

On the Application of Continuous Media Filters over Wireless Networks

Margaritis Margaritidis and George C. Polyzos

Center for Wireless Communications
and
Computer Systems Laboratory
Department of Computer Science and Engineering
University of California, San Diego
La Jolla, California 92093-0114

Abstract –The increasing demand for ubiquitous access to real-time multimedia over the Internet necessitates mechanisms that support various notions of Quality of Service. In this paper we present filtering as a mechanism for tailoring streams to the transmission characteristics of wireless links and the capabilities of the receiving devices, in an attempt to preserve a requested Quality of Service. We analyze the differences between filtering for multicasting and for wireless Web access and outline *MobiWeb*, our proxy and filter based architecture for wireless Web access. We validate our approach with simulation results.

I. INTRODUCTION

The tremendous explosion of the Internet in recent years was accompanied by a strong demand for rich multimedia applications. Text, data, audio and video are the most common media available over the Internet. They demonstrate a diverse set of traffic characteristics and pose different demands on network performance. End users contribute to the Internet heterogeneity by accessing it with devices that have diverse capabilities and use different communication technologies. The diversity of the devices ranges from powerful workstations and desktop PCs, through laptops, to PDAs, two-way pagers and mobile phones. In addition, the variety in communication technologies extends from the rather traditional, such as LANs, T1 lines, modems, ISDN lines, etc., to those more recently deployed, such as cable modems, DSL, cellular modems, wireless LANs, CDPD, and other wireless technologies.

The result of such heterogeneity is that traditional Internet protocols and communication approaches are inadequate to fully support the diverse requirements of the multimedia applications. Therefore, the applications receive services with lower Quality of Service than they demand and can be offered, while the utilization of Internet resources remains sub-optimal.

A solution to the problems of path and device heterogeneity is *filtering* of the data streams. The term filtering describes the transformation of a stream (or a set of streams) into another in order to modify its transmission characteristics and presentation requirements with minimal change to its semantics (i.e., the information content and overall "value" and impact of the original stream).

There are two main motivations for performing such a transformation. First, the path from the source to the destination might be bottlenecked by a link with transmission characteristics that cannot support the original form of the stream. Second, the destination itself might have limited capabilities and be unable to decode or present the stream in the original form. A filter can transform the stream to a new representation that can be efficiently transmitted to the end device and presented to the user.

The use of filters in packet-switching networks was introduced by Pasquale *et al.* [1, 2], as part of their solution for efficiently multicasting continuous media. Since then, filters became a necessary complement of many multimedia multicasting solutions. RSVP [3] and a several techniques based on it [4, 5] have adopted similar

concepts and solutions. *Vgw*, the video gateway for the MBONE, is an example of a multicasting scheme that uses filtering to transcode video between different formats [6]. RTP supports media mixing filters at intermediate nodes [7]. Bhattacharjee *et al.* use "active networking," in which the network provides functions for the end user to exploit; a representative function is the filtering of MPEG video streams [8]. Finally, Yeadon *et al.* presented a complete taxonomy of existing filters for real-time media and an extensive evaluation of their computational and perceptual performance [9].

In the last few years a great push for wireless access to the Internet has been made. Telcos are shifting towards providing integrated data and voice services, abandoning the traditional circuit-switching approach. This migration is expected to peak in a few years with the introduction of 3G wireless. Considering that the number of mobile phone users is approximately twice the number of Internet users today, one can imagine the impact that such a migration might have on the services available over the Internet.

All previous approaches addressing Internet heterogeneity were based on the assumption that data rates were the only characteristic that differentiates end users. With wireless access however, a whole new set of parameters have to be addressed. Mostly unpredictably variable bandwidth, increased delay, significantly high error rates and devices with (possibly very) limited capabilities are a few to mention.

Some of the first filtering approaches addressing wireless access were focused on the filtering process itself, rather than the architectural framework of when, where and how to use it, which is important in situations where mobility is involved [10]. Others considered the dynamic adaptation required in a variable environment, but narrowed their research down to performing only image transcoding [11]. This motivated us to introduce *MobiWeb*, a proxy-based architecture that supports adaptive real-time multimedia applications over wireless links [12, 13, 16]. *MobiWeb* allows adaptive applications to build their own adaptation policies and transparently executes them by applying the appropriate filters on the incoming streams according to the current link conditions.

The recently developed and similarly motivated WAP standard is designed to support transparent access to the Web over a multitude of wireless links [14]. In its current stage of development, however, it doesn't provide support for real-time continuous media, which is the main focus of our work.

II. MULTICAST VS. WIRELESS FILTERING

With the introduction of wireless access, a new approach to where and when to use filtering is needed. Filters were used so far in multicast solutions that assumed heterogeneity only in the available bandwidth. Thus, their major concern was to preserve as much bandwidth as possible from being consumed by the multicasting stream in the Internet infrastructure. In order to achieve that, they tried to move the filter as close to the source as possible. The

propagation of the filter towards the source was stopped when a node was reached with children that demanded to receive the stream with better quality (or an incompatible format) than the filter provided.

In contrast, in the typical mobile/wireless setting, the last hop is a wireless link and it is important to locate the filter as close to the wireless link as possible. The reason for that is the variability of the link and the delay in responding to it. With the filtering mechanism close to the link, it can adapt its behavior according to the current link characteristics, and do so with minimum delay.

There is, however, a trade-off associated with the exact location of the filter. The trend towards pico-cellular architectures increases the frequency of handoffs between cells. Locating the filter just before the final physical hop (e.g., at the base station) means more frequent handoffs, which require migration and setup of the filters at the new cell. Such an operation can prove costly and time consuming and thus it should be avoided by putting the filter earlier in the path. An efficient compromise is to put the filter at the base station controller. This controller is responsible for the management of macro-cells, between which the mobile device hands-off less frequently. Thus the communication overhead for the filter setup is reduced, while the functionality remains virtually intact, since the filter is still fairly close to the limited wireless link.

In addition to assuming only heterogeneity in bandwidth, multicasting solutions have also difficulties dealing with variations in the client's available bandwidth after the initiation of the stream. The high variability of the wireless link characteristics however, makes this an important requirement. The filtering mechanism must change dynamically due not only to limited and variable bandwidth, but also to competition for the available resources with other streams and devices.

Such competition reveals another important consideration; resource reservation and admission control for streams. Even though present trends at the core of the Internet can effectively allow protocols to assume resources to be abundant, this is still not the case with wireless links. Therefore, resource reservation and admission control is still a very important consideration for wireless access to real-time multimedia.

Finally, there is a difference in the migration path that filters follow in the two cases. For multicasting, filters might migrate from the destinations towards the source, moving along the end-to-end path, in order to optimize resource consumption, typically during only the set-up phase, or when there is a change in the composition of the multicast set or the network topology. In the wireless case, however, filters have to move between base station controllers whenever the mobile hands-off to a new cell not covered by this controller. This means that traffic from a given client will now have to compete with a different mix of applications for the available resources in the new cell. Therefore, admission control must be applied again and filtering might have to adjust to the characteristics of the new wireless link.

III. IMPACT OF WIRELESS LINK CHARACTERISTICS ON FILTERING

In addition to where and when to use filtering, wireless access poses yet another novel question: how to use filtering. The reason for such a concern is the different transmission characteristics that a wireless link exposes compared to a wireline one.

In a wireless connection, bandwidth is not only limited, but also variable. Thus, selecting the appropriate filter is not a one-time decision to make at the beginning of the transmission. Instead, a monitoring mechanism should constantly record the link fluctuations and assist the filtering mechanism in adapting accordingly. Such adaptation, in conjunction with the extreme variability of the link, might cause oscillations in the stream's perceptual quality, which is

undesirable for the end user. To avoid this, the monitoring mechanism should also consider the margin of tolerance, before adaptation must be enforced, and the maximum frequency in changes allowed for this stream.

Another factor that differentiates the use of filters over wireless links is the increased delay that these links introduce. In real-time transmission, each frame must reach the client within a given deadline or it cannot be displayed. Moreover, in certain formats, like MPEG and H.261, the inter-coding of frames might exaggerate this problem. Frames that arrive late might have other frames dependent on them for their decoding. Therefore, their delay will cause not one, but several consecutive frames to not be displayed. Such a blackout effect is extremely annoying to the user, who would definitely appreciate some lower quality performance than a total blackout.

Evidently, filtering over wireless links must also be intelligent, in preempting and transmitting the most important frames first. This requires prioritization of packets belonging to different types of frames and, possibly, a different service for each packet-type, tailored to its importance. As an example, consider an MPEG stream with I, P and B frames. Clearly, I frames are the most important ones, since all other frames in a GOP (which typically spans from 4 to 15 frames), depend on the first I frame of the GOP for decoding. Loss or excessive delay of it can cause several consecutive frames to not be displayed. Therefore, I frames must always have greater priority and may need more reliable transmission than other frames.

The reliable transmission of the important frames is most significant over wireless links, since the error characteristics of these channels are dramatically worse than the characteristics of wireline ones. The air interface introduces a considerably increased bit error rate, while the user's mobility causes periods of fading, which create bursts of errors due to limited transmission power and interference from other electromagnetic signals. In both cases, packets are arriving corrupted at the client or are lost during transmission, which usually forces the link layer to apply a recovery mechanism (like FEC or ARQ). This results in less available bandwidth for packets to be transmitted and increased delay before they reach their destination. It is imperative, therefore, to have a monitoring mechanism assisting the filtering process to quickly identify this situation and give precedence to packets with higher priority and more stringent timing constraints.

Finally, of significant importance in the design of filters are the capabilities of the end device that will receive the filtered stream. The multicasting schemes assumed that the end device would be a personal computer (laptop, desktop, or workstation), which would be capable of displaying virtually any kind of incoming stream. The only limiting factor was the quality of the Internet connection. In contrast, small devices, like mobile phones and PDAs, have significant limitations in their capabilities compared to those of a typical PC. Small, possibly monochrome, display, limited audio output, limited CPU power, limited memory and battery power are the most important of them. Filter design for wireless access should aim in meeting those capabilities, instead of just reducing the data rate of the stream. Since standardization in the capabilities of such devices is difficult, a wide range of specialized proprietary filters will probably be needed, each tailored to the limitations of the device it is designed to serve.

IV. TAXONOMY OF FILTERS

Based on the way they manipulate media and according to [2, 9], we can classify filters into four generic categories:

- *Hierarchical filter*: This filter operates on a sub-stream basis, where a flow is split into a number of related sub-streams or

TABLE 1: TRACE RESULTS FROM THE SIMULATION TESTS

Test #	MPEG Stream	Good frames	Corrupted frames	Delayed frames	Dropped frames
1	MPEG Stream 1 with filtering	933	74	58	15
	MPEG Stream 1 w/o filtering	934	251	66	642
2	MPEG Stream 2 with filtering	546	34	31	17
	MPEG Stream 2 w/o filtering	478	3	247	158

scalable extensions. The hierarchical filter is used to select and mix the appropriate sub-streams that match the requirements of the current level of quality of the stream.

- *Frame-dropping filter*: This is a media-discarding filter used to reduce the stream frame rate. The filter is aware of the different types of frames that the stream consists of and discards some, typically the less important ones, in order to achieve media continuity with quality similar (but inferior) to the one defined by the current level of quality.
- *Transcoding filter*: This filter performs conversion from one encoding scheme to another, usually operating in an intermediate representation of the stream. These filters tend to be fairly processing-power intensive.
- *Splitting/Mixing filter*: In situations that the source cannot provide hierarchically encoded streams, the splitting filter can take a single stream and split it into the required sub-stream structure. Conversely, the mixing filter is used to mix streams together or to multiplex audio and video into a single stream, when the coding supports this kind of structure (like MPEG).

These are the most popular categories of filters that are used by Internet applications in multicast scenarios. Because of their wide acceptance, it is fairly easy to become standardized and permanently implemented at wireless network gateways, which can significantly speed up and simplify the filtering process. However, proprietary filters will always be used and having the functionality to support them, like *MobiWeb* does, provides extreme versatility in accommodating present and future applications. As a final solution, in the case that both the client and the wireless network (e.g., gateway or base station controller) are missing the appropriate filter, they can query a filter server that operates as a well-known filter repository. This functionality resembles the way current Web browsers are dealing with requests for unknown plug-ins.

V. SIMULATION RESULTS

To better illustrate the significance of filtering, we performed several simulation experiments with our adaptation framework, *MobiWeb*, using NS-2, the Network Simulator tool [15]. We evaluated the performance of frame-dropping filters using MPEG-1 streams. A series of tests were executed on top of a simulated dynamic wireless environment, with variability in resources due to fluctuations in link quality and handoffs. In addition, we simulated background traffic in the form of constant-bit-rate (CBR) streams that use UDP without filtering and traditional *ftp* traffic over TCP.

We used MPEG streams in our tests for several reasons. MPEG utilizes three frame types with variable and typically very different sizes. Adapting it, therefore, according to the variability of the wireless link is as good of a challenge as adapting any other stream. In addition, MPEG's inter-frame encoding depicts how intelligence associated with the filtering can significantly improve the perceived quality of the stream. Finally, the widespread use of the MPEG




Figure 1: FPS of first test with filtering

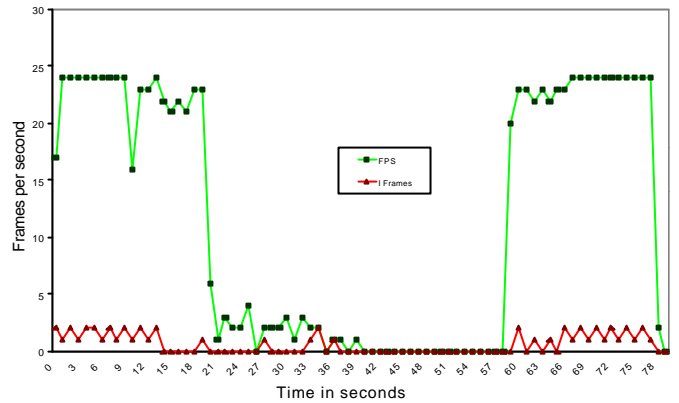


Figure 2: FPS of first test without filtering

format over the Internet today makes it the perfect candidate to evaluate the utility of filtering.

The frame-dropping filter we used decides whether to transmit a frame or not depending on the current link conditions and the importance of the frame. The order of importance is I frames first, then P and finally B. Table 1 shows sample results from the simulation experiments. In both experiments we present the tracing results when the MPEG stream was transmitted with and without filtering.

In the first test one can notice a significant amount of corrupted and dropped frames when filtering is not applied. Due to the inter-frame coding, corrupted frames introduce artifacts in the display, which propagate for several consecutive frames. In addition, dropped frames render other frames (dependent on them for decoding) useless, introducing periods of visible blackout. Furthermore, even though the amount of good frames in both cases is fairly equal, this doesn't translate in the same perceptual quality as well. Figures 1 and 2 show this difference. When filtering is applied, important I frames are receiving the best service, since they are necessary for decoding the other frames. When filtering is not applied, frames are indiscriminately dropped at the base station controller, since the incoming rate overflows the outgoing queue. This case favors the smaller P and mostly B frames, which can more easily be queued and transmitted over the wireless link than I frames. However, the lack of the anchor I frame prohibits their display, resulting in poor perceptual quality for the user.

In the second test we encountered a shifting in the arriving frames from being just corrupted to being delayed in the case where filtering wasn't applied. Apparently, the visual result was even worse since more frames were unable to be displayed and the blackout periods were prolonged. On the other hand, with the presence of filtering the percentages of both delayed and dropped

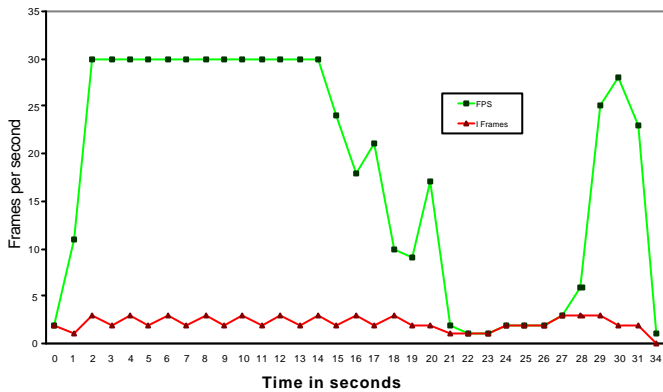


Figure 3: FPS of second test with filtering




Figure 4: FPS of second test without filtering

frames remained in the same small numbers, as they were in the first test. Consequently, the reproduced stream had reduced quality with occasional visible artifacts during link degradation, instead of long disturbing blackout periods. Figures 3 and 4 show again the differences between the two cases.

The tests proved that filtering with *MobiWeb* limited corrupted, delayed and dropped frames to a very small percentage of the total number of frames transmitted. In addition, important frames accepted better QoS, which lead to an overall enhancement in the perceptual quality of the displayed MPEG stream. In this paper we limited the simulation results into a single media type and a single filter type. We intent however to continue our research work by testing in the future a broad selection of media and filters under a variety of wireless link conditions.

VI. CONCLUSIONS

We argued in favor of filtering, as the mechanism for supporting Quality of Service for wireless access to the Internet. We discussed the use of filters in multicasting and wireless access scenarios and we identified the differences in where, when and how to use filtering. Furthermore, we indicated the impact that the peculiarities of the wireless links and the limited capabilities of the end devices have to the filtering process.

The simulation tests showed that filtering allows adaptive real-time streams to achieve decent quality over limited and variable wireless links. Instead of dropping packets indiscriminately, the filtering mechanism favors the transmission of the most important frames, achieving an acceptable quality for the MPEG streams. In addition, the simulation tests showed that the presence of an accompanying management scheme, like *MobiWeb*, proves extremely useful for the

intelligent transmission of real-time multimedia streams over variable wireless links.

ACKNOWLEDGEMENTS

This research was supported by the TRW Foundation, Rockwell Semiconductor Systems, the California MICRO program and the UCSD CWC.

REFERENCES

- [1] J.C.Pasquale, G.C. Polyzos, E.W. Anderson and V.P. Kompella, "The Multimedia Multicast Channel," *Proc. 3rd NOSSDAV*, La Jolla, CA, pp. 197-208, November 1992.
- [2] J.C.Pasquale, G.C. Polyzos, E.W. Anderson and V.P. Kompella, "Filter propagation in dissemination trees: Trading off bandwidth and processing in continuous media networks," *Proc. 4th NOSSDAV*, Lancaster, UK, pp. 269-278, November 1993.
- [3] L. Zhang, S. Deering, D. Estrin, S. Shenker and D. Zappala, "RSVP: a new resource ReSerVation Protocol," *IEEE Network*, vol. 7, no. 5, pp. 8-18, September 1993.
- [4] L.C. Wolf, R.G. Herrtwich and L. Delgrossi, "Filtering Multimedia Data in Reservation-Based Internetworks," *Kommunikation in Verteilten Systemen*, 22-24.02, TU Chemnitz-Zwickau, 1995.
- [5] R. Wittmann, M. Zitterbart, "Towards support for heterogeneous multimedia communications," *Proc. 6th IEEE Computer Society Workshop on Future Trends of Distributed Computing Systems*, Los Alamitos, CA, pp. 336-341, 1997.
- [6] E. Amir, S. McCanne and H. Zhang, "An Application-level Video Gateway," *Proc. of ACM Multimedia*, pp. 255-265, November 1995.
- [7] H. Schulzrinne, S. Casner, R. Frederick and V. Jacobson, "RTP: a transport protocol for real-time applications," *RFC 1889*, Internet Engineering Task Force, January 1996.
- [8] S. Bhattacharjee, K. Calvert and E.W. Zegura, "An Architecture for Active Networking," *Proc. High Performance Networking (HPN'97)*, White Plains, NY, April 1997.
- [9] N. Yeadon, F. Garcia, D. Hutchison and D. Shepherd, "Filters: QoS support mechanisms for multipeer communications," *IEEE JSAC*, vol. 14, no. 7, pp. 1245-1262, September 1996.
- [10] A. Fox, S.D. Gribble, E. A. Brewer and E. Amir, "Adapting to network and client variability via on-demand dynamic distillation," *Proc. 7th ACM Conference on ASPLOS*, October 1996.
- [11] R. Han, P. Bhagwat, R. LaMaire, T. Mummert, V. Perret and J. Rubas, "Dynamic Adaptation in an Image Transcoding Proxy for Mobile WWW Browsing," *IEEE Personal Communications*, vol. 5, no. 6, December 1998.
- [12] M. Margaritidis and G.C. Polyzos, "Application-assisted adaptation of real-time streams over wireless links," *Proc. 2nd UCSD Conference on Wireless Communications*, San Diego, CA, March 1999.
- [13] M. Margaritidis and G.C. Polyzos, "Wireless network support for adaptive real-time applications," *Proc. Advanced Simulation Technologies Conference (ASTC99)*, San Diego, CA, April 1999.
- [14] The WAP Forum, Wireless Application Protocol, <http://www.wapforum.org>
- [15] The VINT Project, "Network Simulator version 2," <http://www-mash.CS.Berkeley.EDU/ns/ns.html>
- [16] M. Margaritidis and G.C. Polyzos, "*MobiWeb*: Enabling Adaptive Continuous Media Applications over Wireless Links," *IEEE International Conference on Third Generation Wireless Communications*, Silicon Valley, San Francisco, California, June 2000.