Improving Mobile Video Streaming with Mobility Prediction and Prefetching in Integrated Cellular-WiFi Networks

Vasilios A. Siris, Maria Anagnostopoulou, and Dimitris Dimopoulos

Mobile Multimedia Laboratory, Department of Informatics
Athens University of Economics and Business, Greece

Abstract. We present and evaluate a procedure that utilizes mobility and throughput prediction to prefetch video streaming data in integrated cellular and WiFi networks. The effective integration of such heterogeneous wireless technologies will be significant for supporting high performance and energy efficient video streaming in ubiquitous networking environments. Our evaluation is based on trace-driven simulation considering empirical measurements and shows how various system parameters influence the performance, in terms of the number of paused video frames and the energy consumption. The proposed approach has been implemented in an Android-based prototype.

1 Introduction

A major trend in mobile networks over the last few years is the exponential increase of powerful mobile devices, such as smartphones and tablets, with multiple heterogeneous wireless interfaces that include 3G/4G/LTE and WiFi. The proliferation of such devices has resulted in a skyrocketing growth of mobile traffic, which in 2012 grew 70%, becoming nearly 12-times the global Internet traffic in 2000, and is expected to grow 13-fold from 2012 until 2017[1]. Moreover, mobile video traffic was 51% of the total traffic by the end of 2012 and is expected to become two-thirds of the world's mobile data traffic by 2017. The increase of video traffic will further intensify the strain on cellular networks, hence reliable and efficient support for video traffic in future networks will be paramount.

The efficient, in terms of both network resource utilization and energy consumption, support for video streaming in future mobile environments with ubiquitous access will require integration of heterogeneous wireless technologies with complementary characteristics; this includes cellular networks with wide-area coverage and WiFi hotspots with high throughput and energy efficient data transfer. Moreover, the industry has already verified the significance of mobile data offloading: globally, 33% of total mobile data traffic was offloaded onto WiFi networks or femtocells in 2012[1].

Our goal is to quantify the improvements for mobile video streaming that can be achieved by exploiting mobility and throughput prediction to prefetch

[1] Source: Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2012-2017, Feb. 6, 2013

video data in local storage of WiFi hotspots, efficiently utilizing the resources of integrated cellular and WiFi networks. Although we consider mobile video data offloading to WiFi hotspots, our results and conclusions are potentially applicable to mobile video offloading to femto or small cell networks, where the backhaul throughput is smaller than the radio interface throughput. The work in this paper is different from our previous work in [13, 11] that considers mobile data offloading for delay tolerant traffic, which requires transferring a file within a time threshold, and delay sensitive traffic, which requires minimizing the file transfer time; unlike these traffic types, video streaming requires a continuous transfer of video data to avoid degradation of user QoE (Quality of Experience), which thus requires a totally different prefetching procedure and evaluation.

2 Related work

Prior work has demonstrated bandwidth predictability for cellular networks [15] and WiFi [8]. Bandwidth prediction for improving video streaming is investigated in [3], and for client-side pre-buffering to improve video streaming in [10]. Both [3, 10] focus on cellular networks, while we consider integrated cellular-WiFi networks. Moreover, our goal is not to develop a new system for mobility and bandwidth prediction, but to evaluate mobility and throughput prediction to prefetch data in order to improve mobile video streaming.

Exploiting delay tolerance to increase mobile data offloading to WiFi is investigated in [1, 6]. The work in [5] applies a user utility model for offloading traffic to WiFi, while [9, 14] investigates usage of multiple heterogeneous wireless interfaces. Our work differs in that we focus on video streaming and exploit prefetching video data in local caches of WiFi hotpots along a vehicle's route.

The feasibility of using prediction together with prefetching is investigated in [2], which does not propose or evaluate specific prefetching algorithms. Prefetching for improving video file delivery in cellular femtocell networks is investigated in [4], and to reduce the peak load of mobile networks by offloading traffic to WiFi hotspots in [7]. Our work differs in that we consider prefetching in WiFi hotpspots along a vehicle's route to improve video streaming.

3 Mobility & throughput prediction for prefetching

Mobility prediction can provide knowledge of how many WiFi hotspots a node (vehicle) will encounter, when they will be encountered, and for how long the node will be in each hotspot's range. In addition to this mobility information, we assume that information on the estimated throughput in the WiFi hotspots and the cellular network, at different positions along the vehicles's route, is also available; for the former, the information includes both the throughput for transferring data from a remote location, e.g., through an ADSL backhaul, and the throughput for transferring data from a local cache. Prefetching can provide gains when the throughput of transferring data from a local cache in the WiFi hotspot is higher than the throughput for transferring data from its original remote server location. This occurs when the backhaul link connecting the hotspot to the Internet has low capacity (e.g., in the case of an ADSL backhaul) or when it is congested; this is likely to become more common as the use of the IEEE 802.11n and 802.11ac standards increases.

The procedure to exploit mobility and throughput prediction for prefetching, Algorithm 1, defines the mobile's actions when it exits a WiFi hotspot, hence has only mobile access (Line 6), and when it enters a WiFi hotspot (Line 9). Mobility and throughput prediction allows the mobile to determine when it will encounter the next WiFi hotspot that has higher throughput than the cellular network's throughput. From the time to reach the next hotspot and the average video buffer playout rate, the mobile can estimate the position that the video stream is expected to reach (offset) when it arrives at the next WiFi hotspot (Line 7). The mobile instructs a local cache in the WiFi hotpot to cache the video stream starting from the estimated offset (Line 8).

Algorithm 1 Using mobility and throughput prediction to prefetch video data

1: **Variables:**
2: R: average video buffer playout rate
3: $T_{\text{next WiFi}}$: average time until node enters range of next WiFi
4: *Offset*: estimated position in video stream when node enters next WiFi hotspot
5: **Algorithm:**
6: **if** node exits WiFi hotspot **then**
7:     $\mathit{Offset} \leftarrow R \cdot T_{\text{next WiFi}}$
8:     Start caching video stream in next WiFi starting from *Offset*
9: **else if** node enters WiFi hotspot **then**
10:     Transfer video data that has not been received up to *Offset* from original location
11:     Transfer video data from local cache
12:     Use remaining time in WiFi hotspot to transfer video data from original location
13: **end if**

When the node enters a WiFi hotspot, it might be missing some portion of the video stream up to the offset from which data was cached in the hotspot; this can occur if, due to time variations, the node reaches the WiFi hotspot earlier than the time it had initially estimated. In this case, the missing data needs to be transferred from the video's original remote location (Line 10), through the hotspot's backhaul link. Also, the amount of data cached in the WiFi hotspot can be smaller than the amount the node could download within the time it is in the hotspot's range. In this case, the node uses its remaining time in the WiFi hotspot to transfer data, as above, from the video's original location (Line 12).

4 Trace-driven evaluation

Our evaluation considers empirical measurements for the throughput of the cellular network and the SNR (Signal-to-Noise Ratio) of WiFi networks along a route between two locations in the center of Athens, Greece. Along the route we embed 2, 4, and 8 WiFi hotspots for different scenarios investigated. The details of the trace-driven simulation environment is presented in an extended version of this paper [12]. Moreover, here we present a small set of our evaluation results.

Number of video streams: Figure 1(a) shows the number of paused frames as a function of the number of video streams, illustrating the high gains that can be achieved with prefetching: For 4 HD streams, prefetching achieves 48% and 76% fewer paused frames compared to when prefetching is not used (i.e., the WiFi

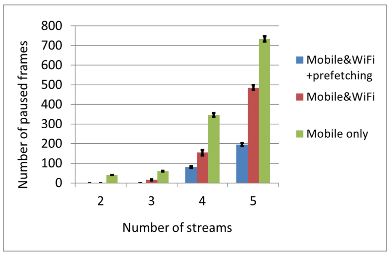

(a) Number of video streams

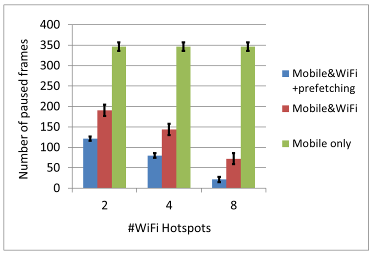

(b) Number of hotspots

Fig. 1. Paused frames for a different number of video streams and hotspots.

network is used opportunistically), and when only the mobile (cellular) network is used. Moreover, prefetching can support 3 HD streams without paused frames, whereas WiFi without prefetching has 15 paused frames and using only the mobile network results in 60 paused frames.

Number of WiFi hotspots: Figure 1(b) shows that the performance for the two mobile & WiFi schemes improves with more hotspots. Moreover, prefetching achieves more than 36% fewer paused frames compared to when prefetching is not used.

Throughput variability: Figure 2(a) shows the influence of the variability of the mobile, WiFi, and ADSL backhaul throughput. A higher throughput variability increases the number of paused frames for all three schemes. Nevertheless, prefetching can achieve more than 45% fewer paused frames compared to the mobile & WiFi scheme without prefetching, and more than 75% fewer paused frames than when only the mobile network is used.

Energy consumption: Figures 2(b) shows the energy consumption for a different number of video streams. Prefetching achieves lower energy consumption compared to the case where WiFi is used without prefetching, which in turn achieves lower energy consumption compared to the case where only the mobile network is used. The gains in terms of energy efficiency are comparatively lower than the gains of increased performance, in terms of a fewer number of paused frames, Figures 1(a). This is due to the significantly higher power consumption per transferred data volume of the mobile network compared to WiFi.

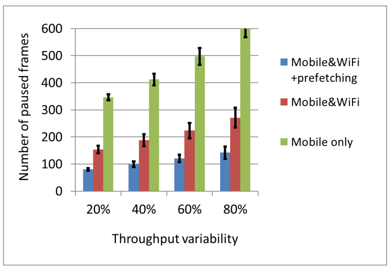

(a) Throughput variability

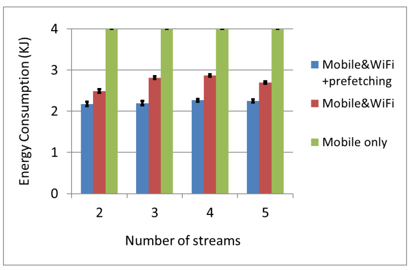

(b) Energy consumption

Fig. 2. Influence of throughput variability and energy consumption.

5 Conclusions and Future Work

We have presented and evaluated a procedure that exploits mobility and throughput prediction to prefetch video data in integrated mobile and WiFi networks, in order to improve mobile video streaming. The procedure can significantly reduce the number of paused frames (QoE), while being robust to time and throughput variability and achieve reduced energy consumption. The proposed approach has been implemented in an Android client for video playout that can stream video data from multiple servers.

Ongoing work includes evaluating the performance of our prototype implementation in a realistic setting. Also, future work includes adopting a standard protocol, based on MPEG-DASH (Dynamic Adaptive Streaming for HTTP), for communication between the mobile client and the video servers or caches.

References

1. A. Balasubramanian, R. Mahajan, and A. Venkataramani. Augmenting mobile 3G using WiFi. In *Proc. of ACM MobiSys*, 2010.
2. P. Deshpande, A. Kashyap, C. Sung, and S. Das. Predictive Methods for Improved Vehicular WiFi Access. In *Proc. of ACM MobiSys*, 2009.
3. K. Evensen et al. Mobile Video Streaming Using Location-Based Network Prediction and Transparent Handover. In *Proc. of ACM NOSDAV*, 2011.
4. N. Golrezaei et al. FemtoCaching: Wireless Video Content Delivery through Distributed Caching Helpers. In *Proc. of IEEE Infocom*, 2012.
5. X. Hou, P. Deshpande, and S. R. Das. Moving Bits from 3G to Metro-Scale WiFi for Vehicular Network Access: An Integrated Transport Layer Solution. In *Proc. of IEEE ICNP*, 2011.
6. K. Lee, I. Rhee, J. Lee, S. Chong, and Y. Yi. Mobile Data Offloading: How Much Can WiFi Deliver? In *Proc. of ACM CoNEXT*, 2010.
7. F. Malandrino et al. Proactive Seeding for Information Cascades in Cellular Networks. In *Proc. of IEEE Infocom*, 2012.
8. A. J. Nicholson and B. D. Noble. BreadCrumbs: Forecasting Mobile Connectivity. In *Proc. of ACM Mobicom*, 2008.
9. P. Rodriguez et al. MAR: a commuter router infrastructure for the mobile internet. In *Proc. of ACM MobiSys*, 2004.
10. V. Singh, J. Ott, and I. Curcio. Predictive Buffering for Streaming Video in 3G Networks. In *Proc. of IEEE WoWMoM*, 2012.
11. V. A. Siris and M. Anagnostopoulou. Performance and Energy Efficiency of Mobile Data Offloading with Mobility Prediction and Prefetching. In *Proc. of IEEE CONWIRE, co-located with IEEE WoWMoM*, 2013.
12. V. A. Siris, M. Anagnostopoulou, and D. Dimopoulos. Improving Mobile Video Streaming with Mobility Prediction and Prefetching in Integrated Cellular-WiFi Networks (Extended version). Available as *arXiv:1310.6171 [cs.NI]*, 2013.
13. V. A. Siris and D. Kalyvas. Enhancing Mobile Data Offloading with Mobility Prediction and Prefetching. In *Proc. of ACM MOBICOM MobiArch Workshop*, 2012.
14. C. Tsao and R. Sivakumar. On effectively exploiting multiple wireless interfaces in mobile hosts. In *Proc. of ACM CoNEXT*, 2009.
15. J. Yao, S. S. Kahnere, and M. Hassan. An Empirical Study of Bandwidth Predictability in Mobile Computing. In *Proc. of ACM WinTech*, 2008.