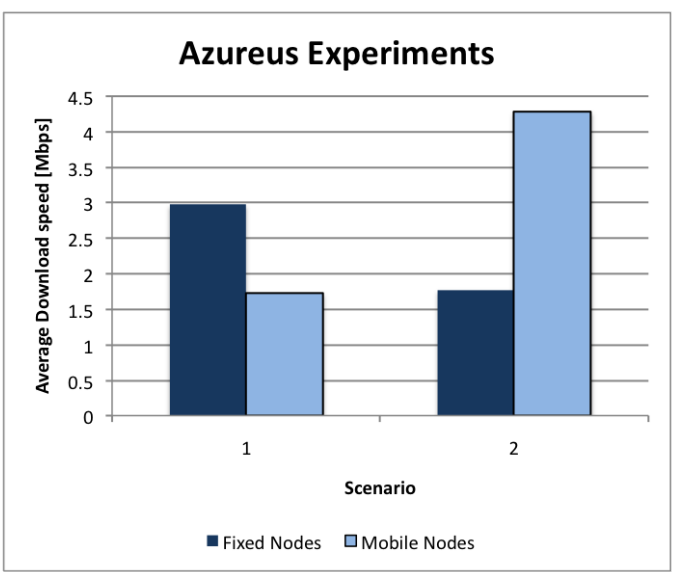

Fig. 1. Average download speed in both scenarios

TABLE II
VIDEO STREAM LOSSES

Parameter	Value
Frame Loss	9.02%
Block Loss	13.41%

REFERENCES

[1] RFC3626, T. Clausen and P. Jacquet, *Optimized Link State Routing Protocol (OLSR)*, 2003, www.olsr.org
[2] www.azureus.com
[3] www.vlc.org
[4] http://dast.nlanr.net/Projects/Iperf/
[5] A Nandan, S Das, G Pau, M Gerla, MY Sanadidi - *Co-operative downloading in vehicular ad-hoc wireless networks* - Wireless On-demand Network Systems and Services, 2005.
[6] BitTorrent official protocol specification www.bittorrent.com/protocol.html
[7] www.sopcast.org

http://www.vehicularlab.org/

C-VeT, The UCLA Campus Vehicular Testbed: Preliminary Measurements and Results

http://www.netlab.cs.ucla.edu/

Eugenio Giordano, Abhishek Ghosh, Giovanni Pau, Mario Gerla

{giordano, aby, gpau, gerla}@cs.ucla.edu

The UCLA Campus Vehicular Testbed C-VeT

Provides:

- A platform to support car-to-car experiments in various traffic conditions and mobility patterns
- A shared virtualized environment to test new protocols and applications
- Remote access to C-VeT through web interface
- Extendible to 1000's of vehicles through WHYNET emulator
- potential integration in the GENI infrastructure

Allows:

- Collection of mobility traces and network statistics
- Experiments on a real vehicular network

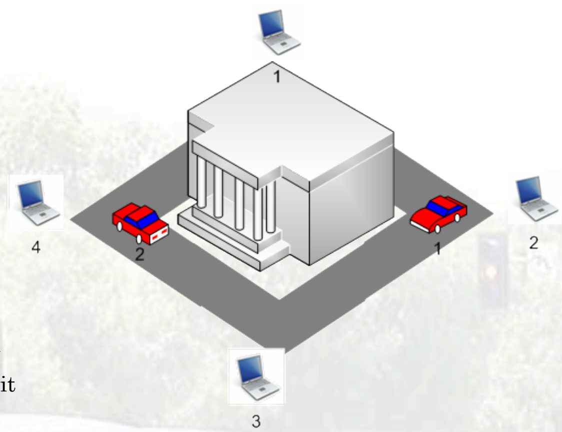

Preliminary Experiments:

- Fixed nodes placed at 4 corners of the Eng IV building at UCLA
- Each fixed node is either 1 hop or 2 hops away from another fixed node
- 2 cars (mobile nodes) mounted with laptops revolve around the building
- Each node (fixed or mobile) uses a 802.11b/g MIMO belkin wireless card
- One of the mobile nodes is mounted with a Microsoft Webcam to transmit video
- Each node runs OLSR Daemon to perform the information exchange and to update the routing tables

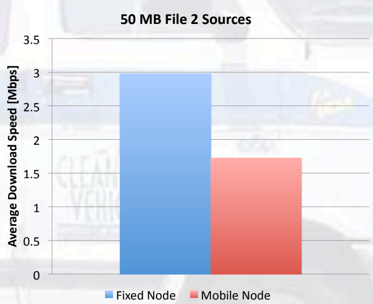

- All fixed nodes and all mobile nodes active
- Fixed nodes 1 and 3 have the 50 MB source file
- Fixed nodes 2 & 4 and mobile nodes 1 & 2 attempt to download this file using Azureus which is a BitTorrent open source Java client with a distributed tracker feature
- Fixed nodes perform better than mobile nodes
- Reason attributed to the fact that fixed nodes are always 1 hop away from both sources

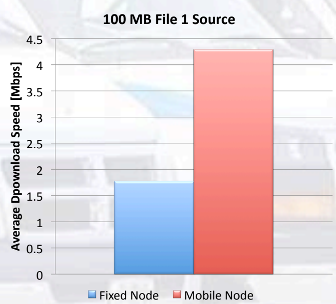

- All fixed nodes active. However mobile node 2 is inactive
- Fixed node 1 has a 100 MB source file
- Fixed nodes 2, 3 & 4 and mobile node 1 attempt to download this file using Azureus
- Now mobile node performs better than fixed nodes
- Reason attributed to the fact that mobile node is 1 hop from the source for half of the time and does not act as a relay. Additionally fixed nodes 2 & 4 always act as relays for the fixed node 3

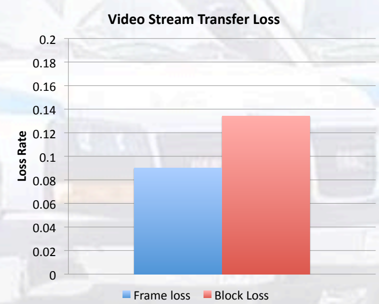

- All fixed nodes active. However mobile node 2 is inactive
- Mobile node 1 transmits live video (Driver's View) using Video Lan Client (VLC) to fixed node 3
- The percentage of loss for both frames and blocks is approximately 10%
- However Such a loss still grants the possibility of actually displaying the video within 2 hops

ISMB Istituto Superiore Mario Boella
http://www.ismb.it/
http://www.5t.torino.it/
SCALABLE NETWORK TECHNOLOGIES
http://www.scalable-networks.com
http://www.esteri.it/
Microsoft Research
http://research.microsoft.com.com