An Experimental Study on Scalable Video Streaming over Hybrid Cellular and Ad Hoc Networks

Ngoc Minh Do[1], Cheng-Hsin Hsu[2], Xin Huang[3], and Nalini Venkatasubramanian[1]
[1]Dept. of Information and Computer Science, University of California, Irvine, USA
[2]Dept. of Computer Science, National Tsing Hua University, Hsin Chu, Taiwan
[3]Deutsche Telekom Inc. R&D Labs, 5050 El Camino Real #221, Los Altos, USA

ABSTRACT

We consider scalable video streaming from a base station over a hybrid cellular and ad hoc network to a large number of mobile devices. While such a system has been recently studied, existing efforts resort to simulations when evaluating their solutions, because there is no public-domain software for setting up a complete hybrid network testbed. In this work, we design and implement a testbed for scalable video streaming over a hybrid network. Our testbed is built on top of a Linux server and multiple Android smartphones. We demonstrate how to use our testbed to evaluate the scheduling algorithms proposed in the literature. We also present a new scheduling algorithm which runs in real-time yet outperforms previous algorithms. We firmly believe that the testbed implementation and our experiences will stimulate future studies in this area.

Categories and Subject Descriptors: C.2.1 [Computer-Communication Networks]: Network Architecture and Design

General Terms: Design, Experimentation

1. INTRODUCTION

Streaming videos over cellular networks suffers from low video quality due to network capacity restrictions. For example, a measurement study [2] shows that an HSDPA cell can only support 6 video streaming users. Modern mobile devices, however, come with multiple network accesses, e.g., cellular and WiFi networks. While cellular networks provide *always-on* connectivity, WiFi interfaces allow mobile devices to form ad hoc networks for higher data rates than most cellular networks. We refer to the combination of these two networks as a *hybrid network* throughout this paper. In this paper, we consider video streaming over a hybrid network to many mobile devices.

Because wireless networks are vulnerable to bandwidth fluctuation, scalable video coding approaches, such as H.264/SVC, have been proposed to cope with this issue. While hybrid network and scalable video coding offer a great opportunity to efficiently deliver video streams to a large number of mobile devices, very few prior studies jointly consider these two techniques. Hua et al. [3] is the first work that considers scalable video streaming over hybrid networks, in which they assume the cellular base stations support multicast and solve the problem of determining the transmission rages of individual layers. In our previous work [1], we study the scheduling problem of scalable video streaming over a hybrid network in order to maximize the overall video quality. We show that the scheduling problem is NP-Hard, and develop an optimal technique, OPT, that uses a Mixed Integer Linear Program (MILP) based algorithm. We also propose a heuristic algorithm, HEU. Neither of the two studies [1,3] validate their solutions via experiments in real testbeds. They had to resort to simulations probably because there is no public-domain software for setting up such an evaluation.

In this work, we design and implement a testbed including a Linux-based server that streams scalable videos to multiple Android smartphones over a hybrid network. We also propose a simpler, more efficient scheduling algorithm called Practical Scheduling Algorithm (PSA). We implement three algorithms: OPT, HEU, and PSA, in the testbed, and conduct experiments to evaluate their performance. Our experimental results from T-Mobile 3G network indicate that: (i) streaming over hybrid networks significantly outperforms cellular-only networks, up to 7.8 dB quality improvement in Peak Signal-to-Noise Ratio (PSNR), and (ii) the proposed PSA runs much faster than HEU presented in [1], yet achieves better video quality than HEU. The testbed implementation and our experiences will stimulate future studies in the area of scalable video streaming over hybrid networks.

2. SCALABLE VIDEO STREAMING OVER A HYBRID NETWORK

We follow the system architecture presented in Do et al. [1]. The hybrid network consists of a streaming server that sends video data to mobile devices over a cellular network. Mobile devices that receive data directly from the base stations are called gateways. Gateways forward the received data to neighboring mobile devices over ad hoc networks. Mobile devices register with the server to receive a video they are interested in. A video is split into video segments with a D-sec duration; a typical value of D is 2. Each segment is encoded into L layers. We refer to a layer of a segment as a *unit*, which is the basic unit for data transfer.

The streaming server runs a scheduling algorithm every $W_r \cdot D$ sec to select units of W segments ($W_r < W$) for streaming during $W_r \cdot D$ sec over individual network links. Before each time the streaming server runs the algorithm to compute a new schedule, all mobile devices report information including their location and unit availability to the server over the cellular network. The streaming server records the data rates at different locations, and builds up a history database [4]. This database is used by scheduling algorithms to look up for the *expected* data rate of links in various locations. The output of the algorithm is a schedule that determines

Middleware Posters'2011, December 12th, 2011, Lisbon, Portugal.

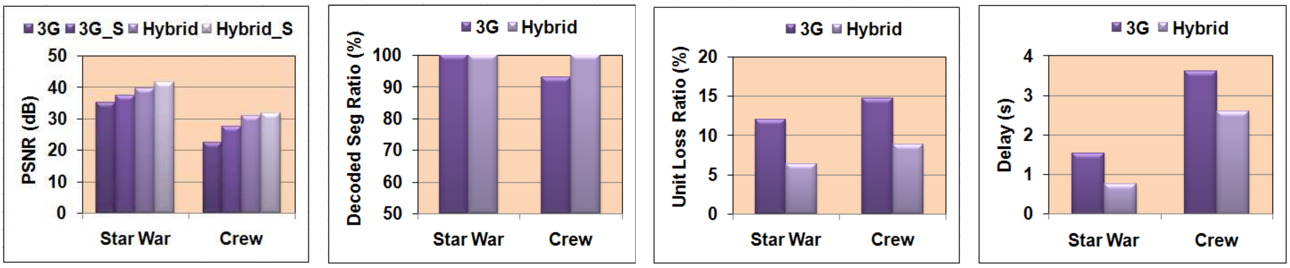

(a) (b) (c) (d)

Figure 1: (a) Quality, (b) decodable segment ratio, (c) unit loss ratio, and (d) delay.

which mobile devices will transmit which units to which mobile devices. The computed schedule is distributed to all mobile devices over the cellular network. The streaming server follows the schedule to send units to gateways. Whenever a gateway or a mobile device receives a unit, it follows the current schedule and forwards that unit to its neighbors over ad hoc links.

We propose a Practical Scheduling Algorithm (PSA) that computes schedules as follows. It sorts units in the scheduling window into a *candidate* list in a descending order of their importance (by layer, by segment, and then by video). PSA iterates through the candidate list and schedules unit transfers. For each unit u, PSA gives ad hoc networks higher priority and transmits u from a mobile device to neighboring mobile devices that do not hold u via a breadth-first tree over ad hoc networks. After considering ad hoc networks, PSA transmits u over cellular network to those mobile devices that haven't showed up on any breadth-first tree. This is done by selecting mobile devices with the highest cellular data rates as gateways. Gateways then forward the units to neighbors via breadth-first trees over ad hoc networks. After scheduling u, PSA moves on to the next unit in the candidate list, until it reaches the end of the list.

3. TESTBED AND SAMPLE EXPERIMENTAL RESULTS

We implemented a Linux streaming server on a 2 GHz Pentium IV machine with 1 GB RAM. The server connects to the Internet via a Fast Ethernet link. We also implemented a streaming client on Android, and deployed it on 5 Android smartphones. This testbed is located at University of California Irvine campus. The smartphones connect to the streaming server over T-Mobile 3G network, and are equipped with GPS readers. A history database containing data rates at locations on our campus was built up at the streaming server. We used H.264/SVC traces available at ASU (`http://trace.eas.asu.edu/videotraces2/cgs`). The server followed the traces to generate UDP video packets. Videos used in our experiments included: *Crew* and *Star War*. We streamed one video per experiment, although our testbed is general and can handle multiple videos concurrently. Each video was split into 2-sec segments, each segment is encoded into 4 layers. We used $W = 10$ and $W_r = W/2$. The initial buffering time is 5 sec. We consider the following metrics: (i) video quality in PSNR, (ii) unit loss ratio, (iii) ratio of missing units, (iv) ratio of undecodable segments due to missing base layer, and (v) unit delay.

We first report the performance of the proposed PSA algorithm. Fig. 1(a) shows the average perceived PSNR. In this figure, 3G, and 3G_S refer to PSNR values received and scheduled over the cellular-only network, respectively, and Hybrid and Hybrid_S are PSNR values received and scheduled over the hybrid network, respectively. With both videos, the hybrid network outperforms the cellular network. The gap between Hybrid and 3G when streaming *Crew* (8.6 dB) is higher than that of *Star War* (4.5 dB) because *Crew* has a higher bit rate. This shows that ad hoc networks significantly improve the video streaming quality, compared to cellular-only networks. Fig. 1(b) partially explains why the cellular network suffered from such a low PSNR when streaming *Crew*: there were only 93% of decodable segments. Fig. 1(c) depicts lost unit ratios in both networks when streaming videos. The number of lost units in the hybrid network is lower since WiFi links have much higher bandwidth than cellular links. This observation also explains lower delays in the hybrid network shown in Fig. 1(d).

Table 1: OPT vs HEU vs PSA Algorithms.

Metrics	OPT	HEU	PSA
PSNR (dB)	16.8	27.5	30.9
Deadline missing unit ratio	55.7	22.4	8.3
Decodable segment ratio	60.6	98.9	100

We also compared PSA against OPT and HEU algorithms. In our experiments, OPT ran on average 1.22 sec to generate a schedule. HEU took shorter time than OPT, but can still take up to 0.50 sec. In contrast, PSA runs in real-time. Table. 1 shows the PSNR, the ratio of missing units, and the ratio of decodable segments when streaming *Crew* over the hybrid network. OPT attempted to optimize the received PSNR at all receivers, but its PNSR is the lowest. This is because OPT took a much longer running time to come up with solutions than the others, and thus cannot keep up with transmissions. This table also shows that PSA results in video quality higher than HEU although PSA has a much shorter running time.

4. CONCLUSION

We considered scalable video streaming over a cellular network with the supports from WiFi ad hoc networks, which run on free spectrums, can be set up everywhere, and with high bandwidth. We designed and implemented a complete testbed for such systems, which can be leveraged by the research community for validating their algorithms in *live* cellular and WiFi networks. We also demonstrated how to use this testbed to evaluate a scheduling algorithm proposed in this work.

5. REFERENCES

[1] N. Do, C. Hsu, J. Jatinder, and N. Venkatasubramanian. Massive live video distribution using hybrid cellular and ad hoc networks. In *Proc. of IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM'11)*, pages 1–9, Lucica, Italy, June 2011.

[2] F. Hartung, U. Horn, J. Huschke, M. Kampmann, T. Lohmar, and M. Lundevall. Delivery of broadcast services in 3G networks. *IEEE Transactions on Broadcasting*, 53(1):188–199, March 2007.

[3] S. Hua, Y. Guo, Y. Liu, H. Liu, and S. Panwar. Scalable video multicast in hybrid 3G/ad-hoc networks. *IEEE Transactions on Multimedia*, 13(2):402–413, April 2011.

[4] A. Nicholson and B. Noble. BreadCrumbs: Forecasting mobile connectivity. In *Proc. of ACM International Conference on Mobile Computing and Networking (MobiCom'08)*, pages 46–57, San Fransico, CA, September 2008.