CROSS-LAYER DESIGN OF AD HOC NETWORKS FOR REAL-TIME VIDEO STREAMING

ERIC SETTON, TAESANG YOO, XIAOQING ZHU, ANDREA GOLDSMITH, AND BERND GIROD
STANFORD UNIVERSITY

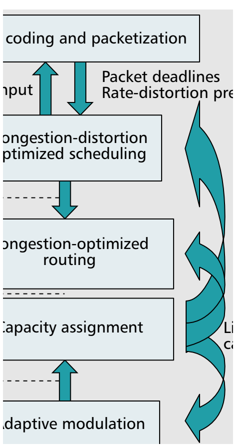

Cross-layer design breaks away from traditional network design where each layer of the protocol stack operates independently. The authors explore the potential synergies of exchanging information between different layers to support real-time video streaming.

This work was supported, in part, by National Science Foundation Grant CCR-0325639.

ABSTRACT

Cross-layer design breaks away from traditional network design where each layer of the protocol stack operates independently. We explore the potential synergies of exchanging information between different layers to support real-time video streaming. In this new approach information is exchanged between different layers of the protocol stack, and end-to-end performance is optimized by adapting to this information at each protocol layer. We discuss key parameters used in the cross-layer information exchange along with the associated cross-layer adaptation. Substantial performance gains through this cross-layer design are demonstrated for video streaming.

INTRODUCTION

An ad hoc wireless network is a collection of wireless nodes that self-configure to form a network without the aid of any established infrastructure. Some or possibly all of these nodes are mobile. These networks are extremely compelling for applications where a communications infrastructure is too expensive to deploy, cannot be deployed quickly, or is simply not feasible. There are numerous potential applications for ad hoc wireless networks, ranging from multihop wireless broadband Internet access, to sensor networks, to building or highway automation, to voice, image, and video communication for disaster areas.

The lack of established infrastructure, the network and channel dynamics, and the nature of the wireless medium offer an unprecedented set of challenges in supporting demanding applications over ad hoc wireless networks. The wireless channel is inherently a broadcast medium, so transmissions from different nodes interfere with each other. The quality of wireless links vary over time and space due to interference, multipath fading, and shadowing. Network conditions are highly dynamic as nodes join and leave the network in an unpredictable manner. Furthermore, as new links are formed and others vanish, routing of traffic from one node to another may frequently change. These daunting challenges have spurred a large body of research on the design, analysis, and fundamental performance limits of ad hoc wireless networks.

For video streaming, high bandwidth requirements are coupled with tight delay constraints as packets need to be delivered in a timely fashion to guarantee continuous media playout. When packets are lost or arrive late, the picture quality suffers as decoding errors tend to propagate to subsequent portions of the video. Due to the high bit rate requirements of video, a media stream may congest the network significantly. Hence, it is imperative to account for the potential impact of each video user on the network statistics and guarantee that the network is not operating beyond its capacity. Unfortunately, most network designs do not provide mechanisms for protocol layers to optimally adapt to underlying channel conditions and specific application requirements. While protocol layering is an important abstraction that reduces network design complexity, it is not well suited to wireless networks since the nature of the wireless medium makes it difficult to decouple the layers. Moreover, meeting the end-to-end performance requirements of demanding applications is extremely challenging without interaction between protocol layers.

We believe that video streaming over ad hoc wireless networks can benefit substantially from a cross-layer design. In this design, interdependencies between layers are characterized and exploited by adapting to information exchanged between layers and building the appropriate amount of robustness into each layer. For example, routing protocols can avoid links experiencing deep fades, or the application layer can adapt its transmission rate based on the underlying network throughput and latency. In this article we explore a new framework for cross-layer design that incorporates adaptation across all layers of the protocol stack: application, transport protocols, resource allocation, and link layer techniques. Link capacities are dynamically reallocated based on link state information, and stream-based multipath source routing and scheduling is performed by balancing network

 1536-1284/05/$20.00

A cross-layer approach to network design seeks to enhance the performance of a system by jointly designing multiple protocol layers. This approach allows upper layers to better adapt their strategies to varying link and network conditions.

congestion and distortion of real-time media streams. This framework includes new techniques to jointly optimize error-resilient source coding, packet scheduling, stream-based routing, link capacity assignment, and adaptive link layer techniques. The optimization is carried out dynamically by all nodes, based on link state communication, to continuously adapt to the changing wireless link conditions and traffic flows. We demonstrate that this cross-layer framework can provide significant performance gains with manageable complexity.

This article starts with a survey of current research in video over wireless networks, and a description of the new cross-layer design framework. To illustrate this framework, several preliminary results from ongoing research are presented, including adaptive link layer techniques, joint capacity and flow assignment for video streaming, congestion-distortion optimized packet scheduling and optimized rate allocation.

Related Work

Supporting multimedia applications over wireless links has been one of the main fields of attention in the networking and video coding communities in the last decade. For such applications, it is well known that the separation principle of source and channel coding put forth in Shannon's information theoretic framework does not hold. Hence, joint source and channel coding techniques have been proposed to overcome the challenges of wireless links. However, these have not yet led to a unified solution to this problem, as explained in [1]. Nevertheless, these considerations have strongly influenced the video community in the design of the new H.264 video coding standard, which incorporates a highly flexible syntax well suited to network transmission [2].

As the number of nodes of a wireless network grows, interference increases, reducing the achievable data rates. In a landmark paper, the capacity of a static wireless ad hoc network is shown to asymptotically vanish as the number of users increases [3]. However, recent results show that in more practical settings, the high data rates and low delay constraints of multimedia applications may be supported [4, 5]. The 802.11 protocol operating in ad hoc mode provides an interesting benchmark against which proposed designs for ad hoc networks may compete. In its simplest form, it allows for one user at a time to be active within a carrier sense region (typically on the order of a few hundred meters). As only active hosts occupy the wireless medium, this medium access control (MAC) protocol comes closer to the achievable capacity than time-division multiple access (TDMA) or frequency-division multiple access (FDMA) systems. In this sense it already incorporates some cross-layering. However, this design is restrictive in more than one way and may have to be surpassed to enable efficient media streaming.

The art of video streaming over ad hoc wireless networks is still in its infancy, especially when addressed via a cross-layer network design. Most of the recent research considers only a subset of layers of the protocol stack. In [4] path diversity in an 802.11 network combined with multistream coding of video is proposed and analyzed. Other cross-layer approaches have been suggested, as in [5], where source, channel coding, packetization, and MAC layer retransmissions are performed together to reach optimized usage of the wireless channel. Power and flow may also be allocated jointly through convex optimization to minimize network congestion [6], and this joint allocation can be combined with MAC layer scheduling [7]. Much work still needs to be done along these directions to identify and exploit cross-layer interactions in real-time media streaming over ad hoc wireless networks.

Cross-Layer Design Framework

A cross-layer approach to network design seeks to enhance the performance of a system by jointly designing multiple protocol layers. This approach allows upper layers to better adapt their strategies to varying link and network conditions. The resulting flexibility helps to improve end-to-end performance given network resources and dynamics. These design concepts are particularly useful for supporting delay-constrained applications such as video.

A cross-layer approach to network design can significantly increase the design complexity. Indeed, protocol layers are extremely useful in allowing designers to optimize a single protocol layer design without the complexity and expertise associated with considering other layers. Thus, cross-layer design should not eliminate the design advantages of layering. Keeping some form of separation, while allowing layers to actively interact, appears to be a good compromise for enabling interaction between layers without eliminating the layering principle. In such a structure each layer is characterized by some key parameters, which are passed to the adjacent layers to help them determine the operation modes that will best suit the current channel, network, and application conditions. In such a design each layer is not oblivious of the other layers, but interacts with them to find its optimal operational point. The main design difficulty in this cross-layer approach resides in characterizing the essential information that should be exchanged between layers. For example, the link layer might be characterized by parameters representing the channel quality, such as signal-to-interference-plus-noise ratio (SINR), or link layer state information such as the bit error rate (BER) or supported data rate. Similarly, the network and MAC layers might exchange the requested traffic rates and supportable link capacities.

Figure 1 shows a block diagram of a new cross-layer design framework with information exchange between the different layers. At the link layer, adaptive techniques are used to maximize the link rates under varying channel conditions. This extends the achievable capacity region of the network. Each point of this region indicates a possible assignment of the different link capacities. Based on link state information, the MAC selects one point of the capacity region by assigning time slots, codes, or frequency bands to each of the links. The MAC layer operates jointly with the network layer to determine the set of network flows that minimize congestion. To find a jointly optimal solution for capacity assignment and network flows, successive suboptimal solutions are exchanged iteratively between these two middle layers tht constitute the core of our cross-layer framework. At the transport layer, congestion-distortion optimized scheduling is performed to control the transmission and retransmission of video packets. Finally, the application layer determines the most efficient encoding rate.

In the remainder of this article we discuss the application of this framework to low-latency video streaming. Specifically, in the next section we present in more detail some adaptive link layer techniques. These techniques determine the maximum rates that can be transmitted between any two nodes in the network. We show how the operating network capacity can be jointly optimized with network flow to improve the end-to-end performance of video streaming. We consider application layer information to determine which packets should be transmitted and when to maximize the decoded video quality while limiting network congestion, and explain how to determine the optimal streaming rate.

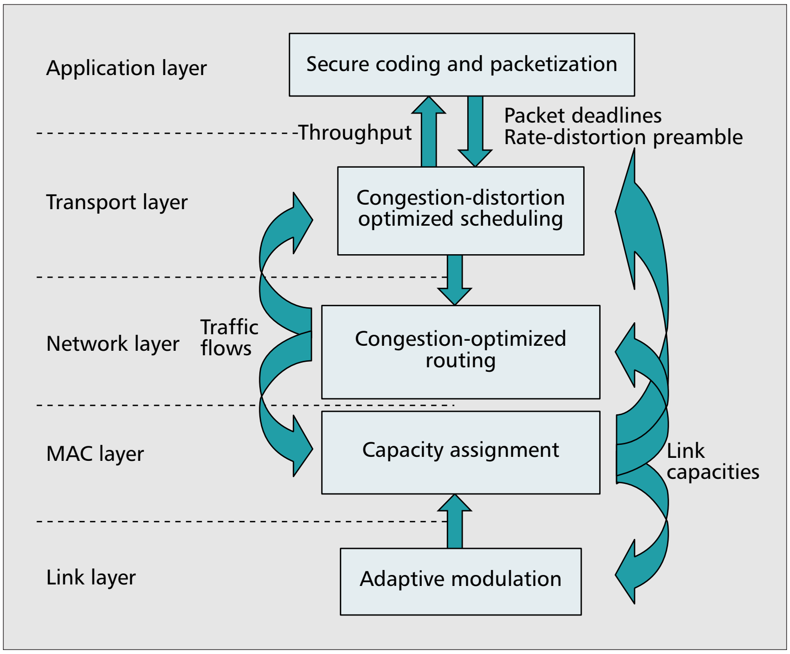

Figure 1. *Cross-layer design framework for low-latency media streaming over ad hoc wireless network.*

Adaptive Link Layer Technique

Wireless link throughput is severely affected by channel impairments such as shadowing, multipath fading, and interference. Adaptive modulation is an efficient technique to improve the data rate by adapting link layer design variables to the variations of the channel environment. These parameters may include modulation, coding, transmitter power, target BER, symbol rate, and combinations of these parameters.

In general, adaptive link layer techniques have not been considered in a cross-layer framework. In this section we consider two adaptive link layer techniques to improve link throughput. We first consider adapting the packet length, given the current SINR and link layer parameters, to optimize throughput. In addition, for a fixed packet length, we consider optimizing link layer parameters such as symbol rate and constellation size for maximal throughput. Indeed, when the packet length is too large, packet error rate (PER) increases, and throughput is limited by frequent retransmissions. On the other hand, the smaller the packet the larger the overhead. The optimal packet length L* which maximizes throughput is given by

$$L^* = \frac{H}{2} + \frac{1}{2}\sqrt{H^2 - \frac{4bH}{\ln(1-P_e)}}, \qquad (1)$$

where H is the size of a packet header,[1] b is the number of bits per symbol, and P_e is the probability of symbol error, which will depend on the modulation type and link SINR [8]. A joint optimization of the packet length and link layer parameters yields the following rule of thumb.

High SINR region: When the channel gain is high, the maximum symbol rate should be used with the highest tolerable constellation size. The corresponding optimal packet length increases as the channel quality improves.

Low SINR region: The error rate should be decreased by adding redundancy in the transmission. One way to accomplish this is to reduce symbol rate; a more practical implementation could use a larger spreading factor or heavier coding. In any case, a packet is transmitted over an extended period of time. The optimal amount of redundancy is such that the optimal symbol error rate in Eq. 1 remains constant.

When all the parameters are optimally chosen, we obtain a point-to-point link throughput given by

$$R = W\log_2\left(1 + \frac{\text{SINR}}{\Gamma}\right), \qquad (2)$$

where W is the channel bandwidth and Γ is a parameter determined by the link layer design. The throughput performance of this adaptive link layer technique is shown in Fig. 2 for uncoded transmission, where a constant gap from Shannon capacity is observed, for different types of channels. This gap could be reduced by using various channel coding techniques.

In a network with more than two nodes, the wireless medium should be shared among multiple sender-receiver pairs through a *transmission strategy* that coordinates medium access. For a given arrangement of nodes an achievable data rate on each link using our adaptive technique is computed by Eq. 2 and the time fraction during which the link will be active under a particular transmission strategy. By considering all the different transmission strategies between senders and receivers, one can compute an achievable capacity region that characterizes the data rates that are simultaneously achievable between each source-destination pair.[2] Upon further inspection, only a limited number of these strategies

[1] H *reflects the frame header as well as the protocol overhead necessary to send a packet in a practical implementation.*

[2] *Synchronization and control overhead would reduce the achievable rates. They are not considered here as we seek to derive a bound on the rate of time-constrained media supportable by the network.*

are relevant, in that they achieve rates that cannot be obtained by time-sharing between other strategies. For an ad hoc network of limited size (e.g., tens of nodes) this region can be computed exactly [9]. Figure 3 shows an example of this achievable capacity region that shows all the rate pairs simultaneously achievable between two specific sender-receiver pairs. Note that the use of a more advanced transmission scheme increases the achievable capacity region. A capacity, as well as a flow, shall be optimally assigned within this region so that traffic requests can be accommodated with minimal congestion. This topic is discussed in the next section.

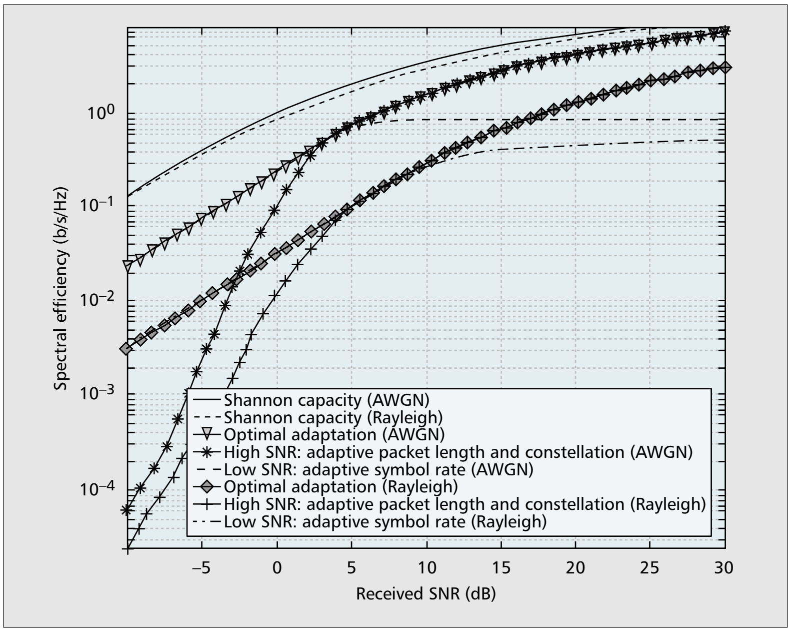

Figure 2. *Throughput performance comparison.*

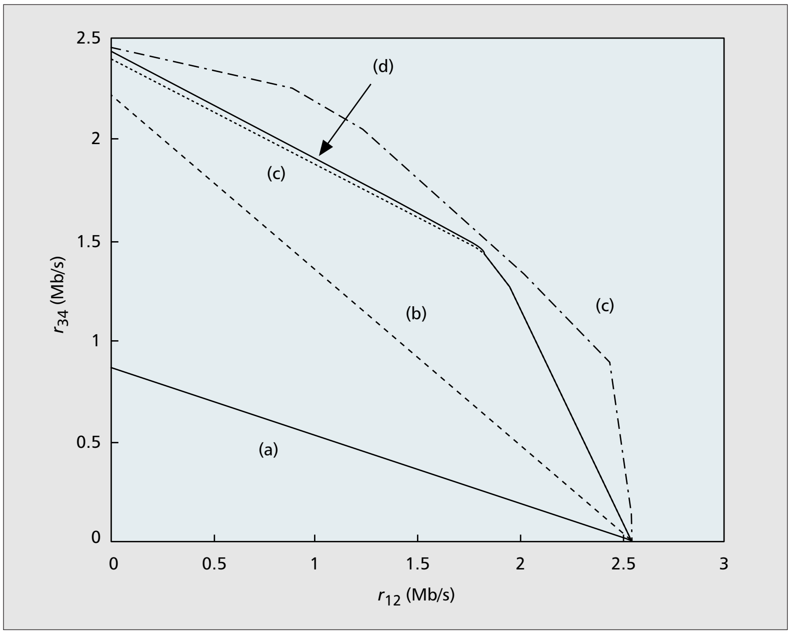

Figure 3. *Capacity region slices of an example ad hoc network: a) single-hop routing, no spatial reuse; b) multihop routing, no spatial reuse; c) multihop routing with spatial reuse; d) two-level power control added to c; e) successive interference cancellation added to c.*

Joint Allocation of Capacity and Flow

Given the network capacity, the network layer assigns traffic flows to each link. The closer the flow is to the link capacity, the higher the latency on that link. The optimal capacity and flow assignment may be computed to minimize any convex or quasi-convex cost function, such as the maximum link utilization over the links of the network,

$$\Delta(\mathbf{C}, \mathbf{f}) = \max_{(i,j)} \frac{f_{ij}}{C_{ij}}. \tag{3}$$

As explained in [10], this metric of network congestion is a quasi-convex function of the network flows and capacities of the network, **f** and **C**. It can be minimized, for example, by a bisection algorithm that involves solving a sequence of convex feasibility problems.

Optimal solutions make efficient use of network resources by favoring good-quality links. Typically, the flow between a sender and a receiver is split among multiple paths. This has some advantages as path diversity may provide higher aggregate data rate through spatial reuse of the wireless spectrum. Multiple routes also have uncorrelated loss patterns; this can be exploited at the application layer. On the other hand, use of a higher number of links creates more contention at the MAC layer, and the complexity of maintaining multiple routes is higher.

To assess the advantages of cross-layer design, we also consider design with oblivious layers, where capacity and flow are optimized independently. In this design bidirectional links are established between neighboring nodes, and the minimum transmission rate supported by these links is maximized. Given this fixed capacity assignment, the flow assignment is computed by minimizing the same objective function as in Eq. 3. Note that for this kind of design some links are assigned wireless capacity, even though they do not have any traffic to support. This results in a waste of network resources, as illustrated in the following.

The effect of multipath routing is vastly different for each of the two designs. For the network with oblivious layers, the capacity assignment is fixed and only depends on topology. Through multipath routing, the data rate allocated to different paths is aggregated, leading to a larger data rate and thus better performance. Interestingly, for cross-layer design, regardless of the number of paths, high rates can always be supported as resources are only allocated to active links. More details are provided in [10].

Figure 4 shows an example of the performance of the cross-layer design and the oblivious layer design for video streaming over three paths between two nodes of a simulated mobile wireless network. In both cases the video stream is encoded at the highest sustainable rate and transmitted over the User Datagram Protocol (UDP). The received video quality is measured in terms of the peak signal-to-noise ratio (PSNR). The advantage of the cross-layer scheme for which only active links are assigned resources is apparent from the curves. The cross-layer design can support a rate up to 1.9 Mb/s, whereas the design with oblivious layers can only hold 250 kb/s of traffic. Consequently, the best achievable video quality with joint assignment is 39 dB in PSNR on average, 6 dB higher than that of independent optimization.

Scheduling and Rate Allocation

In this section we describe the additional improvement that can be obtained through the use of smart packet scheduling, and discuss how to determine the optimal operating rate for streaming.

Congestion-Distortion Optimized Scheduling

Losses on the wireless medium are inevitable due to interference, collisions, and mobility. The absence of a packet at the decoder causes a decoding error, which translates into a quality drop that may propagate to subsequent frames. This effect is illustrated in the sudden quality drops present in Fig. 4, mainly caused by losses on broken links. These drops may be mitigated by smart scheduling at the transport layer.

Traditional transport layer protocols such as TCP provide reliable transmission but are unaware of delay requirements and relative importance of packets. In cross-layer design superfluous transmissions may be avoided by taking into account application layer delay constraints. In [11] an even more advanced technique is proposed, which seeks optimal transmission schedules based on the importance of each packet of a video stream. This type of scheduling aims at maximizing the decoded video quality at the receiver while abiding with a rate constraint. In congestion-limited situations, it is beneficial to use instead a congestion-distortion optimized (CoDiO) scheduler, which limits end-to-end delay [12]. This metric better reflects the impact of a user's transmissions on the congestion of a network. In addition, is inherently adaptive to time-varying network conditions.

The CoDiO scheduler selects the most important packets in terms of video distortion reduction, and transmits them in an order that minimizes the congestion created on the network. For example, I frames are transmitted with priority, whereas B frames might be dropped, and only the most important frames are retransmitted. In addition, the CoDiO scheduler avoids transmitting packets in large bursts as this increases the queuing delay. Gains for this type of scheduler, over a time-varying ad hoc network, are presented in [12].

Determining the Optimal Operating Rate

The optimized flow assignment described earlier is performed based on network capacity and the rate requested by the application, regardless of any latency requirement. For low-latency streaming, when the transmitted rate exceeds a certain threshold, self-congestion causes too much delay in the network to meet the tight delay constraint,

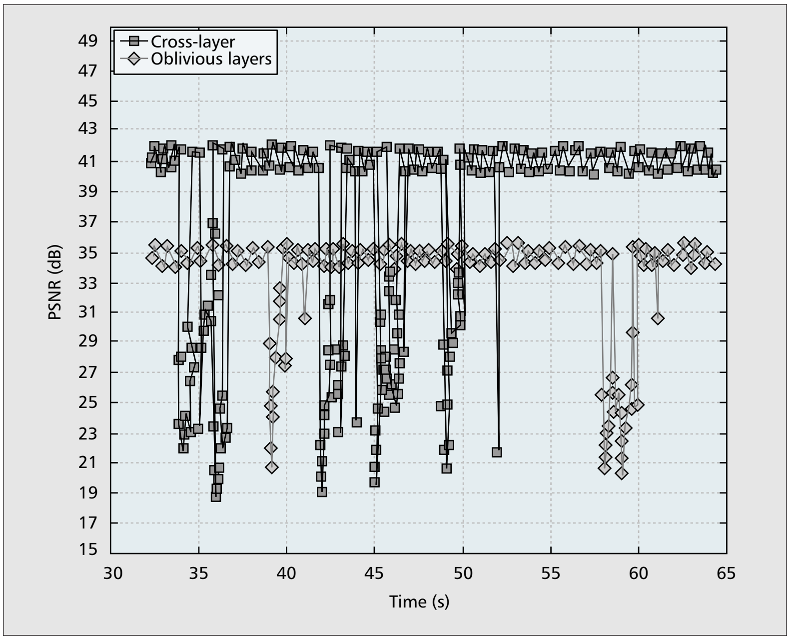

Figure 4. *PSNR performance of video streaming for two different network architectures.*

and the received video quality eventually degrades. With independent layers, the optimal operating rate is determined by algorithms such as TCP-friendly rate control [13] based on end-to-end statistics. In our cross-layer optimization the application layer determines the optimal rate for the video stream, based on rate-distortion characteristics, delay constraints, and current network conditions.

In [14] we derive a model that captures the impact of both encoder quantization and packet losses due to congestion on overall video quality. This model can be used to determine the highest sustainable rate, in conjunction with the optimization described earlier. As an illustration, Fig. 5 shows the decoded video quality for a sequence, encoded at different rates and transmitted over six paths, according to the flow assignment performed earlier. The performance predicted by the model is represented along with experimental results for difference playout deadlines. When the encoding rate approaches the maximum aggregate data rate sustainable by the network (approximately 360 kb/s in this case), the video quality drops, as the end-to-end delay increases due to congestion. When the delay tolerance is small, this performance degradation occurs at rates well below network capacity, as even a slight increase in queuing delay affects the loss rate. In other words, delay constraints reduce the effective network capacity. The highest video quality and the optimal operating rate correspond to the maxima of the curves represented in Fig. 5.

Conclusions

The unique characteristics of wireless ad hoc networks call for new design paradigms that move beyond conventional layering. While joint optimization allows interaction and flexible resource allocation across the network pro-

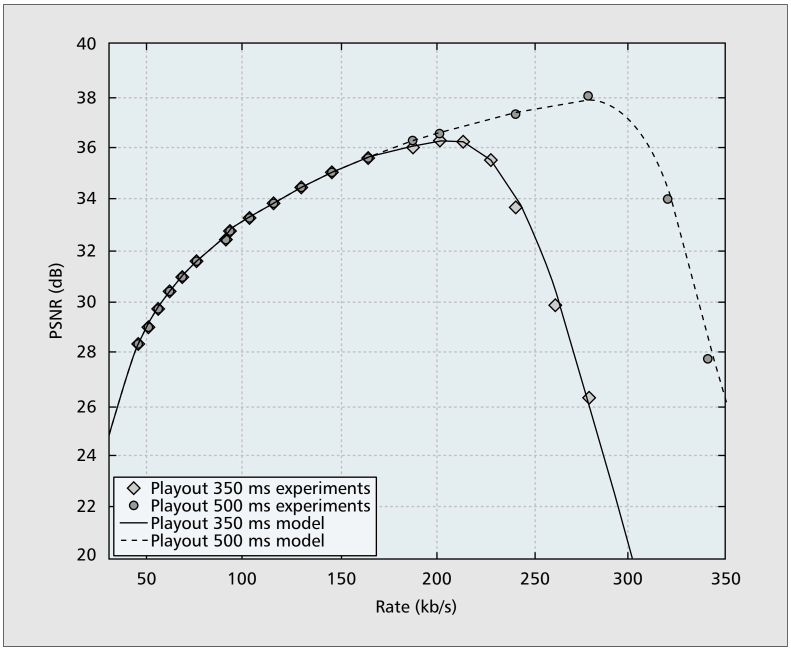

Figure 5. *Rate-PSNR performance for live video streaming of "Foreman" using six-path routing for different playout deadlines.*

tocol stack, the growing complexity may be prohibitive for practical implementation. To strike a balance between performance gains and design complexity, it is important to keep the abstraction of layering while allowing information exchange between adjacent layers. The large gains foreseen by enabling different layers to collaborate are promising for demanding applications such as audio-visual conversations.

We explore a cross-layer design framework for real-time video streaming, which maintains a general layered structure and identifies the key parameters to be exchanged between adjacent layers. In this context adaptive link layer techniques that adjust packet size, symbol rate, and constellation size according to channel conditions are used to improve link throughput, which in turn improves the achievable capacity region of the network. At the MAC and network layers, joint allocation of capacity and flow optimize the supportable traffic rate significantly, and consequently can improve the end-to-end video quality by a wide margin. Smart scheduling at the transport layer further protects the video stream from packet losses and ensures timely arrivals of the video packets without causing excessive network congestion. Knowledge of the video rate-distortion tradeoff and latency requirement at the application layer is used to select the most appropriate source rate for video delivery. Preliminary experimental results illustrate the performance gains achieved by the framework. Future research will focus on achieving similar gains through distributed algorithms while increasing adaptivity, robustness, and scalability.

Many challenging problems lie ahead, and the question of optimal cross-layer design is far from being resolved. While cross-layering provides significant performance advantages, it can also greatly increase complexity, which can make it more difficult to obtain design insights. We expect that the most significant gains will come from interaction between just a few layers, and separation theorems may exist in some settings where decoupling of layers is optimal. These issues present important areas of future research.

References

[1] B. Girod and N. Farber, "Wireless Video," in A. Reibman and M.-T. Sun, Eds., *Compressed Video over Networks*, Marcel Dekker, 1999.

[2] ITU- T Rec. H.264/ISO/IEC 14496-10(AVC), "Advanced Video Coding for Generic Audiovisual Services," 2003.

[3] P. Gupta and P. Kumar, "The Capacity of Wireless Networks," *IEEE Trans. Info. Theory*, vol. 46, no. 2, Mar. 2000, pp. 388–404.

[4] S. Mao *et al.*, "Video Transport over Ad Hoc Networks: Multistream Coding with Multipath Transport," *IEEE JSAC*, vol. 21, no. 10, Dec. 2003, pp. 1721–37.

[5] M. van der Schaar *et al.*, "Adaptive Cross-layer Protection Strategies for Robust Scalable Video Transmission over 802.11 WLANs," *IEEE JSAC*, vol. 21, no. 10, Dec. 2003, pp. 1752–63.

[6] L. Xiao, M. Johansson, and S. Boyd, "Simultaneous Routing and Resource Allocation via Dual Decomposition," *IEEE Trans. Commun.*, vol. 52, no. 7, July 2004, pp. 1136–44.

[7] Y. Wu *et al.*, "Network Planning in Wireless Ad Hoc Networks: a Cross-layer Approach," *IEEE JSAC*.

[8] T. Yoo and A. Goldsmith, "Throughput Optimization using Adaptive Techniques," in preparation.

[9] S. Toumpis and A. J. Goldsmith, "Capacity Regions for Wireless Ad Hoc Networks," *IEEE Trans. Wireless Commun.*, vol. 2, no. 4, July 2003, pp. 736–48.

[10] T. Yoo *et al.*, "Cross-Layer Design for Video Streaming over Wireless Ad Hoc Networks," *Multimedia Sig. Proc. Wksp.*, Sienna, Italy, Oct. 2004, pp. 99–102.

[11] P. A. Chou and Z. Miao, "Rate-Distortion Optimized Streaming of Packetized Media," Microsoft Research tech. rep. MSR-TR-2001-35, Feb. 2001.

[12] E. Setton, X. Zhu, and B. Girod, "Congestion-Optimized Scheduling of Video over Wireless Ad Hoc Networks," *IEEE Int'l. Symp. Circuits and Sys.*, to appear, 2005.

[13] M. Handley *et al.*, "TCP Friendly Rate Control (TFRC): Protocol Specification," RFC 3448, Jan. 2003.

[14] E. Setton, X. Zhu, and B. Girod, "Minimizing Distortion for Multipath Video Streaming over Ad Hoc Networks," *Int'l. Conf. Image Proc.*, Singapore, Oct. 2004, pp. 1751–54.

Biographies

ERIC SETTON (esetton@stanford.edu) received a B.S. degree from Ecole Polytechnique, Palaiseau, France, in 2001 and an M.S. degree in electrical engineering from Stanford University in 2003. He is currently a Ph.D. candidate in the Department of Electrical Engineering of Stanford University, California, and is part of the Image, Video and Multimedia Systems group. Multimedia communication in wired and wireless networks, video compression, and image processing are his main research interests. In 2001 he received the Carnot fellowship and the SAP Stanford Graduate fellowship. In 2003 and 2004 he received the Sony SNRC fellowship. He has spent time in industry in France at SAGEM and in the United States at HP Labs and Sony Electronics. He has four patents pending.

TAESANG YOO [S'04] (yoots@systems.stanford.edu) received a B.S. degree with honors in electrical engineering from Seoul National University, Korea, in 1998, and an M.S. degree in electrical engineering from Stanford Universit in 2003, and is currently working toward a Ph.D. degree at the same university. In summer 2002 he was an intern at Qualcomm, Campbell, California. From 2000 to 2001 he was with Xeline, Seoul, Korea, where he worked on the design of power line communications chipsets. He was also an engineer at DSI, Seoul, Korea, from 1998 to 2000. His research interests include multiple antenna systems, wireless ad hoc networks, and communication theory.

XIAOQING ZHU (zhuxq@stanford.edu) received a B.E. degree in electronics engineering in 2001 from Tsinghua University, China. She received an M.S. degree in electrical engineering in 2002 from Stanford University, California, where she is currently pursuing a Ph.D. degree. Her research interests include video streaming over ad hoc wireless networks and light field compression.

ANDREA GOLDSMITH [F] (andrea@systems.stanford.edu) is an associate professor of electrical engineering at Stanford, and was previously an assistant professor at California Institute of Technology (Caltech). She has also held positions in industry at AT&T Bell Laboratories and Maxim Technologies. Her research includes work in capacity of wireless channels and networks, energy-constrained wireless communications, wireless communications for distributed control, and cross-layer design of wireless networks. She serves on the Board of Governors for the IEEE Information Theory Society, and as Vice-Chair for the Communication Theory Committee of IEEE ComSoc. She received her B.S., M.S., and Ph.D. degrees in electrical engineering from the University of California at Berkeley.

BERND GIROD [F] (bgirod@stanford.edu) is a professor of electrical engineering and (by courtesy) computer science in the Information Systems Laboratory of Stanford University. He was Chaired Professor of Telecommunications in the Electrical Engineering Department of the University of Erlangen-Nuremberg from 1993 to 1999. His research interests are in the areas of networked media systems and video signal compression. Prior visiting or regular faculty positions include MIT, Georgia Tech, and Stanford. He has been involved with several startup ventures as founder, director, investor, or advisor, among them Vivo Software, 8×8), and RealNetworks. He also serves as Chairman of the recently launched Deutsche Telekom Laboratories, Berlin. He received an engineering doctorate from the University of Hannover, Germany, and an M.S. degree from Georgia Institute of Technology.