DISTRIBUTED RATE ALLOCATION FOR VIDEO STREAMING OVER WIRELESS NETWORKS WITH HETEROGENEOUS LINK SPEEDS

Xiaoqing Zhu and Bernd Girod

Information Systems Laboratory, Stanford University, CA 94305, U.S.A.
{zhuxq,bgirod}@stanford.edu

Invited Paper

ABSTRACT

When multiple video streaming sessions are simultaneously transported over a shared wireless network, careful rate allocation is needed for efficient utilization of the network resources, while avoiding network congestion. The problem becomes more challenging when the network comprises links with heterogeneous transmission speeds, and when each video stream has a different utility of its allocated rate. We propose a distributed scheme for congestion-distortion optimized rate allocation among multiple video streams sharing a network. Simulation results are presented for streaming multiple high-definition (HD) video sequences over both single-hop and multi-hop 802.11a networks with heterogeneous link speeds. In comparison with TCP-Friendly Rate Control (TFRC), the proposed cross-layer allocation scheme benefits from being both media-aware and network-aware, and achieves higher average video quality among all streams without incurring excessive network congestion.

1. INTRODUCTION

Simultaneous support of multiple video streaming sessions over a shared wireless network is attractive for many applications, ranging from surveillance camera networks to home entertainment systems. Many technical challenges arise due to the unpredictable nature of wireless links and stringent rate and latency requirements of media streaming. Careful rate allocation is needed to prevent the multiple streaming sessions from congesting the network, while attaining high video quality of all streams.

The problem of multi-stream rate allocation over wireless networks is further complicated by heterogeneity in both the video contents and wireless link speeds. The utility of the allocated rate is different for video streams with different contents: the same rate increase may impact a sequence containing detailed textures and fast motion rather differently than a head-and-shoulder news clip. In addition, wireless links may experience different channel conditions, thus operate at different link speeds. Consequently, packets of the same size sent over a slow link would occupy the shared wireless medium for a longer duration than those sent over a fast link.

In this work, we focus on joint rate allocation for multiple video streams sharing a wireless network consisting of heterogeneous link speeds. The problem is formulated in the convex optimization framework, by minimizing the Lagrangian cost of total video distortion and overall network congestion. We present a distributed solution based on iterative descent among the streams. Cross-layer information exchange between the link state monitor and video rate controller ensures that the allocate rates are updated according to changes in both the video contents and wireless network conditions.

In the next section, we review related work for multi-user rate allocation and performance issues introduced by heterogeneous link speeds in the commonly used 802.11 wireless networks. We then present our model of wireless networks and video distortion-rate (DR) characteristics in Section 3. The proposed distributed rate allocation scheme is described in Section 4. In Section 5, we discuss simulations results of streaming multiple high-definition (HD) video sequences over single-hop and multi-hop 802.11a networks with heterogeneous link speeds.

2. RELATED WORK

Rate allocation among multiple traffic flows over a shared network is an important and well-studied problem. Internet applications typically use the TCP Congestion Control mechanism for regulating traffic outgoing rates [1]. For media streaming applications over UDP, TCP-Friendly Rate Control (TFRC) is a popular choice [2]. The mathematical formulation of the problem, as well as two classes of distributed rate allocation algorithms corresponding to the primal and dual decomposition of the optimization are explained in [3]. Application of such algorithms has been investigated for elastic traffic over the Internet [4].

In the context of wireless networks, many performance issues have been identified when the 802.11 MAC protocol is used for ad hoc networking [5]. In particular, it has been pointed out in [6] that the presence of one stream traveling over a slow link significantly reduces the throughput achieved by other streams over faster links.

Such performance anomaly can be avoided by explicit channel time allocation among the competing traffic streams. A centralized solution is proposed in [7], where the transcoding rate of each outgoing video stream at the common source node is dynamically controlled based on both the rate-distortion characteristics and wireless link speeds of all streams in the system. In [8], a distributed solution is proposed for channel time allocation among video streams originating from different source nodes sharing the network.

3. SYSTEM MODEL

3.1. Wireless network model

Consider a network with a set of wireless links $\mathcal{L}$, and a set of video streams denoted by $\mathcal{S}$. The capacity over Link l is denoted as C_l, and is defined as the maximum throughput over that link, *without* competition or interference from traffic over other links. In other words, C_l is the maximum achievable data rate over Link l when the rest of the network is silent. Typically, the value of C_l depends on the link

This work is partially supported by NSF Grant CCR-0325639, and a gift from Sharp Labs of America, Inc.

transmission speed chosen at the MAC layer (which, in turn, depends on the wireless link quality), the average size of the payload packets, as well as the percentage of MAC and link-layer control overhead. It can be estimated as:

$$C_l = \frac{B_l}{B_l/C_l^0 + T^{overhead}}, \tag{1}$$

where B_l is the average payload packet size, C_l^0 is the bit rate at which the payload is transmitted over link l, termed as *nominal link speed*, and $T^{overhead}$ is the total overhead time in transmitting the accompanying MAC-layer control packets and headers. In wireless networks following the IEEE 802.11 standard, for instance, $T^{overhead}$ is the amount of time spent in transmitting MAC-layer control packets and headers at a *basic* rate typically lower than that used for the payload, resulting in significantly reduced link capacity as compared to the nominal link speed [9].

When multiple links are active in the network, they compete for the shared wireless radio channel, following the rules dictated by the MAC protocol. In general, interference among neighboring links may be incorporated by adjusting the transmission rate of each link, as in the case of CDMA networks. In this work, it is assumed that at any time instance, only one of the competing links is allowed to transmit, achieving capacity C_l. This model is similar to the CSMA/CA mechanism used in 802.11 networks.

We denote *link utilization* as the fraction of time occupied by a link with traffic rate F_l and capacity C_l:

$$u_l = F_l/C_l. \tag{2}$$

This can be further decomposed: $u_l = F_l'/C_l + \sum_s u_l^s$, where the rate of non-video *background traffic* is denoted as F_l' and the contribution of each video stream is represented as u_l^s:

$$u_l^s = \begin{cases} R^s/C_l & \text{Stream } s \text{ traverses Link } l \\ 0 & \text{otherwise} \end{cases}. \tag{3}$$

The set of links that cannot transmit simultaneously with Link l constitutes the *interference set* of Link l, denoted as $\mathcal{L}_l$. For a small network, all links interfere with each other, and have the same interference set $\mathcal{L}$, i.e., the entire network. In this case, the total utilization is constrained to be smaller than unity [1]:

$$u_{\mathcal{L}} = \sum_{l \in \mathcal{L}} u_l < 1. \tag{4}$$

In practice, the total channel time proportion is constrained to be $u_{\mathcal{L}} < \gamma$, where $\gamma < 1$ is an over-provisioning factor. The extra headroom is needed to absorb various effects not included in our model, such as random backoff in a CSMA/CA network to resolve link contention, or inaccurate estimates of instantaneous link capacity.

Congestion over each wireless link is measured as the average delay for all packets traversing that link. Following the classic M/M/1 queuing model, where average packet delay over a single link is inversely proportional to the residual link capacity [10], we propose to model the average packet delay for multiple links competing within the same network as:

$$X_l = \frac{B_l/C_l}{\gamma - u_{\mathcal{L}}}. \tag{5}$$

[1]The network model can be extended to the more general case with multiple over-lapping interference sets, by constraining the total utilization within each interference set $\mathcal{L}_l$ (for Link l) to be less than unity. In the rest of the paper, only the special case where all links share the same interference set $\mathcal{L}$ is discussed, so as to simplify notations.

In other words, congestion over each link X_l is inversely proportional to residual link utilization. Congestion of the entire network is calculated as the sum of all link delays weighted by the traffic flowing across each link:

$$X = \sum_{l \in \mathcal{L}} F_l X_l = \sum_{l \in \mathcal{L}} \frac{B_l u_l}{\gamma - u_{\mathcal{L}}}. \tag{6}$$

Note that, in practice, congestion might be a more complicated function of rate as predicted by the M/M/1 model. Nevertheless, the above expression can be viewed as an approximation of the average link delay, capturing the non-linear increase of delay with total channel time utilization.

3.2. Video distortion model

For each video stream $s \in \mathcal{S}$, its encoded video distortion D^s decreases nonlinearly with increasing encoding rate R^s. Typically, the distortion-rate (DR) characteristic of each stream is convex, and can be fit to a parametric model [11]:

$$D^s(R^s) = D_0^s + \frac{\theta^s}{(R^s - R_0^s)}, \tag{7}$$

where the parameters D_0^s, θ^s and R_0^s depend on the coding scheme and the content of the video. They can be estimated from three or more trial encodings using non-linear regression techniques. To allow fast adaptation of the rate allocation to abrupt changes in the video content, these parameters need to be updated for each group of pictures (GOP) in the encoded video sequence, typically once every 0.5 second.

In practice, only a finite set of encoded bitstreams at various rates might be available for transmission. Therefore, given the optimal rate allocation derived from the parametric model, the final rate allocation decision needs to be chosen as one of the available rates. It also needs to be noted that while we choose to use the parametric model $D^s(R^s)$ for describing the rate allocation algorithm, the optimization framework is general enough to accommodate other convex DR functions, such as a piecewise linear model.

4. DISTRIBUTED RATE ALLOCATION

4.1. Optimization objective

Intuitively, given the same wireless channel conditions at each time instant, network congestion will increase with higher allocated rate R^s. Excessive congestion will, in turn, cause severe degradation of received video quality, as more packets are dropped when they arrive after the playout deadline [12]. On the other hand, decreasing the allocated rate leads to higher video distortion during encoding. We therefore seek to strike a balance between both objectives, and propose to minimize the Lagrangian sum of total video distortion and overall network congestion:

$$\min_{R^s, s \in \mathcal{S}} \quad \sum_{s \in \mathcal{S}} D^s(R^s) + \lambda X. \tag{8}$$

The choice of λ determines the trade-off between video distortion and network congestion. Incorporating (6) and (7), this is equivalent to:

$$\min_{R^s, s \in \mathcal{S}} \quad \sum_{s \in \mathcal{S}} \frac{\theta^s}{R^s - R_0^s} + \lambda \sum_{l \in \mathcal{L}} \frac{B_l u_l}{\gamma - \sum_{l \in \mathcal{L}} u_l}. \tag{9}$$

It can be easily verified that the objective function in (9) is convex in the vector of variables R^s's. Therefore, the optimization can be solved using centralized solutions such as the interior point method [13]. The problem can also be solved in a distributed fashion, thereby avoiding the traffic overhead needed for collecting global information about the DR characteristic of each stream and capacity of every link in the network.

4.2. Distributed solution

From the perspective of Stream s, its contributions to the objective function (9) consist of the following:

$$J^s = D^s(R^s) + \lambda\Big(\sum_{l: u_l^s > 0} \frac{B_l(\overline{u}_l^s + R^s/C_l)}{\overline{\gamma}^s - R^s/\tilde{C}^s} + \sum_{l: u_l^s = 0} \frac{B_l u_l}{\overline{\gamma}^s - R^s/\tilde{C}^s} \Big). \tag{10}$$

In (10), the impact of the allocated rate R^s for Stream s on the overall network congestion is captured by the *equivalent capacity* $\tilde{C}^s$:

$$\frac{1}{\tilde{C}^s} = \frac{1}{\sum_{l: u_l^s > 0} C_l}, \tag{11}$$

total utilization from other traffic flows over each link $\overline{u}_l^s$:

$$\overline{u}_l^s = F_l'/C_l + \sum_{s' \neq s} u_l^{s'}, \tag{12}$$

and maximum allowed utilization given allocation from other traffic flows:

$$\overline{\gamma}^s = \gamma - \sum_l \overline{u}_l^s. \tag{13}$$

Given local knowledge of current video DR function $D^s(R^s)$ and parameters $\tilde{C}^s$, $\overline{u}_l^s$, $\overline{\gamma}^s$ collected along the links traversed by Stream s, the optimal rate R^s can be calculated by minimizing (10). Since the overall cost of a convex function given by (9) is reduced by optimizing over R^s at every step, the global minimum of the cost function can be reached by alternating the rate allocation among the multiple streams.

4.3. Protocol implementation

The design of a distributed rate allocation protocol must consider a number of additional practical requirements. The link capacities may experience abrupt changes, due to fluctuations in the wireless channel conditions. The set of available rates for each video stream is discrete, corresponding to different versions of the encoded bitstreams at different quality levels. In addition, cross-layer information exchange is needed, so that the video source rates can be adapted at the application layer, according to link states reported from the MAC layer. Figure 1 illustrates various components in such a system.

At the MAC layer, a link state monitor keeps an online estimate of the effective capacity C_l and background traffic rate F_l' by logging the arrival and departure of each packet. It also records the *intended* rate allocation R^s advertised by each stream, and calculates the link utilization per stream u_l^s accordingly. Periodic broadcast of link state messages are used to collect the values of C_l and u_l^s from neighboring links in the same interference set.

At the application layer, the video rate controller at the source advertises its intended rate allocation R^s in the video packet header. The link state monitor traversed by the stream then calculates the relevant parameters $\tilde{C}^s$, $\overline{u}_l^s$, $\overline{\gamma}^s$ on behalf of Stream s, based on its local cache of capacity, background traffic rate and per-stream utilization information of all links within its interference set. The destination node extracts such information from the video packet header and reports back to the sender in the acknowledgment packets, so that the video rate controller can re-optimize its intended rate R^s by minimizing (10), with updated link state information.

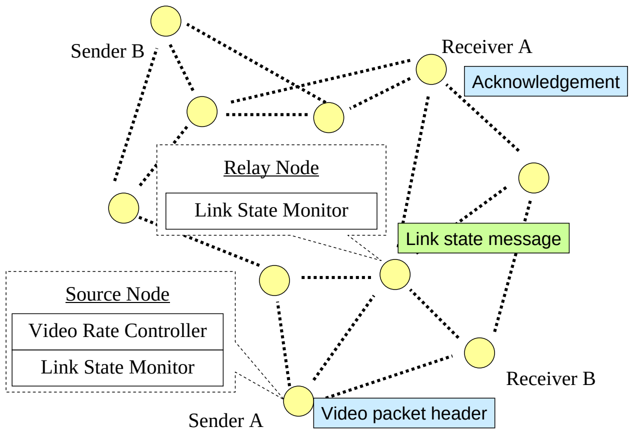

Fig. 1. Cross-layer information exchange among the link state monitor at the MAC-layer and the application-layer video rate controller. Nodes periodically exchange link state messages with its neighbors to collect capacity, flow and utilization of other links within the same interference set. The video packet header contains fields of advertised rate, link capacities and utilization, etc., to facilitate optimal rate allocation at the sender.

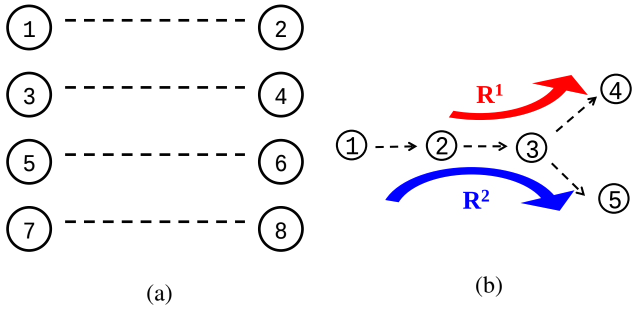

Fig. 2. Topology for network simulations. (a) Single-hop network with 8 nodes and 4 parallel connections. Horizontal distance between neighboring nodes is 30 m; vertical distance between neighboring nodes is 10 m. All nodes are within the transmission range of 53 m from each other. (b) Multi-hop network with 4 links and 2 streams. The distance between adjacent nodes is 50 m. All links are within the same interference set.

5. SIMULATION RESULTS

5.1. Experimental setup

We simulate in ns-2 [14] both single-hop and multi-hop wireless networks, with topologies depicted in Fig. 2. The parameters of SIFS/DIFS/EIFS slot time, random backoff window size and retry limits are chosen according to the specifications of the IEEE 802.11a standard [15]. The transmission power and receiving threshold of the nodes are adjusted to achieve a transmission range of 53 m. The basic rate for header and control packet transmission is set at 6 Mbps, whereas the nominal link speed for payload varies between 6 to 54 Mbps.

Four high-definition (HD) video sequences, *Bigships*, *Crew*, *Cyclists* and *Harbor*, are considered for streaming. The sequences have

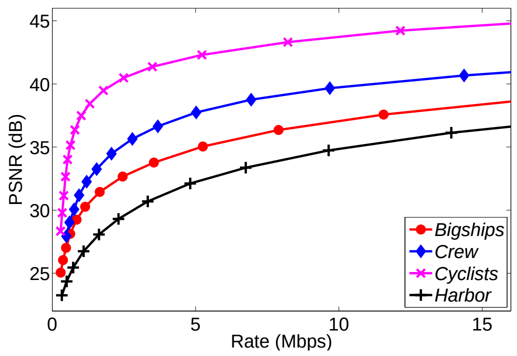

Fig. 3. Rate-PSNR performance of four HD video sequences used in the experiments: *Bigships*, *Crew*, *Cyclists* and *Harbor*.

a spatial resolution of 1280×720 pixels, and frame rate of 60 fps. They are encoded using a fast implementation of the H.264/AVC encoder [16][17] at various quantization step sizes, with GOP length of 30 and IBBP...structure similar to that used in MPEG-2 bitstreams. The trade-off between average encoded video quality and bit rates for the entire sequence lengths are plotted in Fig. 3. Encoded video frames are segmented into packets with maximum size of 1500 bytes, and the transmission intervals of each packet in the entire GOP are spread out evenly, so as to avoid unnecessary queuing delay due to the large sizes of intra coded frames.

Performance of the proposed cross-layer rate allocation scheme is compared against the conventional end-to-end approach based on TCP-Friendly Rate Control (TFRC):

$$R = k \frac{S}{RTT\sqrt{p}}. \tag{14}$$

In (14), the allocated rate R depends on the online estimates of average packet size S, round-trip-time RTT, packet loss rate p and a scaling constant k [2]. In our simulations, end-to-end packet losses almost never occur due to retransmissions by the 802.11a MAC protocol. The rate control formula is therefore simplified into $R = kS/RTT$. The allocated rate at any time instant is constrained to be no more than twice and no less than half of current rate.

Since small acknowledgment (ACK) packets incur much MAC-layer overheads during transmission, the frequency of ACK packets is reduced to once per ten received packets for both schemes. This effectively reduces the channel time utilization of the ACK streams from 10 - 30% (for link speeds of 6 - 54 Mbps) to less than 5 %.

5.2. Two Competing Streams

We first consider the simple case where *Harbor* is streamed over a fast link with nominal speed of 54 Mbps, while *Cyclists* is streamed over a slow link of 6 Mbps. Figure 4 shows the traces of the estimated link capacity, allocated video rate, corresponding video quality in PSNR and packet delivery delay of each stream, resulting from the cross-layer allocation in comparison with TFRC. At time 5.0 second, when the second stream becomes active, the cross-layer video rate controller for the first stream reacts by reducing its allocated rate. Similarly, when *Harbor* has finished transmission, rate and quality of *Cyclists* is updated to a higher level with cross-layer allocation. In both cases, the rate allocation from the cross-layer scheme converges within 0.5 - 1 s, whereas allocation from TFRC tends to change more slowly, with greater fluctuations in allocated rate and packet delay.

It is also interesting to note that during the period when both video streams are present, the optimal allocated rate of each stream also depends on its DR characteristics: higher rate is allocated to the more complex *Harbor* sequence, whereas the relatively static *Cyclists* sequence is streamed at a lower rate with higher quality. In comparison, the allocated rates from TFRC tend to be similar during the period when the slow link is competing with the fast links, as the observed round-trip-time for *Harbor* over the 54 Mbps link is increased due to the presence of *Cyclists*.

When the link speed of the second link is varied from 6 Mbps to 54 Mbps, the allocated video rates and resulting PSNR values comparing the cross-layer and TFRC schemes are plotted in Fig. 5, averaged over 120 seconds of simulation after the allocations have stabilized. Allocated rates to both streams increase with the speed of the slower link. The gap in the allocated rates is smaller from TFRC and diminishes when both links operate at the same speed. The rates from cross-layer allocation, in contrast, is media-aware, and favors the faster link supporting the more complex *Harbor* sequence. Consequently, the video quality gap between the two sequences remains around 8 dB for TFRC, yet is less than 4 dB for cross-layer allocation. Improvement in average video quality ranges from 1.0 dB with both links at 54 Mbps to 2.0 dB with the slower link at 6 Mbps.

5.3. Multiple Streams

In the next set of experiments, multiple video streams are transported over parallel links. All links except one have nominal speed of 54 Mbps, supporting *Harbor*. The slower link speed varies between 6 Mbps and 54 Mbps, supporting *Cyclists*. The allocation results in the case of three competing streams are shown in Fig. 6. Similar to the previous scenario of two video streams, the allocated rates and quality of the more difficult streams traversing faster links are higher from cross-layer allocation, at the cost of lower quality and rate for the less demanding stream over the slower link. The quality gap between the video sequences is reduced from 8.4 -9.3 dB with TFRC to 3.8 -5.6 dB with cross-layer allocation. Improvement in the average video quality ranges from 0.7 dB with all three links at 54 Mbps to

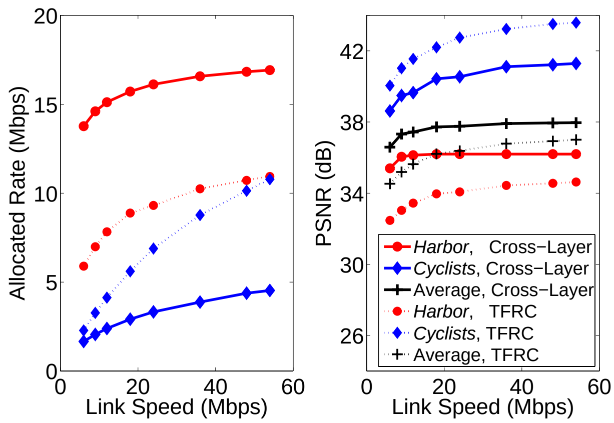

Fig. 5. Allocated video rate and quality in PSNR resulting from the TFRC and cross-layer allocation schemes. *Harbor* is streamed over the first link with nominal speed of 54 Mbps. The second link supporting *Cyclists* has a link speed varing from 6 Mbps to 54 Mbps. In all experiments, λ = 0.0001 in cross-layer allocation and k = 1.0 in TFRC. The results are averaged over 120 seconds after allocation has stabilized.

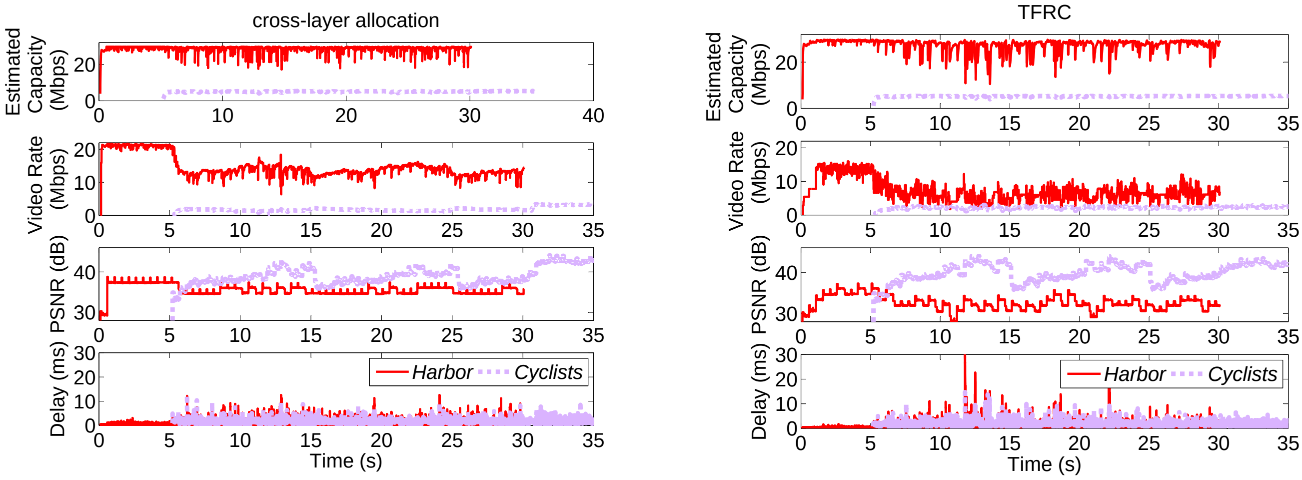

Fig. 4. Traces of link capacity, allocated video rate, resulting PSNR and packet delay when *Harbor* is streamed over a fast link with 54 Mbps nominal speed and *Cyclists* is streamed over a slow link of 6 Mbps. The cross-layer allocation results are plotted on the left with $\lambda = 0.00005$; results from TFRC are plotted on the right with $k = 1.0$.

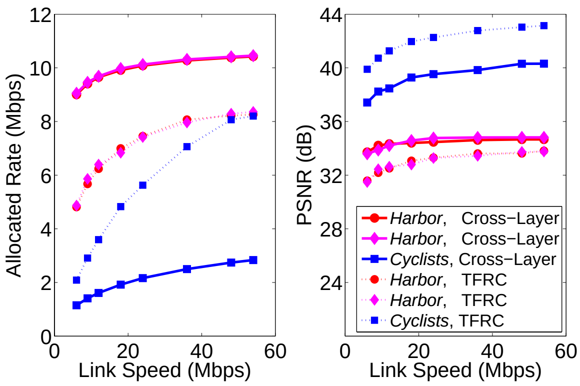

Fig. 6. Allocated video rate and PSNR resulting from the TFRC and cross-layer allocation schemes. *Harbor* is streamed over the first and second link, respectively, with nominal speed of 54 Mbps. The third link supporting *Cyclists* has a link speed varying from 6 Mbps to 54 Mbps. In all experiments, $\lambda = 0.0001$ in cross-layer allocation and $k = 1.0$ in TFRC. The results are averaged over 120 seconds after allocation has stabilized.

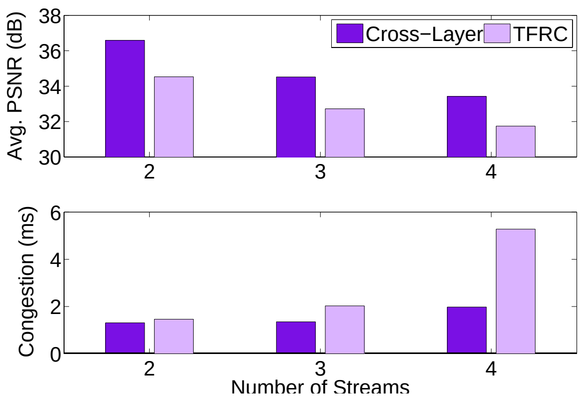

Fig. 7. Average video quality in PSNR and overall network congestion resulting from TFRC and cross-layer allocation, as the number of video streams changes from 2 to 4. The slowest link has nominal speed of 6 Mbps. In all experiments, $\lambda = 0.0001$ for cross-layer allocation and $k = 1.0$ for TFRC. The results are averaged over 120 seconds after allocation has stabilized.

1.8 dB with slowest link at 6 Mbps.

Figure 7 compares the two allocation schemes in terms of average video quality and overall network congestion, with nominal speed of the slowest link at 6 Mbps. The average video quality decreases for both schemes, as the number of competing streams increases within the same wireless network. Nevertheless, the cross-layer allocation achieves higher average video quality and lower network congestion, as compared to TFRC.

5.4. Multi-hop Network

We now consider the more challenging case where the video streams are sent over a multi-hop network as shown in Fig. 2 (b). Figure 8 (a) compares the cross-layer allocation against TFRC in terms of allocated video rate and resulting PSNR, as *Crew* is steamed over the 2-hop connection and *Harbor* is streamed over a 3-hop connection, with nominal speed of the link from Node 1 to 2 varying from 6 Mbps to 54 Mbps. Fig. 8 (b) shows the comparison for *Bigships* over the 2-hop connection and *Cyclists* over the 3-hop connection. In both cases, the cross-layer allocation is adapted according to the DR characteristics of both streams, therefore it results in higher average video quality over TFRC, which yields almost identical allocations regardless of media content. The performance gain in terms of average video quality in PSNR is 1.6 - 2.1 dB in the case of *Crew* and *Harbor*, and is 0.7 - 1.1 dB with *Bigships* and *Cyclists*.

6. CONCLUSIONS

We propose a distributed rate allocation scheme for multi-user video streaming over a shared wireless network. The trade-off between average video quality and overall network congestion is optimized, by

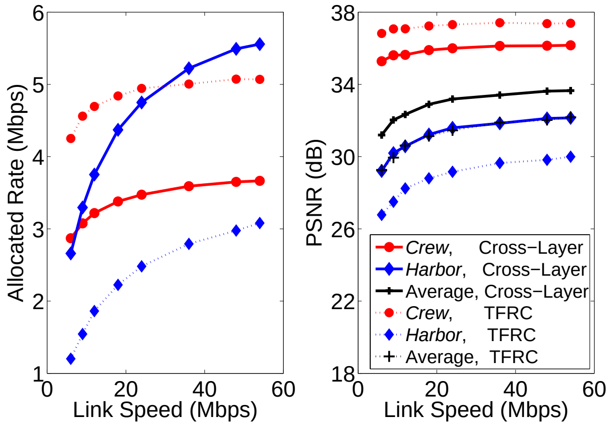

(a)

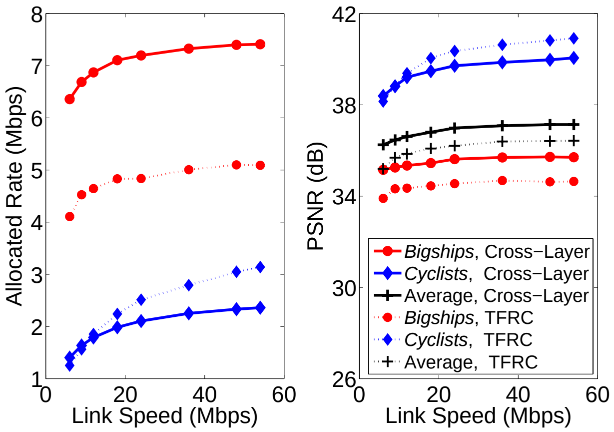

(b)

Fig. 8. Average video rate and PSNR resulting from the TFRC and cross-layer allocation schemes over the multi-hop network depicted in Fig. 2 (b). The nominal link speed of the first link (Node 1 to 2) varies from 6 Mbps to 54 Mbps. (a) *Crew* is streamed over the 2-hop path, *Harbor* is delivered over the 3-hop path. 54 Mbps. (b) *Bigships* is streamed over the 2-hop path, *Cyclists* is delivered over the 3-hop path. In all experiments, λ = 0.0001 in cross-layer allocation and k = 1.0 in TFRC. The results are averaged over 120 seconds after allocation has stabilized.

allowing cross-layer information exchange between the link state monitors at the MAC layer and the video rate controller at the application layer. Compared to TFRC, the proposed allocation scheme benefits from explicit knowledge of both the video DR characteristics and wireless link capacities. In the case where different video contents are streamed over links of different speeds, the proposed cross-layer allocation can achieve more balanced quality among the video streams, with higher average video quality and lower network congestion than TFRC. Performance gain ranges between 0.7 dB to 2.1 dB in terms of average video quality in PSNR, depending on the number of participating video streams, RD characteristic of each stream, as well as heterogeneity in the wireless link speeds.

7. REFERENCES

[1] V. Jacobson, "Congestion avoidence and control," *Proc. SIGCOMM'88*, vol. 18, no. 4, Aug. 1988.

[2] S. Floyd and K. Fall, "Promoting the use of end-to-end congestion control in the Internet," *IEEE/ACM Trans. on Networking*, vol. 7, pp. 458–472, Aug. 1999.

[3] F. Kelly, A. Maulloo, and D. Tan, "Rate control for communication networks: Shadow prices, proportional fairness and stability," *Journal of Operations Research Society*, vol. 49, no. 3, pp. 237–252, 1998.

[4] R. J. La and V. Anantharam, "Utility-based rate control in the Internet for elastic traffic," *IEEE Trans. on Networking*, vol. 10, no. 2, pp. 272–285, Oct. 2002.

[5] C. Chaudet, D. Dhoutaut, and I.G. Lassous, "Performance issues with IEEE 802.11 in ad hoc networking," *IEEE Communications Magazine*, vol. 43, no. 7, pp. 110–116, July 2005.

[6] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," *Proceedings INFOCOM'03, Twenty-Second Annual Joint Conference of the IEEE Computer and Communications Societies, San Francisco, U.S.A.*, vol. 2, pp. 836–843, Mar. 2003.

[7] M. Kalman and B. Girod, "Optimal channel-time allocation for the transmission of multiple video streams over a shared channel," *Proc. IEEE International Workshop on Multimedia Signal Processing (MMSP'05), Shanghai, China*, Oct. 2005.

[8] X. Zhu, P. van Beek, and B. Girod, "Distributed channel time allocation and rate adaptation for multi-user video streaming over wireless home networks," *IEEE International Conference on Image Processing (ICIP'07), (accepted)*, sept 2007.

[9] Cali, M. Conti, and E. Gregori, "IEEE 802.11 wireless LAN: Capacity analysis and protocol enhancement," *Proceedings INFOCOM'98, Seventeenth Annual Joint Conference of the IEEE Computer and Communications Societies, San Francisco, U.S.A.*, vol. 1, pp. 142–149, Mar. 1998.

[10] L. Kleinrock, *Queuing Systems, Volume II: Computer Applications*, Wiley Interscience, New York, USA, 1976.

[11] K. Stuhlmüller, N. Färber, M. Link, and B. Girod, "Analysis of video transmission over lossy channels," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 6, pp. 1012–32, June 2000.

[12] X. Zhu, E. Setton, and B. Girod, "Congestion-distortion optimized video transmission over ad hoc networks," *EURASIP Journal of Signal Processing: Image Communications*, vol. 20, no. 8, pp. 773–783, Sept. 2005.

[13] N. Z. Shor, "Minimization Methods for Non-differentiable Functions," *Springer Series in Computational Mathematics. Springer*, 1985.

[14] "NS-2," http://www.isi.edu/nsnam/ns/.

[15] *IEEE Standard for Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, P802.11*, Nov. 1997.

[16] ITU-T and ISO/IEC JTC 1, *Advanced Video Coding for Generic Audiovisual services, ITU-T Recommendation H.264 - ISO/IEC 14496-10(AVC)*, 2003.

[17] "x.264," http://developers.videolan.org/x264.html.