

# Distributed Media-Aware Rate Allocation for Video Multicast over Wireless Networks

Xiaoqing Zhu, Thomas Schierl, *Member, IEEE*, Thomas Wiegand, *Fellow, IEEE*,  
and Bernd Girod, *Fellow, IEEE*

**Abstract**—A unified optimization framework for rate allocation among multiple video multicast sessions sharing a wireless network is presented. Our framework applies to delivery of both scalable and non-scalable video streams. In both cases, the optimization objective is to minimize the total video distortion of all peers without incurring excessive network utilization. Our system model explicitly accounts for heterogeneity in wireless link capacities, traffic contention among neighboring links, as well as different video rate-distortion (RD) characteristics. The proposed distributed rate allocation scheme leverages cross-layer information exchange between the media access control and application layers to achieve fast convergence at the optimal, media-aware allocation. Performance of the proposed media-aware rate allocation protocol is compared against a heuristic scheme based on TCP-friendly rate control (TFRC). In network simulations of standard-definition video streaming over single or multiple multicast trees, the proposed scheme consistently achieves higher overall video quality than the TFRC-based heuristics. When delivering scalable streams, the flexibility of per-peer rate adaptation inside each multicast tree yields a further slight improvement in overall video quality over multicast of non-scalable streams.

**Index Terms**—Cross-layer design, distributed rate allocation, scalable video coding (SVC), video multicast, wireless 802.11 networks.

## I. INTRODUCTION

**A**N EFFICIENT way for simultaneously serving the same video content to a group of interested receivers is via multicast delivery. This can be achieved either at the network layer via Internet Protocol (IP) multicast, or at the application layer by logical relays [2]. With the advance of wireless networking technologies, wireless video multicast offers a viable

solution for many applications, e.g., broadband multimedia content sharing within a residential community.

Rate allocation for video multicast is necessary for accommodating heterogeneity in receiver capabilities, while avoiding network congestion. In wireless networks, this problem is further complicated by traffic contention among neighboring links and by heterogeneity in the link speeds. For instance, in IEEE 802.11 wireless networks with contention-based media access control (MAC) protocols, packets of the same size sent over a slow link would occupy the shared wireless channel for a longer duration than those sent over a fast link [3]. Moreover, the rate utility varies among video streams carrying different contents; the impact of the same rate increase on an action movie sequence may be rather different from that on a head-and-shoulder news clip. Therefore, when multiple video multicast sessions share the same wireless network, it is important for the rate allocation scheme to account for heterogeneity in both wireless links and video rate-distortion (RD) characteristics.

Our earlier work has studied a distributed media-aware rate allocation scheme for multiuser video streaming over wireless networks [4]–[6]. The scheme aims at minimizing the total video distortion of all participating streams without incurring excessive network utilization. The optimal solution is achieved in a distributed manner, by maintaining and exchanging congestion prices locally among neighboring wireless nodes at the MAC layer while adapting the video rate of each stream at the application layer. In this paper, we propose to extend the same set of design principles to video multicast over wireless networks. We focus on application-layer multicast, where pre-encoded video contents can be shared over wireless networks consisting of static nodes. In addition to the challenges faced by unicast streaming, rate allocation schemes for video multicast need to efficiently accommodate peers experiencing heterogeneous link qualities. It is also important to obtain periodic feedback from children peers for constant update of rate allocation decisions, while avoiding explosion of acknowledgement packets across the multicast delivery tree.

We address both considerations in our design of a unified rate allocation optimization framework. In particular, we introduce per-hop acknowledgements between each parent and its children, carrying reports of congestion prices as the main form of feedback. Our framework also accommodates scenarios of both non-scalable and scalable video multicast. When delivering non-scalable video streams, congestion prices are

Manuscript received September 25, 2009; revised August 11, 2010; accepted October 16, 2010. Date of publication May 2, 2011; date of current version September 2, 2011. An early version of this work was presented at the SPIE Visual Communications and Image Processing Conference, San Jose, CA, in January 2008 [1]. This paper was recommended by Associate Editor W. Gao.

X. Zhu is with the Advanced Architecture and Research Group, Cisco Systems, Inc., San Jose, CA 95134 USA (e-mail: zhuxq@stanfordalumni.org).

T. Schierl is with the Fraunhofer Institute for Telecommunications-Heinrich Hertz Institute, Berlin 10587, Germany (e-mail: thomas.schierl@hhi.fraunhofer.de).

T. Wiegand is with the Image Communications Laboratory, Berlin Institute of Technology, Berlin 10623, Germany, and the Fraunhofer Institute for Telecommunications-Heinrich Hertz Institute, Berlin 10587, Germany (e-mail: thomas.wiegand@tu-berlin.de).

B. Girod is with the Information Systems Laboratory, Stanford University, Stanford, CA 94305 USA (e-mail: bgirod@stanford.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCSVT.2011.2129690

accumulated recursively along each multicast tree before being sent back to the root peer. When delivering scalable video streams, congestion prices are accumulated around the neighborhood of each wireless link to facilitate graceful quality reduction at intermediate nodes *inside each multicast tree*. In both scenarios, each sending (root or relay) peer calculates the optimal video rate based on both the accumulated congestion price and the video RD characteristic. The final allocated video rates hence achieve the minimum total video distortion of all peers in all multicast sessions, without overloading the network.

The rest of this paper is organized as follows. Section II reviews related work in video multicast. Section III presents the wireless network model and the video RD model used in our problem formulations. Sections IV and V explain how to perform rate allocation for multicast delivery of scalable and non-scalable video streams, respectively. Section VI compares the performance of the proposed media-aware rate allocation protocol against a heuristic scheme based on TCP-friendly rate control (TFRC) in various network simulation scenarios.

## II. RELATED WORK

Regardless of whether video multicast is implemented at the network layer or at the application layer, the system needs to address the issue of receiver heterogeneity. For a single multicast session, rate allocation among the receivers is a well-studied optimization problem [7], [8]. Given non-scalable compressed media streams, the sender needs to adapt its outgoing video rate to accommodate the receiver with the lowest bandwidth, as in [9]. Alternatively, one can extend the single-stream adaptation approach to organize receivers into separate trees according to their bandwidth types, as described in the destination set group protocol [10]. Given a scalable representation of the compressed video stream, the approach of receiver-driven layered multicast (RLM) allows each receiver to adjust its subscription to video quality layers based on its available bandwidth [11]. The layered structure of the media stream is also well suited to dynamic optimization of the unequal error protection strength among the packets [12]–[14].

For rate allocation across different multicast sessions, most solutions typically aim at achieving some level of fairness among video sessions and TCP flows [15], [16]. The optimization framework in [17] investigates rate allocation for receivers with generic rate utilities over a network, without catering to the specific characteristics of video streams and wireless networks.

Recent studies have tackled the additional challenges arising from multicast over wireless networks. In [18] and [19], for instance, the RLM scheme is combined with distributed source coding and forward error correction techniques to achieve better error-resiliency. For more robustness against topological changes in a wireless network, various multipath routing schemes have been extended to construction of multiple multicast trees, in order to leverage the structured redundancy in multiple-description coded video streams [20]–[22]. Studies have also indicated the benefit of cross-layer design for a single wireless video multicast session, for instance via joint

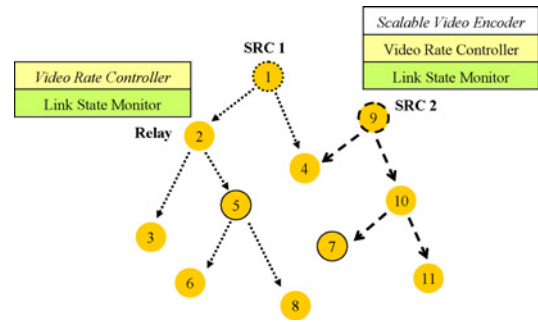


Fig. 1. Overview of a wireless video multicast system, where two video multicast sessions share a wireless network. Each video stream is delivered to all participating peers over an application-layer multicast tree. Acknowledgment packets are sent on a per-hop basis, from each receiver to its parent peer. Media-aware rate allocation is achieved via interactions between MAC-layer link state monitors (LSMs) and application-layer video rate controllers (VRCs). The modules containing italic fonts are only needed when delivering scalable video streams.

optimization of power allocation, multicast tree construction, and rate allocation [23], [24].

## III. SYSTEM MODEL

Fig. 1 presents an overview of a wireless video multicast system. The wireless nodes self-organize into a network and deliver video streams along application-layer multicast trees. In the following, we first introduce the wireless network model and video RD model for formulating the multicast rate allocation problem. We then describe the key components comprising the proposed media-aware multicast rate allocation protocol.

### A. Wireless Network Model

We index each wireless link with  $l$  and each video multicast session with  $s$ . The set of all links within a wireless network is denoted as  $\mathcal{L}$ , the set of all multicast sessions,  $\mathcal{S}$ . The application-layer multicast tree  $\mathcal{T}^s$  for Session  $s \in \mathcal{S}$  consists of all links traversed by that video stream. Link  $l$  *succeeds* Link  $l'$  in multicast tree  $\mathcal{T}^s$ ,  $l' \Rightarrow l$ , if the source of Link  $l$  is the destination of Link  $l'$ . In Fig. 1, for instance, Link  $5 \rightarrow 8$  succeeds Link  $2 \rightarrow 5$  in the first multicast tree.

We assume in this paper that transmission errors are effectively remedied by error control techniques such as adaptive modulation, channel coding, and retransmissions [25]. Consequently, the wireless channel exhibits time-varying throughput over each link, with no residual packet losses observed at the application layer [26], [27].<sup>1</sup> We define throughput  $C_l$  as the maximum achievable data rate over Link  $l$ , when the rest of the network is *not* transmitting. Total traffic rate over that link is  $F_l = F_l' + \sum_{s:l \in \mathcal{T}^s} R_l^s$ , including both the rate of non-video traffic  $F_l'$  and the allocated rates  $R_l^s$ 's of all video sessions traversing that link.

In wireless networks with contention-based MAC protocols such as IEEE 802.11 [28], nearby links  $l$  and  $l'$  typically

<sup>1</sup>While this assumption is valid for wireless networks with static nodes in our setting, it may not hold in networks with mobile nodes and frequent topological changes. We leave the exploration of mobile networks as future research.

cannot transmit at the same time, and are deemed to *interfere with* each other. Interference among links is denoted as  $l \bowtie l'$ , and is reciprocal. In other words,  $l' \bowtie l$  if and only if  $l \bowtie l'$ . For Link  $l_{i \rightarrow j}$ , we define the set of links whose source or destination node is within transmission range of either Node  $i$  or Node  $j$  as its *interference set*  $\mathcal{L}_l$ . By definition,  $l' \bowtie l, \forall l' \in \mathcal{L}_l$ .

We define link utilization as the fraction of time during which each link is active:  $u_l = F_l/C_l$ . It is obvious that the same traffic rate incurs higher utilization over a slower link than over a faster link. Consequently, total utilization within each interference set  $\mathcal{L}_l$  is constrained by

$$\tilde{u}_l = \sum_{l' \in \mathcal{L}_l} u_{l'} < \gamma \quad (1)$$

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 traffic contention over the shared wireless media, or inaccurate estimates of instantaneous link throughput.

### B. Video RD Model

The RD tradeoff of each video stream is described using a parametric model [29] as follows:

$$D^s(R) = \frac{\theta^s}{R - R_0^s} + D_0^s \quad (2)$$

where the parameters  $D_0^s$ ,  $R_0^s$ , and  $\theta^s$  can be fitted to empirical RD data points using nonlinear regression techniques. They are updated for every group of pictures (GOP) in the pre-encoded video stream.

For non-scalable video streams, the video rate over each multicast tree is constrained to the rate chosen by the root peer:  $R_l^s = R^s, \forall s \in \mathcal{S}$ . Rate adaptation is achieved by means of bitstream switching among multiple pre-encoded versions of the video sequence.

For scalable video streams, the video rate can be determined separately for each peer inside each multicast tree, so long as the rate of a child does not exceed that of its parent:  $R_l^s \leq R_{l'}^s$  for  $l' \Rightarrow l$ . Rate adaptation can be performed at the root and at each relaying peer, by means of sending or forwarding subsets of enhancement-layer video packets, a mechanism to be explained in greater detail in Section V.

## IV. RATE ALLOCATION FOR NON-SCALABLE STREAMS

### A. Optimization Objective

The goal of multicast rate allocation is to maximize the overall viewing experience of all participating peers without incurring excessive network utilization. Based on the subjective viewing test results described in [30], we choose to minimize the weighted sum of mean-squared error (MSE) distortion of all streams as follows:

$$\min \sum_{s \in \mathcal{S}} w^s N^s D^s(R^s) \quad (3)$$

$$\text{s.t.} \quad R^s > R_{min}^s, \quad \forall s \in \mathcal{S} \quad (4)$$

$$R^s < R_{max}^s, \quad \forall s \in \mathcal{S} \quad (5)$$

$$\tilde{u}_l < \gamma, \quad \forall l \in \mathcal{L}. \quad (6)$$

In (3), the contribution from Stream  $s$  is weighted by both the number of peers in its multicast tree  $N^s$  and a user-specified weighting factor  $w^s$ .<sup>2</sup> Presence of the user-specified weighting factors allows greater flexibility in our scheme to accommodate different user preferences or priorities. For instance, a preferred user is able to obtain higher rate and quality for a rather simple video sequence by assigning a higher weighting factor to its multicast stream. In addition, (4) and (5) bound the allocated rate within the range provided by available encoded bitstreams. Constraint (6) guards against excessive total utilization within each interference set.

One can easily verify that the optimization in (3)–(6) has a convex objective function with linear constraints. We only consider the case when the problem is feasible, and assume that the case where streams at minimum rates still violate the total channel utilization constraint is handled separately by some access control mechanism. If all link states and all video RD parameters were available at a central entity, the optimal solution could be calculated using standard numerical techniques such as the interior point method [31]. However, the overhead in collecting such global information may not scale well with growing network size or stream density. In practice, therefore, a distributed solution is preferable.

### B. Distributed Solution

We now show how the pricing-based rate control algorithm proposed in [32] can be extended to solve the multicast rate allocation problem in (3)–(6). The original optimization objective of maximizing generic logarithmic utility functions is substituted with minimizing video-specific parametric RD functions. In addition, the resource constraints now correspond to total channel utilization within each interference set, instead of total rate over individual links in a wired network. The algorithm can be decomposed into the following two iterative steps.

1) *Congestion Price Update*: A non-negative congestion price  $\lambda_l$  is associated with each interference set  $\mathcal{L}_l$ . Its value is updated periodically, according to instantaneous residual rate over Link  $l$  as follows:

$$\lambda_l(t) = \max[\lambda_l(t - \tau) + \kappa(\tilde{u}_l - \gamma)C_l, 0]. \quad (7)$$

In (7),  $\tau$  indicates the price update interval and  $\kappa$  is a scaling factor controlling the update step sizes.<sup>3</sup>

For  $\lambda_l > 0$ , the price update is proportional to instantaneous *excess* total utilization over  $\mathcal{L}_l$  and link throughput  $C_l$ . The intuition behind this is that  $\lambda_l$  increases if total channel time utilization  $\tilde{u}_l$  temporarily exceeds the specified limit  $\gamma$  in order to induce rate reduction by all streams affecting  $\mathcal{L}_l$ . Conversely, as long as  $\tilde{u}_l$  is below the target  $\gamma$ , the corresponding congestion price should keep decreasing to encourage higher rates from all contributing streams.

<sup>2</sup>Note that the the proposed optimization objective tends to allocate lower rates for streams traversing multiple congested interferences sets, similar as TCP's bias against streams experiencing long round-trip-times.

<sup>3</sup>It will become clear later in (8) and (9) that the unit of the congestion price  $\lambda$  should be MSE distortion divided by rate. Consequently, the unit of  $\kappa$  should be MSE distortion divided by rate squared. From now on, we will denote the unit of  $\lambda$  with MSE/(Mb/s), and that of  $\kappa$  with MSE/(Mb/s)<sup>2</sup>.

2) *Video Rate Update*: The rate of each multicast tree is updated as follows:

$$R^s = \begin{cases} R_{opt}^s = R_0^s + \sqrt{\frac{w^s N^s \theta^s}{\Lambda^s}} & R_{min}^s \leq R_{opt}^s \leq R_{max}^s \\ R_{min}^s & R_{opt}^s < R_{min}^s \\ R_{max}^s & R_{opt}^s > R_{max}^s \end{cases} \quad (8)$$

where  $D^s(R)$  follows the parametric video RD model (2) and

$$\Lambda^s = \sum_{l \in \mathcal{T}^s} \frac{\sum_{l' \in \mathcal{L}_l} \lambda_{l'} C_{l'}}{C_l}. \quad (9)$$

The value of  $\Lambda^s$  summarizes the congestion contribution of all wireless links affected by the multicast tree. In the special case where the multicast tree consists of only a single link  $l$ ,  $\Lambda^s$  reduces to  $\lambda_l$ . In general, the contribution from each link in the multicast tree  $\lambda_l$  is inflated by a factor of  $1 + \sum_{l' \neq l, l' \in \mathcal{L}_l} C_{l'} / C_l$ . Note that the value of  $\Lambda^s$  can be calculated in a recursive manner, whereby each node reports to its parent the accumulated congestion price  $\Lambda_l^s$  of its own subtree as follows:

$$\Lambda_l^s = \begin{cases} \frac{\sum_{l' \in \mathcal{L}_l} \lambda_{l'} C_{l'}}{C_l} + \sum_{l': l \Rightarrow l'} \Lambda_{l'}^s, & N_l^s > 0 \\ \frac{\sum_{l' \in \mathcal{L}_l} \lambda_{l'} C_{l'}}{C_l}, & N_l^s = 0. \end{cases} \quad (10)$$

Here  $N_l^s$  indicate the number of children receiving Stream  $s$  from the peer at the destination of Link  $l$ .

As illustrated in Fig. 2, the accumulated congestion price  $\Lambda^s$  determines the slope of the linear term counterbalancing the RD tradeoff. Intuitively, a higher video rate is encouraged when the accumulated congestion price  $\Lambda^s$  is lower. When the network becomes congested, increase in congestion prices at bottleneck interference sets leads to a higher accumulated price  $\Lambda^s$  and subsequently lower allocated video rates. In addition, the allocated video rate is confined within the range of  $[R_{min}^s, R_{max}^s]$ , as dictated by (4) and (5).

The iteration between (7) and (8) naturally constitutes a distributed algorithm: the congestion price update only requires local observation of  $\tilde{u}_l$  while the video rate update depends only on the end-to-end accumulated price  $\Lambda^s$  and the video RD parameters of individual streams. At each iteration step, the computational complexity for both congestion price update and video rate update stays constant for each peer, irrespective of the number of multicast sessions involved or network size.

At equilibrium, the congestion prices satisfy as follows:

$$\begin{cases} \lambda_l > 0, & \tilde{u}_l = \gamma, \text{ or} \\ \lambda_l = 0, & \tilde{u}_l < \gamma. \end{cases} \quad (11)$$

In other words, only congestion prices of fully utilized interference sets have strictly positive values. This corresponds to the Karush-Kuhn-Tucker conditions for (6) in the original optimization problem, thereby guaranteeing optimality of the solution [33].

### C. Rate Allocation Protocol

The proposed media-aware multicast rate allocation protocol comprises two key components, as shown in Fig. 1: LSMs at the MAC layer and VRCs at the application layer. The LSM at each wireless node maintains updated congestion price




Fig. 2. Allocated video rate  $R_{opt}^s$  is determined by the accumulated congestion price  $\Lambda^s$ , by the weight of importance  $w^s$ , and by the video RD function  $D^s(R)$ .

information based on local observations of link throughput and traffic rates. The VRC at the root of each multicast tree is in charge of video rate update, whereas each intermediate node mainly generates per-hop acknowledgement packets and caches accumulated congestion prices for its subtree. Cross-layer information exchange between the LSMs and VRCs is achieved by granting intermediate wireless nodes access to a set of special video packet header fields, as listed in Table I.<sup>4</sup>

The LSM of each wireless node continuously tracks the throughput and background traffic rates of its outgoing links. Upon relaying of a video packet, it is allowed to read out the advertised video rate  $R^s$  from the special header of traversing video packets. The local congestion price  $\lambda_l$  is periodically updated based on the total channel time utilization within local interference sets. The LSM also periodically exchanges link state messages with neighboring nodes, containing entries listed in Table II.

For delivery of non-scalable video streams, the VRC is needed only at the root of each multicast tree. The convergence criterion for the rate allocation process is that the calculated rate fluctuation between consecutive allocations is smaller than the rate difference between adjacent available rate points. At the same time, the received accumulated congestion price should not exhibit an increasing or decreasing trend. When the allocation has converged to the optimal value  $R^s$ , the VRC chooses the quality level  $k$  of the next GOP such that  $R_k^s \leq R^s < R_{k+1}^s$ ,  $1 \leq k \leq K$ . The actual bitstream switch occurs later, when the first frame in the next GOP is transmitted. Meanwhile, each VRC at an intermediate node keeps a local cache of congestion prices reported by ACK packet headers from all its children. It then reports the sum according to (10) as the accumulated congestion price of its *subtree* in the ACK packet to its parent.

Note from Table I that the size of the video packet header is constant, and can be easily accommodated by 5 bytes with sufficient accuracy. On average, the protocol overhead introduced by additional fields in the video packet header and ACK streams increases on the order of  $O(\bar{S} \cdot \bar{P})$ , where  $\bar{S}$  is the average number of multicast sessions traversing each interference set and  $\bar{P}$  is the average number of links traversed by each multicast stream in each interference set. While more frequent transmission of ACK packets facilitates faster convergence, it also increases the protocol overhead. In this paper, we choose the acknowledgement frequency to be once

<sup>4</sup>This paper focuses on the conceptual design of the rate allocation protocol. In a real implementation, it is possible to map these special fields as extensions of existing transport protocol headers, for instance, as supported by RTP [34] header extensions.

TABLE I  
FIELDS FOR CROSS-LAYER INFORMATION EXCHANGE IN THE VIDEO  
PACKET HEADER

Symbol	Content	Size	Range
$s$	Tree/Session ID	1 byte	0–255
$R^s$	Advertised rate	2 bytes	$0-65535 \times 10^{-3}$ Mb/s
$\Lambda^s$	Congestion price	2 bytes	$0-65535 \times 10^{-2}$ MSE/(Mb/s)

TABLE II  
ENTRIES IN THE LINK STATE MESSAGE

Symbol	Content	Size	Range
$l$	Link ID	1 byte	0–255
$C_l$	Link throughput	2 bytes	$0-65535 \times 10^{-3}$ Mb/s
$u_l$	Link utilization	1 byte	0–255%
$\lambda_l$	Congestion price	2 bytes	$0-65535 \times 10^{-3}$ MSE/(Mb/s) <sup>2</sup>

per video frame for a good balance between protocol overhead and convergence speed.

According to Table II, each entry in the link state message requires 6 bytes to ensure sufficient accuracy. Size of the link state message is approximately proportional to the number of entries in the update, hence proportional to the number of links the node is incident with. On average, the overhead increases on the order of  $O(\bar{N} \cdot \bar{L})$ , where  $\bar{N}$  is average number of nodes within each interference set, and  $\bar{L}$  is average number of links incident to each node. The frequency of the link state message exchanges determines the tradeoff between protocol overhead and its responsiveness. More frequent message exchanges incur less delay in keeping the link utilization information consistent among all neighboring nodes, at the cost of higher channel time utilization.

## V. RATE ALLOCATION FOR SCALABLE STREAMS

In a network containing wireless links with heterogeneous speeds, it is preferable to allow rate adaptation at intermediate relay peers *inside each multicast tree*. Otherwise the received video qualities would be limited by the peer experiencing the slowest link. We adopt the scalable video coding (SVC) extension of the H.264/MPEG-4 AVC standard for this purpose [35], [36].

### A. Rate Adaptation with SVC

Fig. 3 illustrates the structure of an SVC stream with a GOP length of four frames. The frames labeled T0 form the first temporal layer; the frames labeled T1 and T2 form the second and the third temporal layers, respectively. Each frame is encoded into one base-layer picture and one enhancement-layer picture. Note that the base-layer pictures labeled with either T1 or T2 are predicted from the enhancement-layer pictures of neighboring frames. This approach, also known as medium granularity scalability (MGS), provides a balance between high coding efficiency of the base-layer pictures and mismatch error control when the reference enhancement-layer pictures are missing [36].




Fig. 3. Structure of a GOP in an SVC stream encoded with MGS-SNR and temporal scalability.

For a stream with  $M$  temporal layers,  $M + 1$  rate points can be attained by sequentially dropping enhancement-layer pictures based on their temporal priorities. The first rate point corresponds to streaming base-layer pictures only for all frames, the second rate point includes also enhancement-layer pictures belonging to the first temporal layer, and so on. The highest rate point includes both base-layer and enhancement-layer pictures for all frames. The RD tradeoff achieved by this approach can be approximated by the parametric model (2).

### B. Optimization Objective

Similar as in the case of multicast of non-scalable video streams, the goal of the rate allocation scheme is to minimize the total video distortion of all participating peers without incurring excessive network utilization as follows:

$$\min \sum_{s \in \mathcal{S}} \sum_{l \in \mathcal{T}^s} w_l^s D^s(R_l^s) \quad (12)$$

$$\text{s.t.} \quad R^s > R_{min}^s, \quad \forall s \in \mathcal{S} \quad (13)$$

$$R^s < R_{max}^s, \quad \forall s \in \mathcal{S} \quad (14)$$

$$\bar{u}_l < \gamma, \quad \forall l \in \mathcal{L} \quad (15)$$

$$R_l^s \leq R_{l'}^s, \quad l' \Rightarrow l, \quad \forall l \in \mathcal{L}, \forall s \in \mathcal{S}. \quad (16)$$

In (12), the user-specified weighting factor is denoted by  $w_l^s$  for each peer in each session, and allows greater flexibility for the scheme to accommodate users with different preferences or priorities. Constraints (13) and (14) bound the allocated rate within the range of each scalable video stream. Constraint (15) guards against excessive total utilization within each interference set. Constraint (16) states that the allocated rate for a peer cannot exceed the rate of its parent.

### C. Distributed Solution

The optimal distributed solution can, again, be achieved by iteratively updating local congestion prices while adapting the rate of each peer as follows:

$$\lambda_l(t) = \max [\lambda_l(t - \tau) + \kappa(\bar{u}_l - \gamma)C_l, 0] \quad (17)$$

$$R_l^s = \begin{cases} R_{l,opt}^s = R_0^s + \sqrt{\frac{w_l^s \theta^s}{\Lambda_l}}, & R_{l,min}^s \leq R_{l,opt}^s \leq R_{l,max}^s \\ R_{l,min}^s, & R_{l,opt}^s < R_{l,min}^s \\ R_{l,max}^s, & R_{l,opt}^s > R_{l,max}^s \end{cases} \quad (18)$$

where the accumulated congestion price  $\Lambda_l$  accounts for contributions for all links within the interference set of Link  $l$  as follows:

$$\Lambda_l = \frac{\sum_{l' \in \mathcal{L}_l} \lambda_{l'} C_{l'}}{C_l}. \quad (19)$$



In (17), the price update interval is  $\tau$  and the price update scaling factor is  $\kappa$ . Following the same intuition as in (7), the congestion price update is proportional to instantaneous excess total utilization over  $\mathcal{L}_l$  and link throughput  $C_l$ . A temporarily over congested interference set leads to increased value of  $\lambda_l$ , whereas temporarily under-utilized interference set leads to decreased value of  $\lambda_l$ . In (18), a higher video rate is encouraged if the observed accumulated congestion price  $\Lambda_l$  is low. Conversely, increase in congestion price  $\Lambda_l$  leads to reduced video rate for all streams affecting the interference set  $\mathcal{L}_l$ . The allocated rate is further bounded by the minimum and the maximum available rates of each stream, and by the rate allocated to the parent peer:  $R_{l,min}^s = R_{min}^s$ ,  $R_{l,max}^s = \min[R_{max}^s, R_l^s]$ , for  $l' \Rightarrow l$ .

As in the case of non-scalable video multicast, the proposed distributed algorithm requires simple calculations for congestion price update and video rate update. At each iteration step, the computational complexity at each peer stays constant irrespective of the number of participating multicast sessions or network size.

#### D. Rate Allocation Protocol

For delivery of scalable video streams, the LSMs follow exactly the same procedures as in the case of multicast of non-scalable video streams. On the other hand, video rates can be adapted on a per-hop basis along the multicast distribution tree. At the beginning of each GOP, each VRC informs subsequent children peers of the updated RD parameters of the video stream via preamble packets. It records the advertised rate of its parent upon receipt of a video packet and extracts the value of accumulated congestion price  $\Lambda_l$  upon receipt of an ACK packet from one of its children. The optimal rate  $R_l^s$  is then calculated according to (18) for each child peer. Same as in non-scalable video multicast, the convergence criterion is that fluctuation in the calculated rate between consecutive allocations is smaller than the rate difference between adjacent available rate points without exhibiting an increasing or decreasing trend. Upon convergence of allocation, the VRC then forward video packets selectively to match the allocated rate.

Similar to multicast of non-scalable video streams, overhead of the rate allocation scheme introduced by the additional video packet header scales linearly with the average number of multicast trees and network size, on the order of  $O(\bar{S} \cdot \bar{P})$ . The overhead introduced by link state message exchanges scales with network density and network size, on the order of  $O(\bar{N} \cdot \bar{L})$ .

## VI. SIMULATION RESULTS

### A. Simulation Setup

We evaluate performance of the proposed media-aware multicast rate allocation protocol in ns-2 [39] simulations over a small wireless network shown in Fig. 1. The transmission power and receiving threshold of the nodes are adjusted to achieve a transmission range of 55 m. Parameters of SIFS/DIFS/EIFS slot time, random backoff window size, and retry limits are chosen according to specifications of the IEEE




Fig. 4. Rate-peak signal-to-noise ratio (PSNR) tradeoff curves of the six test video sequences: *City*, *Crew*, *Downtown*, *Harbor*, *Ice*, and *Terrace*. The sequences have spatial resolutions of  $704 \times 576$  pixels and temporal frequencies of 30 f/s. They are encoded into (a) non-scalable streams by  $\times 264$  [37] and (b) scalable streams by SVC JSVM 8.8 [38]. Note that the available rates for the non-scalable streams go beyond the range of the graph shown in (a).

802.11a standard [28]. The basic rate for header and control packet transmissions is set to 6 Mb/s, whereas the nominal link speed for payload transmissions varies between 6 and 54 Mb/s.<sup>5</sup>

Six standard-definition (SD) video sequences are considered for streaming. The sequences have spatial resolutions of  $704 \times 576$  pixels and temporal frequencies of 30 f/s.

For delivery of non-scalable streams, the video sequences are encoded using  $\times 264$  [37], a fast implementation of the H.264/AVC standard [40]. The GOP length is 30 frames; the GOP structure is IBBPBBP..., similar to that used in MPEG-2 streams. Different RD tradeoff points are obtained by varying the quantization parameter (QP) from 20 to 49.

For delivery of scalable streams, the test sequences are encoded into scalable video streams using the SVC JSVM 8.8 reference software [38]. Each frame is encoded into one base-layer picture with QP of 36 and one enhancement-layer picture with QP of 30. The GOP length is 32 frames, corresponding to six temporal layers and seven RD tradeoff points.

Fig. 4 shows the rate-PSNR tradeoff of all six sequences for both non-scalable and scalable representations. Note that the adoption of hierarchical B frames in the scalable streams yields slightly higher coding efficiency than the non-scalable streams. Encoded video frames (or base-layer and enhancement-layer pictures) are further segmented into network packets with maximum size of 1500 bytes. Packet transmission intervals are evenly spread out within each GOP to avoid unnecessary queuing delays due to intra-coded frames.

All simulations presented in this paper adhere to the same set of protocol parameter choices: the target utilization is  $\gamma = 85\%$ , the price update scaling factor is  $\kappa = 1.0 \text{ MSE}/\text{Mb/s}^2$ , and the price update interval is  $\tau = 10 \text{ ms}$ . To balance between fast convergence and protocol overhead, the video acknowledgment frequency is set as one ACK per frame and the LSM message exchange interval is chosen as 20 ms. The weight of importance is set as unity for all peers in all multicast trees. The playout deadline for all video streams is chosen as 1 s.

### B. TFRC-Based Heuristic Scheme

Performance of the media-aware rate allocation protocol is compared against a TFRC-based heuristic scheme. We

<sup>5</sup>Note that the notion of link speed does not factor in MAC-layer overheads such as random backoff and guard time slots between adjacent transmissions, therefore is typically significantly higher than the effective link throughput observed by the video stream.

further enhance TFRC by introducing a *virtual* random-early-detection (RED) mechanism at the relay nodes. Such mechanism helps to avoid quality degradation caused by actual packet losses due to queue overflow. The relay node monitors its queue size, and calculates the probability for randomly marking a packet as *virtual loss* according to the same principles in RED queues [41]. Such virtual loss is marked in the the 1-bit explicit-congestion-notification field in the IP packet header [42]. The receiver calculates the average percentage of marked packets as *virtual* packet loss ratio  $p$  and reports this information back to the sender via ACK packets.

For delivery of non-scalable video streams, each relay peer averages the measured round trip times  $T_l^s$  and virtual packet loss ratios  $p_l^s$  reported by all its children, before feeding back such information to its parent as follows:

$$\tilde{p}_l^s = \begin{cases} \sum_{l':l \Rightarrow l'} \tilde{p}_{l'}^s, & N_l^s > 0 \\ p_l^s, & N_l^s = 0 \end{cases} \quad (20)$$

$$\tilde{T}_l^s = \begin{cases} \sum_{l':l \Rightarrow l'} \tilde{T}_{l'}^s, & N_l^s > 0 \\ T_l^s, & N_l^s = 0 \end{cases} \quad (21)$$

where  $N_l^s$  indicates the number of children receiving Stream  $s$  from the peer at the destination of Link  $l$ . The allocated rate is then calculated as a function of average end-to-end round-trip time  $\tilde{T}^s$  and virtual packet loss ratio  $\tilde{p}^s$  over the entire multicast tree [43] as follows:

$$R^s = k \frac{B}{\tilde{T}^s \sqrt{\tilde{p}^s}} \quad (22)$$

with scaling factor  $k$  and average packet size  $B$ .

For delivery of scalable video streams, the allocated rate is calculated on a per-hop basis as a function of the measured round trip time  $T_l^s$  and virtual packet loss ratio  $p_l^s$  over the connection between each peer and its parent as follows:

$$R_l^s = k \frac{B}{T_l^s \sqrt{p_l^s}}. \quad (23)$$

The rate is further constrained by the rate of the incoming stream, i.e., the rate allocated to the parent peer, to ensure that  $R_l^s \leq R_{l'}^s$  for  $l' \Rightarrow l$ .

### C. Allocation for Non-Scalable Streams

1) *Single Multicast Tree*: We first consider the simple case of a single multicast session over the first seven-node multicast tree shown in Fig. 1. All links are within the same interference set. Fig. 5 shows the traces of estimated link capacities, accumulated congestion price  $\Lambda^s$  over the entire multicast tree, allocated video rates, corresponding video qualities in PSNR, and end-to-end packet delivery delay experienced by Peers 2 and 6. In this experiment, the outgoing link speed of Peer 5 initially starts as 54 Mb/s, then drops to 12 Mb/s between time  $t = 10$  s to  $t = 30$  s. All other links operate at 54 Mb/s. It can be observed that the congestion prices and the allocated video rates quickly reach convergence both before and after the link speed drop at Peer 5. During the period of low link speed at Peer 5, the allocated rate and resulting video quality of the entire multicast tree are reduced accordingly, as a consequence of the increased accumulated congestion price. Since the proposed scheme mindfully avoids network




Fig. 5. Traces of estimated link capacities, accumulated congestion prices, allocated video rates, corresponding video qualities, and packet delivery delay. The *City* sequence streams over the first multicast tree shown in Fig. 1. The outgoing link speed of Peer 5 initially starts as 54 Mb/s, and drops to 12 Mb/s between time  $t = 10$  s to  $t = 30$  s. All other links operate at 54 Mb/s.




Fig. 6. Average video quality measured as PSNR of the average video distortion of all peers. One sequence streams over the first multicast tree shown in Fig. 1. The outgoing link speed of Peer 5 varies from 6 Mb/s to 54 Mb/s while all other links operate at 54 Mb/s. (a) *City*. (b) *Crew*. (c) *Downtown*. (d) *Harbor*. (e) *Ice*. (f) *Terrace*.

congestion by limiting the utilization within each interference set, it also leads to low packet delivery delays, well below the latency requirement of 1 s.

Fig. 6 compares the media-aware allocation against the TFRC-based heuristic scheme in terms of average video quality, measured as PSNR of the average video distortion of all peers over the entire multicast tree. As the outgoing link speed of Peer 5 varies from 6 Mb/s to 54 Mb/s, both schemes achieve similar results in that they both try to maximize the allocated rate without overloading the network. Since the TFRC-based heuristic allocation tends to fluctuate over time, it results in a slightly lower average video quality than the proposed media-aware scheme.

2) *Multiple Multicast Trees*: Next, we consider the scenario of two video sequences streaming over two multicast trees, respectively, as depicted in Fig. 1. The outgoing link speed of Peer 5 varies between 6 Mb/s and 54 Mb/s while all other links operate at 54 Mb/s. Fig. 7 compares the media-aware allocation and the TFRC-based heuristics scheme for three video sequence pairs. We compare the allocated video rate and quality of *each* multicast tree, as well as the average video quality for all peers. Note that the media-aware allocation outperforms the TFRC-based heuristic scheme in terms of the overall video quality, measured as PSNR of the average video distortion of all peers in *both* trees. The performance

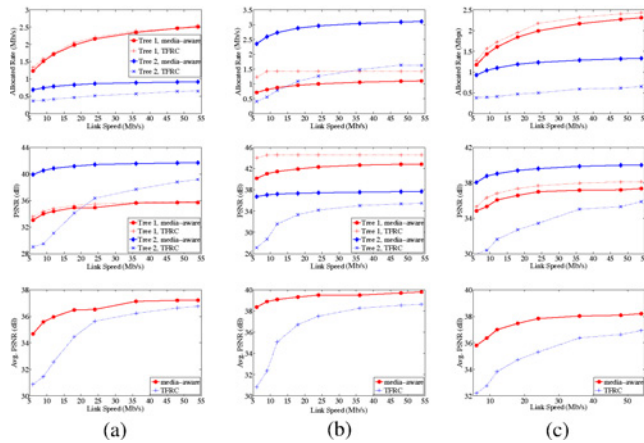


Fig. 7. Allocated video rate to each multicast tree, average video quality of each multicast tree, and average video quality measured as PSNR of the average video distortion of all peers. Two sequences stream over two multicast trees as shown in Fig. 1. The outgoing link speed of Peer 5 varies from 6 Mb/s to 54 Mb/s while all other links operate at 54 Mb/s. (a) *City* versus *Downtown*. (b) *Downtown* versus *Crew*. (c) *Terrace* versus *Ice*.

gain is more pronounced when there is greater heterogeneity in the wireless link speeds. When the outgoing link speed of Peer 5 is 6 Mb/s, for instance, the improvement in overall video quality ranges between 3.3 dB and 7.2 dB in PSNR.

#### D. Allocation for Scalable Streams

1) *Single Multicast Tree*: When delivering a scalable video streams, rate allocation for peers in the same tree can be different to accommodate heterogeneity in their link speeds. This is reflected in the traces in Fig. 8. In this experiment, the outgoing link speed of Peer 5 initially starts as 54 Mb/s, then drops to 12 Mb/s between time  $t = 10$  s to  $t = 30$  s. All other links operate at 54 Mb/s. It can be noted that both congestion prices and allocated video rates rapidly converge within several seconds, both before and after the drop in link speed at Peer 5. During the period when its link speed is at 12 Mb/s, the throughput of Links  $5 \rightarrow 6$  and  $5 \rightarrow 8$  are reduced accordingly. Consequently, Peer 5 receives higher accumulated congestion prices for both its children, hence the allocated rates for Peers 6 and 8 are lower than allocations to other peers in the tree. Similar as in the case of non-scalable video multicast, the proposed scheme also benefits from proactive avoidance of congestion, hence is able to maintain low packet delivery delays, well below the latency requirement of 1 s.

Fig. 9 compares the media-aware allocation against the TFRC-based heuristic scheme in terms of the video quality received by each peer. As the outgoing link speed of Peer 5 varies from 6 Mb/s to 54 Mb/s, the media-aware scheme increases its allocation in accordance with increase in the wireless link speed, consistently allocating higher rates to peers receiving over faster links. In contrast, the TFRC-based heuristics tends to allocate similar rates to all peers within the network. When the two outgoing links of Peer 5 operate at the lowest speed, the network only has sufficient resource to accommodate the *Harbor* sequence at its lowest rate, therefore results from both schemes coincide. On the contrary, when all links are operating at the highest speed,




Fig. 8. Traces of estimated link capacities, accumulated congestion prices, allocated video rates, corresponding video qualities, and packet delivery delay. The scalably-encoded *Harbor* sequence streams over the first multicast tree shown in Fig. 1. The outgoing link speed of Peer 5 initially starts as 54 Mb/s, and drops to 12 Mb/s between time  $t = 10$  s and  $t = 30$  s. All other links operate at 54 Mb/s.




Fig. 9. Video quality of each peer when the scalably encoded *Harbor* sequence streams over the first multicast tree in Fig. 1. The outgoing link speed of Peer 5 varies from 6 Mb/s to 54 Mb/s; all other links operate at 54 Mb/s. (a) Peer 2. (b) Peer 4. (c) Peer 3. (d) Peer 5. (e) Peer 6. (f) Peer 8.

fluctuations in the observed end-to-end round trip times lead to network under-utilization for TFRC-based heuristics. The media-aware scheme therefore achieves higher overall video quality, measured in PSNR of the average video distortion of all peers. The performance gain ranges between 0.1 dB and 1.9 dB in PSNR for various video sequences and link speeds, as shown in Fig. 10. Note also that the video quality from media-aware allocation for most sequences tends to saturate at the maximum level of the scalable stream, as the outgoing link speed of Peer 5 increases to 54 Mb/s. The only exception is *Harbor*, which carries the most demanding RD characteristics and a higher maximum rate than what the network can support with all links operating at 54 Mb/s.

2) *Multiple Multicast Trees*: Similar observations hold for two video sequences over two multicast trees. Fig. 11 compares the overall video quality for all peers in both multicast trees as achieved by the media-aware scheme and






Fig. 10. Average video quality measured as PSNR of the average video distortion of all peers. One scalably encoded video sequence streams over the first multicast tree shown in Fig. 1. The outgoing link speed of Peer 5 varies from 6Mb/s to 54Mb/s, while all other links operate at 54Mb/s. (a) *City*. (b) *Crew*. (c) *Downtown*. (d) *Harbor*. (e) *Ice*. (f) *Terrace*.




Fig. 11. Average video quality measured in PSNR of the average video distortion of all peers. Two scalably encoded video sequences stream over two multicast trees as shown in Fig. 1. The outgoing link speed of Peer 5 varies from 6Mb/s to 54Mb/s, while all other links operate at 54Mb/s. (a) *City* versus *Downtown*. (b) *Downtown* versus *Crew*. (c) *Terrace* versus *Ice*.

by the TFRC-based heuristics. In addition to adapting the allocated video rates to their respective RD characteristics, the media-aware scheme also allocates higher video rates for peers receiving over faster links. The TFRC-based heuristics, in contrast, is oblivious to difference in individual link speeds since the observed virtual packet loss ratios and round trip times are affected by competitions in packet transmissions over all neighboring links. As the outgoing link speed of Peer 5 increases from 6Mb/s to 54Mb/s, the average video quality resulting from the media-aware allocation increases accordingly. It outperforms the TFRC-based heuristics by 0.7–1.3 dB in PSNR of average video distortion across all peers.

*E. Comparison of Scalable and Non-Scalable Multicast*

Finally, we study the effect of scalable representation for video multicast by comparing the results from the previous two sections. For completeness, our comparison also include additional results of non-scalable multicast of scalable streams. The three alternatives are summarized as follows.

- 1) *Scalable stream, scalable multicast (SSSM)*: Multicast of scalable video streams encoded using the SVC JSVM 8.8 reference software [38]. Rate adaptation is performed at each peer inside each multicast tree, following (18). Simulation results are the same as those presented in Section VI-D.
- 2) *Scalable stream, non-scalable multicast (SSNM)*: Multicast of scalable video streams encoded using the SVC JSVM 8.8 reference software [38]. Rate adaptation is performed at the root peer only, following (8) for each entire multicast tree.
- 3) *Non-scalable stream, non-scalable multicast (NSNM)*: Multicast of non-scalable video streams encoded using  $\times 264$  [37], a fast implementation of the H.264/AVC

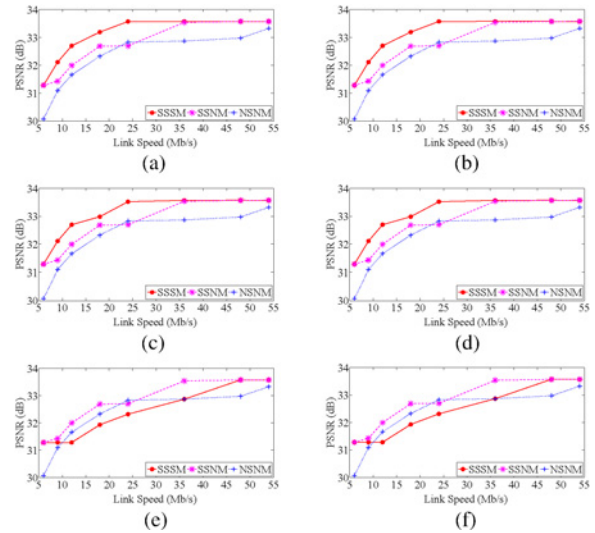


Fig. 12. Comparison of SSSM, as well as SSNM and NSNM in terms of video quality of each peer. The *Harbor* sequence streams over the first multicast tree shown in Fig. 1. The outgoing link speed of Peer 5 varies from 6Mb/s to 54Mb/s, while all other links operate at 54Mb/s. (a) Peer 2. (b) Peer 4. (c) Peer 3. (d) Peer 5. (e) Peer 6. (f) Peer 8.

standard [40]. Rate adaptation is performed at the root peer only, following (8) for each entire multicast tree. Simulation results are the same as those presented in Section VI-C.

Comparisons between SSSM and SSNM highlight the potential benefit of scalable multicast from the additional flexibility of rate adaptation inside each multicast tree. Performance difference between SSNM and NSNM, on the contrary, can be attributed to the different RD tradeoffs achieved by the two encoding schemes.

1) *Single Multicast Tree*: We first compare the three alternatives in the simple case of streaming the *Harbor* sequence over a single multicast tree. Fig. 12 plots the video quality of each peer against the link speed of Peer 5. Comparing SSSM against SSNM, it can be noted that scalable multicast tends to improve the video quality of peers with higher link speeds, at the expense of lowering the video quality of peers with lower link speeds. Such difference is more obvious at intermediate link speeds of Peer 5. Overall, the NSNM scheme achieves similar results to SSNM, with slightly lower video qualities due to less efficient RD performance of the non-scalable stream.

Fig. 13 summarizes the comparison for simulations from all six sequences. Contrasting SSSM against SSNM, it can be noted that scalable multicast slightly improves the average video quality. Due to limited rate and quality range of the scalable video streams, the performance gains are only observable when streaming more demanding sequences *Crew* and *Harbor*, or when the link speed of Peer 5 is at 6Mb/s for *City* and *Terrace*. For the same reason, the NSNM scheme benefits from a wider choice of available rates and qualities from multiple non-scalable streams, therefore attains higher average video quality when streaming the less demanding *Downtown* and *Ice* sequences, or when supporting *City* and *Terrace* at higher link speeds of Peer 5.




Fig. 13. Comparison of SSSM, as well as SSNM and NSNM in terms of average video quality of all peers. One sequence streams over the first multicast tree shown in Fig. 1. The outgoing link speed of Peer 5 varies from 6 Mb/s to 54 Mb/s, while all other links operate at 54 Mb/s. (a) *City*. (b) *Crew*. (c) *Downtown*. (d) *Harbor*. (e) *Ice*. (f) *Terrace*.




Fig. 14. Comparison of SSSM, as well as SSNM and NSNM in terms of average video quality of all peers. Two video sequences stream over two multicast trees as shown in Fig. 1. The outgoing link speed of Peer 5 varies from 6 Mb/s to 54 Mb/s, while all other links operate at 54 Mb/s. (a) *City* versus *Downtown*. (b) *Downtown* versus *Crew*. (c) *Terrace* versus *Ice*.

2) *Multiple Multicast Trees*: Fig. 14 further compares the three schemes for two video streams over two multicast trees. The benefit of scalable multicast can be observed for two out of three sequence pairs. The performance gain is more obvious at lower link speeds of Peer 5. As the network becomes more loaded with two streams, the benefits of both SSSM and SSNM over NSNM start to prevail. For instance, the gain in overall video quality is up to 1.2 dB in PSNR, for *City* versus *Downtown*. In the case of *Downtown* versus *Crew*, on the contrary, NSNM still achieves a higher average video quality available from a higher-quality version of the non-scalable streams.

## VII. CONCLUSION

This paper presented a unified framework for media-aware rate allocation among multiple video multicast sessions sharing a wireless network. For delivery of both scalable and non-scalable video streams, the proposed scheme aimed at minimizing the total video distortion of all peers while limiting network utilization. Our distributed solution leveraged cross-layer information exchange between MAC-layer LSMs and application-layer VRCs, to achieve fast convergence at the optimal allocated video rates. The resulting allocation was both media-aware and network-aware. Video streams with more demanding RD characteristics tend to be allocated higher rates. In scalable multicast, peers receiving over faster links get higher rates than peers over slower links. Performance of the proposed media-aware allocation is compared against a heuristic scheme based on TFRC, in simulations of SD video streaming over single or multiple multicast trees. For delivery of both scalable and non-scalable streams, the media-aware


allocation consistently outperforms TFRC-based heuristics in terms of average video quality of all peers. The additional flexibility of per-peer rate adaptation in scalable multicast yields a further slight improvement in overall video quality. On the contrary, the limited rate and quality range of a scalable stream tends to limit the improved efficiency of scalable multicast, when compared with non-scalable multicast of non-scalable streams.

One main limitation of the proposed scheme is that it relies on cooperative behavior from all participating multicast sessions to achieve the desired rate allocation results. In reality, however, video streams sharing a common wireless network may follow different rate adaptation schemes, for instance based on TCP-friendly allocation. It therefore remains an intriguing future research topic to augment the design of the proposed media-aware rate allocation scheme for coexistence with other TCP-based flows.

## REFERENCES

- [1] X. Zhu, T. Schierl, T. Wiegand, and B. Girod, "Video multicast over wireless mesh networks with scalable video coding (SVC)," *Proc. SPIE VCIP*, vol. 6822, pp. 1–8, Jan. 2008.
- [2] A. Ganjam and H. Zhang, "Internet multicast video delivery," *Proc. IEEE*, vol. 93, no. 1, pp. 159–170, Jan. 2005.
- [3] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in *Proc. 22nd Annu. Joint Conf. IEEE Comput. Commun. Soc. INFOCOM*, vol. 2, Mar. 2003, pp. 836–843.
- [4] X. Zhu and B. Girod, "Distributed rate allocation for multi-stream video transmission over ad hoc networks," in *Proc. IEEE ICIP*, vol. 2, Sep. 2005, pp. 157–160.
- [5] X. Zhu, P. van Beek, and B. Girod, "Distributed channel time allocation and rate adaptation for multi-user video streaming over wireless home networks," in *Proc. IEEE ICIP*, Sep. 2007, pp. 69–72.
- [6] X. Zhu and B. Girod, "Distributed media-aware rate allocation for wireless video streaming," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 20, no. 11, pp. 1462–1474, Nov. 2010.
- [7] T. Jiang, E. W. Zegura, and M. Ammar, "Inter-receiver fair multicast communication over the Internet," in *Proc. ACM 9th Int. Workshop*, Jun. 1999, pp. 103–114.
- [8] H. Yousefi-zadeh, H. Jafarkhani, and A. Habibi, "Layered media multicast control (LMMC): Rate allocation and partitioning," *IEEE/ACM Trans. Netw.*, vol. 13, no. 3, pp. 540–553, Mar. 2005.
- [9] J.-C. Bolot and T. Turletti, "A rate control mechanism for packet video in the Internet," in *Proc. IEEE Annu. Joint Conf. IEEE Comput. Commun. Soc. (INFOCOM)*, Jun. 1994, pp. 1216–1223.
- [10] S. Y. Cheung, M. H. Ammar, and X. Li, "On the use of destination set grouping to improve fairness in multicast video distribution," in *Proc. IEEE Annu. Joint Conf. IEEE Comput. Commun. Soc. (INFOCOM)*, Mar. 1996, pp. 553–560.
- [11] S. McCanne, V. Jacobson, and M. Vetterli, "Receiver-driven layered multicast," in *Proc. SIGCOMM*, vol. 26, no. 4, Aug. 1996, pp. 117–130.
- [12] X. Li, S. Paul, and M. Ammar, "Layered video multicast with retransmissions (LVMR): Evaluation of hierarchical rate control," in *Proc. IEEE Annu. Joint Conf. IEEE Comput. Commun. Soc. (INFOCOM)*, vol. 3, Mar. 1998, pp. 1062–1072.
- [13] P. Chou, A. Wang, and S. Mehrotra, "Error control for receiver-driven layered multicast of audio and video," *IEEE Trans. Multimedia*, vol. 3, no. 1, pp. 108–122, Jan. 2001.
- [14] W. Tan and A. Zakhor, "Video multicast using layered FEC and scalable compression," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 11, no. 3, pp. 373–386, Mar. 2001.
- [15] X. Li, S. Paul, and M. Ammar, "Multi-session rate control for layered video multicast," *Proc. SPIE Multimedia Comput. Netw.*, vol. 3654, pp. 175–189, Jan. 1999.
- [16] L. Vicisano, L. Rizzo, and J. Crowcroft, "TCP-like congestion control for layered multicast data transfer," in *Proc. IEEE Annu. Joint Conf. IEEE Comput. Commun. Soc. (INFOCOM)*, Mar. 1999, pp. 996–1003.

- [17] K. Kar, S. Sarkar, and L. Tassiulas, "Optimization based rate control for multirate multicast sessions," in *Proc. 20th Annu. Joint Conf. IEEE Comput. Commun. Soc. (INFOCOM)*, vol. 1, Apr. 2001, pp. 123–132.
- [18] A. Majumdar and K. Ramchandran, "Video multicast over lossy channels based on distributed source coding," in *Proc. IEEE ICIP*, vol. 5, Oct. 2004, pp. 3093–3096.
- [19] J. Wang and K. Ramchandran, "Receiver-driven multicast over wireless with distributed source coding and FEC," in *Proc. IWCMC*, Aug. 2007, pp. 600–605.
- [20] S. Mao, X. Cheng, Y. T. Hou, and H. D. Sherali, "Multiple description video multicast in wireless ad hoc networks," in *Proc. Int. Conf. Broadband Netw.*, 2004, pp. 671–680.
- [21] W. Wei and A. Zakhor, "Multipath unicast and multicast video communication over wireless ad hoc networks," in *Proc. Int. Conf. Broadnets*, Oct. 2004, pp. 496–505.
- [22] W. Wei and A. Zakhor, "Multiple tree video multicast over wireless ad hoc networks," *IEEE Trans. Circuits, Syst. Video Technol.*, vol. 17, no. 1, pp. 2–15, Jan. 2007.
- [23] A. Kovacs and I. Godor, "Cross-layer optimized wireless multicast for layered media," *Comput. Netw.*, vol. 53, no. 7, pp. 1062–1072, May 2009.
- [24] Y. He, I. Lee, and L. Guan, "Optimized video multicasting over wireless ad hoc networks using distributed algorithm," *IEEE Trans. Circuits, Syst. Video Technol.*, vol. 19, no. 6, pp. 796–807, Jun. 2009.
- [25] D. Wu, T. Hou, and Y.-Q. Zhang, "Scalable video coding and transport over broadband wireless networks," *Proc. IEEE, Special Issue Multi-Dimensional Broadband Wireless Technol. Applicat.*, vol. 89, no. 1, pp. 6–20, Jan. 2001.
- [26] T. Stockhammer, H. Jenkac, and G. Kuhn, "Streaming video over variable bit-rate wireless channels," *IEEE Trans. Multimedia*, vol. 6, no. 2, pp. 268–277, Apr. 2004.
- [27] Y. Li, A. Markopoulou, N. Bambos, and J. Apostolopoulos, "Joint power-playout control for media streaming over wireless links," *IEEE Trans. Multimedia*, vol. 8, no. 4, pp. 830–843, Aug. 2006.
- [28] IEEE 802.11 WG, *IEEE Standard for Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, Nov. 1997.
- [29] K. Stuhlmüller, N. Färber, M. Link, and B. Girod, "Analysis of video transmission over lossy channels," *IEEE J. Select. Areas Commun.*, vol. 18, no. 6, pp. 1012–1032, Jun. 2000.
- [30] X. Zhu and B. Girod, "Subjective evaluation of multi-user rate allocation for streaming heterogeneous video contents over wireless networks," in *Proc. IEEE ICIP*, Oct. 2008, pp. 3092–3095.
- [31] N. Z. Shor, *Minimization Methods for Non-Differentiable Functions*. New York: Springer, 1985.
- [32] F. P. Kelly, A. Maulloo, and D. Tan, "Rate control for communication networks: Shadow prices, proportional fairness and stability," *J. Oper. Res. Soc.*, vol. 49, no. 3, pp. 237–252, Mar. 1998.
- [33] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [34] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, "RTP: A transport protocol for real-time applications," document RFC 3550 (Standard), IETF, Jul. 2003.
- [35] *Advanced Video Coding for Generic Audiovisual Services (SVC Extensions)*, ITU-T Rec. H.264-ISO/IEC 14496-10, ITU-T and ISO/IEC JTC 1, Jul. 2007.
- [36] H. Schwarz, D. Marpe, and T. Wiegand, "Overview of the scalable video coding extension of H.264/AVC," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 17, no. 9, pp. 1103–1120, Sep. 2007.
- [37] VideoLan Organization. *x264* [Online]. Available: <http://developers.videolan.org/x264.html>
- [38] Joint Video Team (JVT). *Joint Scalable Video Model (JSVM) Reference Software* [Online]. Available: [http://iphome.hhi.de/imagecom/\\_G1/savce/index.htm](http://iphome.hhi.de/imagecom/_G1/savce/index.htm)
- [39] *The Network Simulator NS-2* [Online]. Available: <http://www.isi.edu/nsnam/ns>
- [40] *Advanced Video Coding for Generic Audiovisual Services*, ITU-T Rec. H.264-ISO/IEC 14496-10, ITU-T and ISO/IEC JTC 1, May 2003.
- [41] S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance," *IEEE/ACM Trans. Netw.*, vol. 1, no. 4, pp. 397–413, Aug. 1993.
- [42] K. K. Ramakrishnan, S. Floyd, and D. Black, *The Addition of Explicit Congestion Notification (ECN) to IP*, document RFC 3168 (Proposed Standard), Sep. 2001.
- [43] S. Floyd and K. Fall, "Promoting the use of end-to-end congestion control in the Internet," *IEEE/ACM Trans. Netw.*, vol. 7, no. 4, pp. 458–472, Aug. 1999.



**Xiaoqing Zhu** received the B.E. degree in electronics engineering from Tsinghua University, Beijing, China, in 2001, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 2002 and 2009, respectively.

She is with the Advanced Architecture and Research Group, Cisco Systems, Inc., San Jose, CA. She held summer internships at IBM Almaden Research Center, San Jose, in 2003, and at the Sharp Laboratories of America, Camas, WA, in 2006.

Her current research interests include wireless video networking, distributed resource allocation, and performance enhancement for Internet video delivery.

Dr. Zhu was awarded the Stanford Graduate Fellowship from 2001 to 2005, and received the Best Student Paper Award in ACM Multimedia in 2007.



**Thomas Schierl** (M'07) received the Dipl.-Ing. (M.Sc.) degree in computer engineering in 2003 and the Dr.-Ing. (Ph.D.) degree in electrical engineering in 2010, both from the Berlin Institute of Technology, Berlin, Germany.

He is currently the Head of the Multimedia Communications Group in the Image Processing Department of Fraunhofer Heinrich Hertz Institute (HHI) and responsible for multimedia and network-related research and standardization activities of HHI. He and his group typically contribute to standardization

organizations including IETF, MPEG, 3GPP, and DVB. He is co-authoring various IETF RFCs and ISO/IEC MPEG standards as well as various ACM and IEEE research papers on multimedia transmission. In 2007, he visited the group of Prof. B. Girod at Stanford University, Stanford, CA. His current research interests include mobile media content delivery and multimedia cloud computing.



**Thomas Wiegand** (M'05–SM'08–F'11) received the Dipl.-Ing. degree in electrical engineering from the Technical University of Hamburg-Harburg, Hamburg, Germany, in 1995, and the Dr.-Ing. degree from the University of Erlangen-Nuremberg, Erlangen, Germany, in 2000.

He is a Professor with the Department of Electrical Engineering and Computer Science at the Berlin Institute of Technology, Berlin, Germany, chairing the Image Communication Laboratory, and jointly heading the Image Processing Department of the

Fraunhofer Institute for Telecommunications-Heinrich Hertz Institute, Berlin. He joined the Heinrich Hertz Institute in 2000 as the Head of the Image Communication Group in the Image Processing Department. From 1993 to 1994, he was a Visiting Researcher with Kobe University, Kobe, Japan. In 1995, he was a Visiting Scholar with the University of California at Santa Barbara, Santa Barbara. From 1997 to 1998, he was a Visiting Researcher with Stanford University, Stanford, CA, and served as a Consultant to 8x8, Inc., Santa Clara, CA. From 2006 to 2008, he was a Consultant to Stream Processors, Inc., Sunnyvale, CA. From 2007 to 2009, he was a Consultant to Skyfire, Inc., Mountain View, CA. Since 2006, he has been a Technical Advisory Board Member of Vido, Inc., Hackensack, NJ. Since 1995, he has been an active participant in standardization for multimedia with successful submissions to ITU-T VCEG, ISO/IEC MPEG, 3GPP, DVB, and IETF. His current research interests include video processing and coding, multimedia transmission, as well as computer vision and graphics.

Prof. Wiegand received the SPIE VCIP Best Student Paper Award in 1998. In October 2000, he was appointed as the Associated Rapporteur of ITU-T VCEG. In December 2001, he was appointed as the Associated Rapporteur/Co-Chair of the JVT. In February 2002, he was appointed as the Editor of the H.264/MPEG-4 AVC video coding standard and its extensions (FRExt and SVC). From 2005 to 2009, he was the Co-Chair of MPEG Video. In 2004, he received the Fraunhofer Award and the ITG Award of the German Society for Information Technology. The projects that he co-chaired for development of the H.264/AVC standard have been recognized by the 2008 ATAS Primetime Emmy Engineering Award and a pair of NATAS Technology and Engineering Emmy Awards. In 2009, he received the Innovations Award of the Vodafone Foundation, the EURASIP Group Technical Achievement Award, and the Best Paper Award of IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY. In 2010, he received the Eduard Rhein

Technology Award. In 2011, he received the Best Paper Award of EURASIP and was elected as fellow of the IEEE. He was a Guest Editor for the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY for its Special Issue on the H.264/AVC Video Coding Standard in July 2003, its Special Issue on Scalable Video Coding-Standardization and Beyond in September 2007, and its Special Section on the Joint Call for Proposals on High Efficiency Video Coding (HEVC) Standardization. From 2006 to 2010, he was an Associate Editor of IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY.



**Bernd Girod** (M'80–SM'97–F'98) received the Engineering Doctorate degree from the University of Hannover, Hannover, Germany, and the M.S. degree from the Georgia Institute of Technology, Atlanta.

Since 1999, he has been a Professor of electrical engineering and (by courtesy) computer science with the Information Systems Laboratory, Stanford University, Stanford, CA. Previously, he was a Professor of telecommunications with the Department of Electrical Engineering, University of Erlangen-Nuremberg, Erlangen, Germany. His current research interests include video compression and networked media systems. He has published over 400 conference and journal papers, and five books.

Prof. Girod was the recipient of the EURASIP Signal Processing Best Paper Award in 2002, the IEEE Multimedia Communication Best Paper Award in 2007, the EURASIP Image Communication Best Paper Award in 2008, and the EURASIP Technical Achievement Award in 2004. As an entrepreneur, he has been involved with several startup ventures as Founder, Director, Investor, or Advisor, and among them are Polycom (Nasdaq:PLCM), Pleasanton, CA, Vivo Software, Waltham, MA, 8x8 (Nasdaq:EGHT), Santa Clara, CA, and RealNetworks (Nasdaq:RNWK), Seattle, WA. He is a fellow of the EURASIP and a member of the German National Academy of Sciences (Leopoldina).