Performance Evaluation of Simulcast vs. Layered Multicasting over Best-effort Networks

Christos Bouras, Apostolos Gkamas and Georgios Kioumourtzis
University of Patras, Computer Engineering and Informatics Department
&
Research Academic Compute Technology Institute
Rion, GREECE
E-mail: bouras@cti.gr, gkamas@cti.gr, gkioumou@ceid.upatras.gr

Abstract: **Multi-rate multicast schemes can be broadly classified into two categories. In layered multicast, a video file is transmitted by a base layer, which contains the most important features of the video. Additional layers, called enhancement layers, contain data that refine the quality of the base layer. In simulcast, the video file is transmitted by replicated layers that contain the same content at different quality. The benefits of layered multicast versus simulcast are still under question, as layered multicast presents higher complexity and more challenging deployment than simulcast. In this paper, two multi-rate multicast control schemes are compared. The layered multicast SMCC congestion control against our proposed solution for simulcast transmission, named ASSP. We compare the two schemes under a controlled simulation environment with the network simulator software (ns-2) by taking into account the evaluation criteria in RFC 5166. The results demonstrate that both SMCC and ASSP are TCP-friendly while SMCC seems to suffers from small oscillations of the transmission rate. In network topologies with low complexity ASSP consumes no more bandwidth than SMCC for the transmission of the different simulcast streams, while being a simpler solution than the more complicated SMCC.**

1. INTRODUCTION AND RELATED WORK

The heterogeneity requirements of networks and users' devices require that video streaming applications should be adaptive. The current advances of Scalable Video Coding (SVC) [1] in which each video file is encoded into one base layer and several enhancement layers fit well with layered transmission schemes. In contrast, simulcast is referenced to the transmission of a number of independent streams with the same content that differ in quality and hence in bandwidth requirements. The advantage of having different versions of the same content is that it does not require the more sophisticated scalable encoders and does not incur the extra overhead due to scalable encoding. The drawback is that the multiple versions of the same multimedia information are transmitted over the network in parallel (something that might be consider as waste of bandwidth) so that users can choose the appropriate version at any given time. A second major issue related to the transmission of multimedia data is "TCP-friendliness", so that TCP-based applications will not starve. The research community has provided a sufficient number of new proposals on simulcasting ([2], [3], [4]). A representative is the Destination Set Grouping (DSG) [5] in which the source transmits three streams at respectively low, medium and high quality data of the same video content. A receiver subscribes to a stream that is closer to its requirements. SRAMT-S [6] is a later simulcast solution in which the transmission rate of each simulcast stream is adaptive. In SRAMT-S, the adaptation strategy is based on joint sender-receiver bandwidth estimation. The sender adjusts the transmission rate of each simulcast stream based on these estimations.

In the layered multicast area, the first practical adaptation protocol was proposed in [7]. In this protocol, known as receiver-driven layered multicast (RLM) the sender transmits each video layer over a separate multicast group. Receivers periodically join a higher layer's group to explore the available bandwidth. However, the drawback in this approach is that one receiver's join experiments can introduce packet losses experienced by other receivers that share the same bottleneck link. Later works ([8], [9]) try to mitigate the effects of long IGMP leave latencies due to frequent changes of a receiver subscription level, which leave the network in a congested state. Fine-grained layered multicast [10] addresses the drawbacks of cumulative layering regarding the coarse-grained adaptation of receivers.

Simulcasting and layered multicast have been compared in many different contexts, including IP-layer multicasting [3], and TCP-friendliness [4]. Recent studies [11] on IPTV networks indicate that in some cases layered multicast is less efficient than simulcast if for most of the channels one resolution is needed. On the other hand the transmission of a big number of simulcast streams with the same content waste network bandwidth. Therefore, it is important to minimize the number of the transmitted simulcast streams in order to better utilize the available network resources. This can be feasible only if the streams are adaptive, so that they can serve a large number of users with similar receiving capabilities.

We present in this work a simulation-based comparison of a layered multicast scheme, named SMCC (Smooth Multi-rate Multicast Congestion Control) [12] against ASSP (Adaptive Smooth Simulcast Protocol) [13]. ASSP is our proposed simulcast solution in which except for the small number of required streams it demonstrates smooth transmission rates and TCP-friendliness. Both of these solutions make use of “dynamic” layers in which the transmission rate of each layer is defined by the “slowest receiver” in the multicast stream.

The rest of this paper is organized as follows: In the next section, we briefly describe the functionality of SMCC and ASSP. In section 3, we present the comparison of the two solutions based on simulation results with the network simulation software (ns-2) [14]. We conclude our paper in section 4.

2. PROTOCOLS OVERVIEW

SMCC is a multiple rate equation-based congestion control algorithm for cumulative layered multicast that employs TFMCC [15] as the primary underlying control mechanism for each layer. Since each layer is controlled independently by TFMCC, the properties of TFMCC hold for all participants in the multicast session in any given layer. As such, the layer rates are both dynamic and adaptive.

SMCC combines the benefits of TFMCC (smooth rate control, equation-based TCP friendliness) with the scalability and flexibility of multiple rates to provide a multiple rate multicast congestion control policy. In SMCC, receivers cumulatively subscribe to appropriate layers based on theirs estimated rate using the TCP throughput equation also employed in TFMCC. In addition to the TFMCC functionality, SMCC provides a new additive increase join attempt to avoid abrupt rate increases when a receiver attempts to subscribe to an additional layer. However, the calculated throughput using equation-based methods may not provide a sufficiently accurate indication to decide when to join the next layer. To avoid these problems, the receiver in SMCC joins the next layer through the join attempt when its calculated throughput is in the range of the next layer rate. Ultimately, the smooth rate change of SMCC is ideally suited to streaming multimedia applications.

ASSP is a new multi-rate transport protocol for simulcast transmission over best-effort networks. The key attributes of ASSP are: a) TCP-friendly behavior, b) adaptive per-stream transmission rates, c) adaptive scalability to large sets of receivers, and finally, d) smooth transmission rates that are suitable for multimedia applications. The building block of ASSP is based on our previous work that comprises a single-rate multicast congestion control protocol named ASMP [16]. Performance evaluation results [17] indicate that ASMP is a serious competitor to well-known single rate schemes, TFMCC and PGMCC [18]. ASSP is the extension of ASMP from the single-rate multicast congestion control schemes to simulcast transmission. As a result, the transmission of each stream, in the context of ASSP, is based on the underlying congestion control mechanism. The ASSP itself is responsible for handling all the issues related to simulcast transmission as well as the management and synchronization of the multiple multicast streams.

ASSP exploits the concept of “smooth transmission” to avoid high oscillations of the transmission rates of each individual stream, and minimizes the join and leave attempts. This is an important attribute as frequent join and leave requests introduce network congestion due to long IGMP leave latencies. Furthermore, frequent join and leave requests lead to high oscillations that cause instability. Another important attribute of ASSP is “TCP-friendliness” as each individual receiver calculates a TCP-friendly bandwidth share with the use of the TCP analytical model. Lastly, ASSP is a pure end-to-end solution that does not require any network support except for IP-multicast. Therefore, ASSP can be easily deployed over different administrative domains. More details on the functionality of ASSP along with preliminary design experiments can be found in [13].

3. PERFORMANCE EVALUATION

For the purpose of this work, we integrate the publicly available codes of SMCC [19] and ASSP [20] into ns-2 and conduct our experiments under a controlled environment.

Performance evaluation criteria of transport protocols are discussed in the recent RFC 5166 [21]. Therefore, by taking into account RFC 5166 we investigate the performance of SMCC and ASSP in the following areas:

- The network utilization level and the fairness of each protocol towards TCP traffic.
- The fairness among flows of the same protocol.
- The delay and the packet loss ratio that are introduced by the two protocols.
- The stability of each protocol under changing conditions, especially their ability to prevent oscillations.

For our experiments, we use the same network topologies with identical attributes to obtain a complete and fair comparison between SMCC and ASSP. The codes, simulation scripts and results are publicly available in [20].

3.1 Network utilization and TCP-friendliness

To investigate the ability of each protocol to exploit network resources and their TCP-friendliness we use the simulation scenario in Figure 1. In this topology, we set up three bottleneck links, each with different capacity. Therefore, receivers are connected with links that differ in capacity. Each bottleneck link is shared by two multicast receivers and one TCP flow. C0 to C4 stand for the network routers. We set up Drop Tail queues in the routers and set the one-way delay in all paths to 32 ms. The packet size for all flows is set to 1000 bytes to obtain a fair comparison. We run the following simulations for ASSP and SMCC. In the ASSP case, the sender transmits three multicast streams with the following limits: Stream1 (300 Kb/s - 2 Mb/s), Stream2 (2 Mb/s - 5 Mb/s), and Stream3 (5 Mb/s - 10 Mb/s). At start time, all ASSP receivers join the lowest capacity stream,

Stream1. For the SMCC simulations, we set the transmission rate of the base layer to 1 Mb/s.

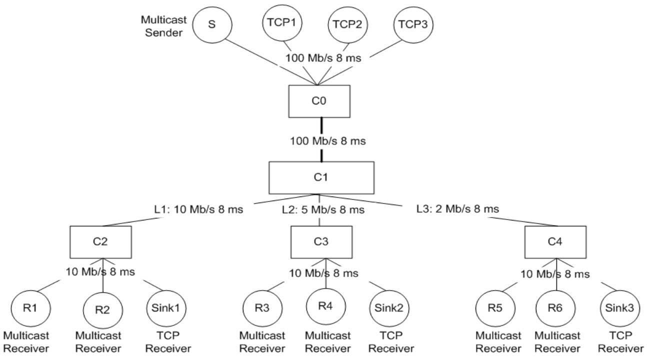

Figure 1 - Bottleneck scenario

We can see (Figures 2 and 3) how the two protocols adjust their transmission rates to satisfy the receiving requirements in the different bottleneck links. The transmission rates of both protocols are dynamically adjusted with respect to the observed congestion in the three different links. With a visual comparison, we observe that ASSP presents more stable behavior than SMCC. This however, will be discussed in the next sections in which we investigate the stability of the two protocols in paths with varying RTT values. In Table 1 are given the simulation statistics of the different flows in terms of the achieved throughput. In a "perfect world", network resources should be equally shared by multicast and TCP flows and network utilization would reach up to one hundred per cent. ASSP presents lower performance in the higher capacity link (L1) while SMCC has lower performance in the lower capacity link (L3). We also observe that SMCC present higher bandwidth requirements in the link connecting C0 to C1 for the transmission of the three streams than ASSP. SMCC calculates the maximum transmission rate of the different multicast streams based on the transmission rate of the base layer. Under SMCC, the maximum cumulative sending rates are defined as follows:

$$B_i = 2^i \cdot B_0 \text{ for } i \geq 1 \tag{1}$$

where, B_0 is the transmission rate of the base layer and B_i is the transmission rate of layer i. To further investigate the impact of the initial rate for the base layer in SMCC we run various simulations and summarize the results in Table 2 (all values in Mb/s). We measure the bandwidth requirements for the three multicast streams in the link connecting routers C0 and C1. SMCC bandwidth consumption depends on the transmission rate of the base layer and varies from 6.49 Mb/s to 7.18 Mb/s. SMCC's performance is highly coupled to the initial setting of the base layer. Therefore, it is important for SMCC to choose an optimal value for the base layer during the initiation of the multicast session. This area requires a deeper investigation on how to define such an algorithm.

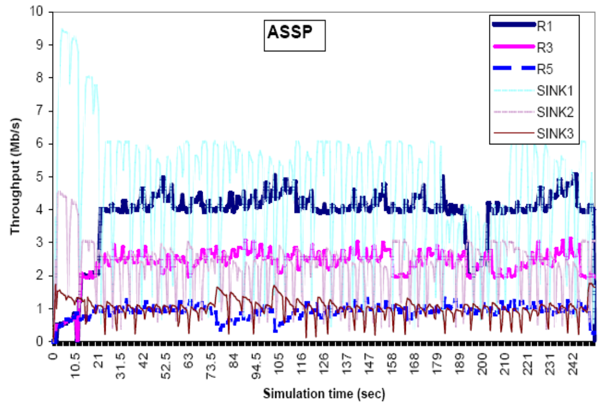

Figure 2 - Transmission rates of ASSP flows

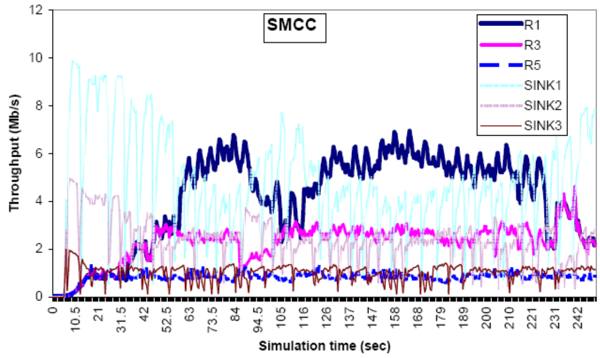

Figure 3 - Transmission rates of SMCC flows

Table 1 –Simulation statistics

		Average Throughput (Mb/s)	Link Utilization (%)	TCP Average Throughput (Mb/s)	BW requirements between C0 and C1 (Mb/s)
SMCC	L1	4.28	90.6	4.78	7.4
	L2	2.28	93.2	2.38	
	L3	0.84	92.5	1.01	
ASSP	L1	3.88	84.4	4.56	7.1
	L2	2.37	92.8	2.27	
	L3	0.92	97.0	1.02	

ASSP on the other hand, depends on the settings for the upper and lower thresholds of the simulcast streams. The upper and lower thresholds however, simply map the capacity of the bottleneck links in the path between the sender and the groups of multicast receivers. This seems to be a simpler task if we consider that receivers can inform the sender about their receiving capabilities (depending on their access link for example) during the initiation of the multicast session. A low complexity algorithm can define the optimal number of simulcast streams that have to be transmitted in order to satisfy the majority of receivers in the multicast group. This is however left as a future work. The measured Jain's fairness index [22] for both protocols in all the bottleneck links with competing TCP flows is 0.99. Therefore, ASSP and SMCC proved to be TCP-friendly as almost equally share networks recourses with the TCP flows.

Table 2 –SMCC link utilization under different transmission rates for the base layer

Base Layer (Mb/s)	Average Throughput (Mb/s)		Link Utilization (%)	BW requirements between C0 and C1 (Mb/s)
BL=0.5	L1	3.69	73.8	6.49
	L2	1.96	78.4	
	L3	0.84	84.0	
BL=1.0	L1	4.11	82.2	7.18
	L2	2.26	90.4	
	L3	0.81	81.0	
BL=2.0	L1	4.30	86.0	6.81
	L2	1.67	66.8	
	L3	0.84	84.0	

3.2 Intra-protocol fairness

In this simulation scenario, we investigate the fairness between flows of the same protocol with the network topology in Figure 4. In this scenario, two multicast senders transmit at opposite directions to different user groups. We summarize the simulation results in Table 3. We observe that both ASSP and SMCC present high level of intra-fairness as the available bandwidth in the bottleneck links is fairly distributed between the two multicast sources. However, the two protocols present different levels of link utilization. ASSP seems to be more efficient than SMCC except for the case in the low capacity link.

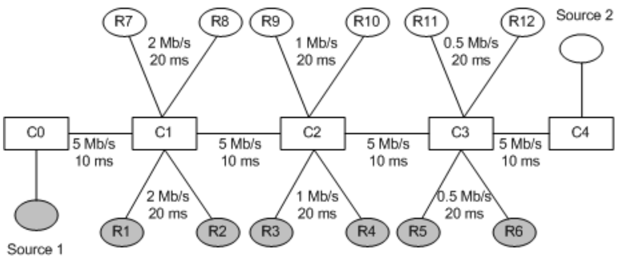

Figure 4 – Network topology for protocol intra-fairness

3.3 Delay and Packet Loss Ratio Measurements – Minimizing Oscillations

We conduct a number of simulations with different RTTs in the network topology of Figure 1. Both SMCC and ASSP have similar performance in terms of the average delay although ASSP outperforms SMCC in all cases (Table 4). Our assessment is that both protocols present high performance with delay values that are within acceptable rates. Although multimedia applications are known as "packet loss tolerant applications", we cannot underestimate the effects of packet losses on any application type. Therefore, packet loss ratio is an important attribute for any multimedia transmission protocol.

Table 3 – Intra-protocol fairness simulation statistics

Link capacity	ASSP Average Throughput (Mb/s)		Jain's Index	SMCC Average Throughput (Mb/s)	Jain's Index
L1 (2 Mb/s)	Source 1	0.684	0.999	0.379	0.991
	Source 2	0.671		0.454	
L2 (1 Mb/s)	Source 1	0.346	0.994	0.379	1.0
	Source 2	0.403		0.379	
L3 (0.5 Mb/s)	Source 1	0.234	0.999	0.443	0.999
	Source 2	0.243		0.436	

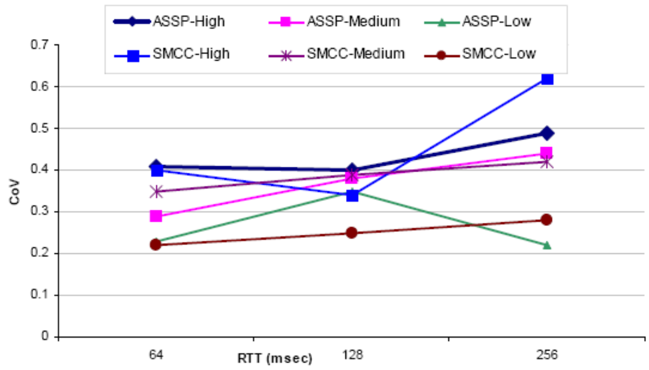

Figure 5 - CoV of average throughput over different RTTs

Table 4 – Simulation statistics over different RTTs

	RTT (msec)	Average end-to-end delay (msec)	Loss ratio (%)
SMCC	64	91	0.08
	128	116	0.05
	256	169	0.03
ASSP	64	86	0.2
	128	100	0.16
	256	157	0.2

We observe that, ASSP presents higher loss ratio than SMCC in all cases, although the values are very low for both protocols. The reason behind ASSP's lower performance is the higher feedback intervals of the RTCP [23] protocol compared to those of the TFMCC. TFMCC's sender requires receiving one feedback report from the group representative on every RTT while ASMP's feedbacks are regulated by the RTCP protocol with higher intervals. To measure the stability of the two protocols we use the coefficients of variation (CoV)[1] of the throughput values over the different RTTs and plot the results in Figure 5. We observe that SMCC presents higher oscillations than ASSP when the RTT values increase.

4. CONCLUSION AND FUTURE WORK

[1] Coefficient of Variation (CoV) is the standard deviation divided by the mean.

In this work we presented a simulation-based comparison of two proposed schemes for multi-rate multicast transmission. The performance of SMCC was highly coupled to the initial settings of the base layer while the ASSP control protocol needed to know in advance the available bandwidth of user groups behind the same bottleneck link. Both protocols proved to be TCP-friendly as equally shared network resources with TCP flows. The Jain's fairness index was 0.99 in all simulation scenarios. ASSP outperformed SMCC in terms of link utilization and scaled better in the high capacity links in almost all of the simulation sets. Delay measurements showed that both protocols had a good performance while ASSP was more efficient than SMCC. SMCC responded better to network congestion than ASSP whose feedback functions depended on the high intervals of the RTCP protocol. As a result, ASSP was "slower" than SMCC in responding to network congestion and presented higher loss ratio. Simulations with different RTTs showed that ASSP seemed to be more stable than SMCC.

A deeper investigation is needed on how to define the algorithms for optimal performance for both protocols. A possible solution for SMCC would be on-the-fly selection of the transmission rate for the base layer, as network conditions change in an unpredictable and dynamic manner. On the other hand, ASSP should also be able to optimize the number of the available streams based on network conditions.

A more complete performance evaluation based not only on network-centric metrics but also on specific multimedia-based quality metrics would provide a better insight on the advantages and drawback of the two tested solutions. All the above areas are left for future work.

REFERENCES

[1] ITU-T Recommendation H.264, "Advanced video coding for generic audiovisual services" July 2007.

[2] T. Jiang, E. W. Zegura, and M. H. Ammar, "Inter-receiver fair multicast communication over the Internet," in Proc. NOSSDAV'99, Jun. 1999.

[3] T. Kim and M. H. Ammar, "A comparison of layering and stream replication video multicast schemes," in Proc. NOSSDAV'01, Jun. 2001.

[4] P. de Cuetos, D. Saparilla, and K.W. Ross, "Adaptive streaming of stored video in a TCP-friendly context: multiple versions or multiple layers," in Proc. Packet Video Workshop, Apr. 2001.

[5] X. Li, M. Ammar, and S. Paul, "Video Multicast over the Internet," IEEE Network Magazine, vol. 13, no.2, Apr. 1999, pp. 46-60.

[6] C. Bouras, A. Gkamas, "SRAMT-S: A hybrid sender and receiver-based adaptation scheme for TCP friendly multicast transmission using simulcast approach", 1st IFIP Workshop on Internet Technologies, Applications and Social Impact (WITASI-02), Wroclaw, Poland, 2002.

[7] S. McCanne, V. Jacobson, and M. Vetterli, "Receiver-Driven Layered Multicast," Proc. ACM Sigcomm Conf., ACM Press, 1996, pp. 117-130.

[8] J.W. Byers et al., "FLID-DL congestion control for layered multicast", in Proceedings of NGC, 2000.

[9] J. Byers, G. Kwon, "STAIR: Practical aimed multirate multicast congestion control", in: Proceedings of NGC, 2001.

[10] J. Byers, M. Luby, M. Mitzenmacher, " Fine-grained layered multicast", in: Proceedings of INFOCOM 2001, Volume 2, Issue, 2001 Page(s):1143 – 1151.

[11] Avramova, Z. et al, "Comparison of Simulcast and Scalable Video Coding in Terms of the Required Capacity in an IPTV Network", in Proceedings of PV 2007, Lausanne, Switzerland, November 2007.

[12] G.-I. Kwon, J. Byers, "Smooth multirate multicast congestion control", in: Proceedings of IEEE INFOCOM, 2003.

[13] C. Bouras, A. Gkamas, G. Kioumourtzis, "Adaptive Smooth Simulcast Protocol for Multimedia Transmission", IEEE Symposium on Computers and Communications (ISCC'09), July 5 - 8, 2009, Sousse, Tunisia.

[14] "The network simulator ns-2" [Online]. Available: http://www.isi.edu/nsnam/ns/

[15] RFC 4654, "On TCP-friendly Multicast Congestion Control (TFMCC)", J. Widmer, M. Handley, August 2006.

[16] C. Bouras, A. Gkamas, G. Kioumourtzis, "Adaptive Smooth Multicast Protocol for Multimedia Data Transmission", 2008 International Symposium on Performance Evaluation of Computer and Telecommunication Systems – SPECTS 2008, Edinburgh, UK, 16 - 18 June 2008.

[17] C. Bouras, A. Gkamas, G. Kioumourtzis, "Comparison of Single-Rate Multicast Congestion Control Protocols vs. ASMP", 16th Annual Meeting of the IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS)– MASCOTS 2008, Baltimore, MD, USA, 8-10 Sep. 2008.

[18] L. Rizzo, "pgmcc: A TCP-friendly single-rate multicast congestion control scheme," in Proc. of ACM SIGCOMM '00, 2000.

[19] [Online].Available: http://cs-people.bu.edu/guin/smcc.html

[20] [Online].Available: http://ru6.cti.gr/ru6/ns_rtp_home.php

[21] RFC 5166, S. Floyd, "Metrics for the Evaluation of Congestion Control Mechanisms", Network Working Group, March 2008.

[22] R. Jain, D. Chiu, and W. Hawe, "A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Systems," DEC Research Report TR-301, Tech. Rep., Sept 1984.

[23] RFC 3550, RTP: A Transport Protocol for Real-Time Applications, H. Schulzrinne, S. Casner, R. Frederick, V. Jacobson, July 2003.