Title
Competitive Equilibrium Bitrate Allocation for Multiple Video Streams

Journal
IEEE Transactions on Image Processing, 19(4)

ISSN
1057-7149 1941-0042

Authors
Tiwari, M.
Groves, T.
Cosman, P. C.

Publication Date
2010-04-01

DOI
10.1109/TIP.2009.2038777

Peer reviewed

Competitive Equilibrium Bitrate Allocation for Multiple Video Streams

Mayank Tiwari, *Student Member, IEEE*, Theodore Groves, and Pamela C. Cosman, *Fellow, IEEE*

***Abstract*—We consider the problem of simultaneous bitrate allocation for multiple video streams. Current methods for multiplexing video streams often rely on identifying the relative complexity of the video streams to improve the combined overall quality. In such methods, not all the videos benefit from the multiplexing process. Typically, the quality of high motion videos is improved at the expense of a reduction in the quality of low motion videos. In our approach, we use a competitive equilibrium allocation of bitrate to improve the quality of all the video streams by finding trades between videos across time. A central controller collects rate-distortion information from each video user and makes a joint bitrate allocation decision. Each user encodes and transmits his video at the allocated bitrate through a shared channel. The proposed method uses information about not only the differing complexity of the video streams at every moment but also the differing complexity of each stream over time. Using the competitive equilibrium bitrate allocation approach for multiple video streams, simulation results show that all the video streams perform better or at least as well as with individual encoding. The results of this research will be useful both for ad hoc networks that employ a cluster head model and for cellular architectures.**

***Index Terms*—Competitive equilibrium, Edgeworth box, H.264/AVC, Pareto optimality, rate control, video compression.**

I. Introduction

With the advancement of digital video technology in recent years, there has been an enormous surge in the amount of video data shared across networks. In many cases, a transmission link is shared by more than one video stream. Applications where multiple compressed video streams are transmitted simultaneously through a shared channel include direct broadcast satellite (DBS), cable TV, video-on-demand services, and video surveillance. Some more common applications are YouTube and instant video streaming by content providers, such as Netflix, where multiple streams are transmitted simultaneously, and in many cases, these streams share a common transmission channel. In such cases, it has been shown that joint bitrate allocation schemes for multiple streams can perform better than an equal bitrate allocation.

For some existing methods of transmitting multiple video streams [1]–[4], improving the overall or average quality is the goal. Overall quality improvement can be achieved by exploiting the relative complexity of different video streams at every moment. However, not all video streams may benefit from the multiplexing process. Generally, the quality of high complexity videos improves at the expense of reduced quality of low complexity videos. Three transform domain multiplexing methods were discussed in [1]: MINAVE, MINVAR, and S-MINVAR. The overall quality when averaged across all the videos was maximized in the MINAVE method but at the expense of reducing the quality for some videos. At the cost of a peak signal-to-noise ratio (PSNR) reduction, the MINVAR method in [1] was proposed to reduce the frame level video quality variance between various video streams. Further, using an encoder buffer, S-MINVAR in [1] was proposed to reduce the quality variance across both videos streams and across frames. While the bit allocation algorithm in [1] is a good method for minimizing the quality variance between the videos, it comes at the expense of substantially reducing the average video quality, and also not all the users improve their individual video quality.

Mechanism-based resource allocation for multiple video streams was studied in [5] and [6]. These methods use a central controller for resource allocation. In [5], a Groves mechanism was used to control a network comprised of selfish users. Under the mechanism, a user's cost for receiving his share of resources depends on the information transmitted to the central controller, and it was shown that a user will report his true values for receiving his allocated share of resources. However, the overall quality of the system was improved at the expense of reductions of video quality for some users. A bandwidth resource allocation procedure using the Nash and Kalai-Smorodinsky bargaining solutions was proposed in [6] for multiple collaborative users, and some important properties of the bargaining solution were presented for effective multimedia resource allocation. It was shown in [6] that the utility of all the users increased compared to a disagreement point defined by an initial resource allocation such that the initial utility is zero for all users. However, not all the available resource is allocated at this disagreement point. However, if the disagreement point is defined by any arbitrary allocation of the *total* available resource, then the method given in [6] will converge only to the disagreement point.

Manuscript received May 20, 2009; revised November 03, 2009. First published December 18, 2009; current version published March 17, 2010. This work was supported in part by the Office of Naval Research, in part by the National Science Foundation under Grant 0635165, in part by the Center for Wireless Communications at UC San Diego, and in part by the UC Discovery Grant program. Part of this work was presented at the IEEE 42nd Asilomar Conference on Signals, Systems and Computers, October 2008. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Antonio Ortega.

M. Tiwari and P. C. Cosman are with the Department of Electrical and Computer Engineering, University of California, San Diego, La Jolla, CA 92093-0508 USA (e-mail: mayank@ece.ucsd.edu; pcosman@ucsd.edu).

T. Groves is with the Department of Economics, University of California, San Diego, La Jolla, CA 92093-0508 USA (e-mail: tgroves@ucsd.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/TIP.2009.2038777

All the above methods use a central controller for resource allocation. Some decentralized resource allocation methods were proposed in [7] and [8]. The Vickrey-Clarke-Groves mechanism was used to allocate resources in [7] and a pricing mechanism for resource allocation was used in [8]. In both these methods, the emphasis was on inducing truth-telling behavior from various users based on a mechanism that adjusts the trade-off between the value of additional resources and a cost levied to induce honest information transmittal. It is interesting to note that, in [8], a very special utility function is assumed for all users where a user's utility is a linear function of the video quality and "money", which implies every user values an incremental change in his video quality exactly the same—a very strong and arguably unrealistic assumption. Moreover, in [8], an improvement in each user's utility was shown compared to the initial utility of zero, however, not all users would improve their utility if the total amount of resources were initially allocated, similar to [6].

None of the methods discussed above aims to improve the quality of *all* the users individually. If the alternative is an equal distribution of the total available resource, a user may see his video quality degrade if he participates in any one of the resource allocation methods described above. Thus, it is reasonable to expect that a user might choose not to participate in a specific resource allocation process if he is unsure of his video quality improvement. The user might decide he will be better off acting individually and receive a fixed share of the resource.

In this paper, the goal is that no video stream will suffer quality degradation by participating in the multiplexing process, compared to independent encoding. As we will see later, while the method does not offer a guarantee that no video will suffer a quality degradation from multiplexing compared to individual encoding, in practice a quality degradation is extremely unlikely, and did not occur at all in our experimental cases. We use a competitive equilibrium approach for bitrate allocation among various video users based on their video complexity. A competitive equilibrium consists of an allocation and a price vector at which (a) each user's allocated consumption vector maximizes his utility given a budget constraint defined by the equilibrium prices and his initial wealth, and (b) at the equilibrium prices, the aggregate supply of each resource that was endowed initially to the users equals the aggregate demand for it.

For video multiplexing, we define the utility in terms of mean squared error (MSE), and the resource at any time-slot is the available bitrate. If the video complexity is high then more of the resource is required to attain some level of utility compared to the amount required for achieving the same level of utility for a less complex video. The equilibrium price is the rate of exchange between current bitrate and expected future bitrate. As will be discussed later, it reflects the number of current bits a user must give up to get one expected bit some time in the future or equivalently, how many current bits a user can get by giving up one expected bit some time in the future. The price varies with the relative video complexity between the current time-slot and future time-slots. If the average complexity of all the videos in the current time-slot is greater than the expected average complexity in the future, then the bits at the current time-slot are more valuable and the price will be greater, and vice versa if the average complexity of all the videos in the current time-slot is less than the expected average complexity of the videos in the future.

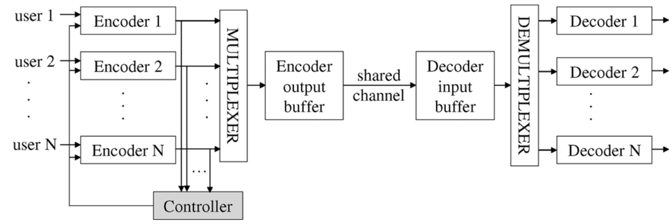

Fig. 1. Bitrate allocation for multiple video streams using a central controller.

The general equilibrium approach views the economy as a closed and interrelated system in which the equilibrium values for all the variables are determined simultaneously. Our method for implementation of competitive equilibrium selects an expected efficient, or Pareto optimal, allocation of bitrate for multiple videos. At a Pareto optimal solution, there is no alternative way of allocating resources that makes some users better off without making some other users worse off. By computing the expected competitive allocation in the Edgeworth box, a common tool in economics for equilibrium analysis, we find a point where all users, in expectation, perform better or at least as well as what they could achieve independently. This method exploits gains in quality that can be achieved by trading bits across time rather than merely reallocating bits among the video streams at any time, as is done with current methods of multiplexing. In our preliminary work [9], the competitive equilibrium at any time-slot was solved by reducing the entire video sequence to two equal length time-slots, one for the current time-slot and one for the average of all the remaining time-slots. In this paper, we trade the bits for the current time-slot against all the remaining time-slots [with the same expected rate-distortion (RD) curve] and this gives more flexibility in trading.

Fig. 1 shows a general block diagram for joint bitrate allocation for multiple video users. Each user passes RD information to a central controller. The central controller computes the competitive equilibrium for the users simultaneously and sends the allocated bitrate information to the corresponding encoder. Each encoder uses this information to encode his video. Encoded bitstreams are multiplexed and transmitted through a shared channel to the decoder. At the decoder, the bitstreams are demultiplexed and sent to the corresponding decoder to decompress the bitstream and display. We operate the competitive equilibrium bitrate allocation at the level of a group of pictures (GOP). However, it can be operated at other granularities.

In this work, we ignore any incentive the users might have to communicate false information in order to acquire additional resources. For example, a user with low complexity video in the current time-slot and high complexity in the next time-slot might overstate his current complexity relative to the next time-slot, thereby lowering the relative bitrate price in the next time-slot, and, thus, enabling more favorable allocations in the current and next time-slots. The mechanism we discuss here is susceptible to this type of strategic manipulation.

However, such incentive issues are not relevant for applications in which separate users share the same objective (i.e., form a team [10], for example, as in a military scenario). Also, as is extensively discussed in the economic literature, the advantages to such strategic manipulation become vanishingly small as the number of users being multiplexed increases. The larger the number of users, the less any individual user can affect the computed competitive equilibrium prices. Therefore, we assume all the users inform the controller of their true video RD characteristics. The controller then uses the competitive equilibrium approach to determine the bitrate allocation.

The rest of the paper is organized as follows. Section II describes the Edgeworth box for illustrating competitive equilibria. Section III describes competitive equilibrium bitrate allocation methods for multiple video streams. Results are given in Section IV, and Section V concludes the paper.

II. Edgeworth Box for Competitive Equilibrium

In this section, we briefly describe a competitive equilibrium and its Edgeworth box representation. Interested readers are encouraged to read [11] for further details.

The *Edgeworth box* [12] is a graphical tool for exhibiting Pareto optimal allocations and illustrating a competitive (Walrasian) equilibrium in a pure exchange economy [11], in which no production is possible and the commodities that are ultimately consumed are those that individual users possess as initial endowments. The users trade these endowments among themselves in a market for mutual advantage. For the competitive equilibrium analysis, we start with an example using two users who exchange quantities of two goods with each other for their mutual advantage. This simple case is amenable for graphical analysis using the tool known as an Edgeworth Box. Later, for multiplexing many video streams, we will apply the theory of competitive equilibrium for exchanging bitrate to improve the quality of all the video streams.

Consider two users ($i = 1, 2$) and two goods ($j = 1, 2$). User i's consumption vector is $x_i = (x_i^1, x_i^2)$, i.e., user i's consumption of good j is $x_i^j \geq 0$. Each user i is initially endowed with an amount $c_i^j \geq 0$ of good j. The total endowment of good j in the economy is denoted by $c^j = c_1^j + c_2^j$, assumed to be strictly positive. An allocation $x \in \mathbf{R}_+^4$ is an assignment of a non-negative consumption vector to each user: $x = \{x_1, x_2\} = \{(x_1^1, x_1^2), (x_2^1, x_2^2)\}$. We say that an allocation is *nonwasteful* and *feasible* if $x_1^j + x_2^j = c^j$ for all goods j (the total consumption of each good is equal to the economy's aggregate endowment of it).

In the Edgeworth box, user 1's quantities are measured with the southwest corner as the origin (O_1), shown in Fig. 2. User 2's quantities are measured using the northeast corner as the origin (O_2). For both the users, the horizontal dimension measures quantities of good 1 and the vertical dimension measures quantities of good 2 from their respective origins. The width and height of the box are c^1 and c^2, the economy's total endowment of goods 1 and 2. The initial endowment point is given by $c = \{(c_1^1, c_2^1), (c_1^2, c_2^2)\}$. Any point x in the box represents a division of the total endowment between users 1 and 2, as shown in Fig. 2.

User i's wealth is defined by the market value of his goods endowed initially. Suppose users can buy or sell these goods in the market for prices p^1 and p^2. In general equilibrium theory, the wealth of a user is derived internally by the value of the

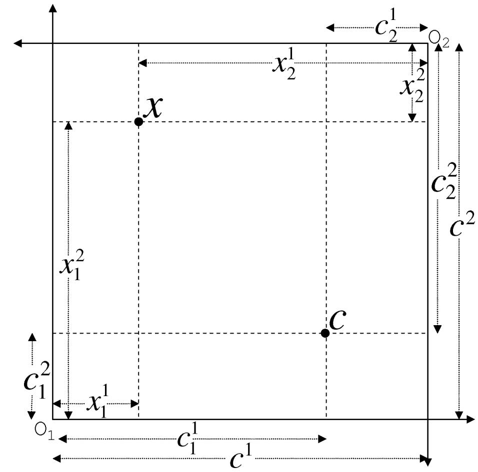

Fig. 2. Edgeworth box for two users and two goods.

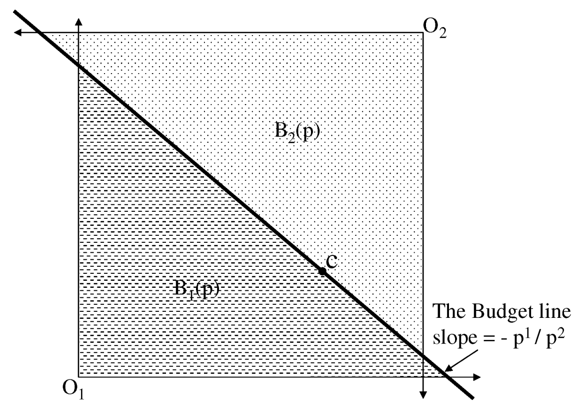

Fig. 3. Budget sets in an Edgeworth box for two users and two goods.

prices. For any price system $p = (p^1, p^2)$ and initial endowment, the budget set for user i is

$$B_i(p) = \left\{x_i \in \mathbf{R}_+^2 : p^1.x_i^1 + p^2.x_i^2 \leq p^1.c_i^1 + p^2.c_i^2\right\}. \quad (1)$$

The budget set for user i is the set of all consumption vectors x_i which user i can afford at price p. The budget sets for two users in an Edgeworth box are shown in Fig. 3. A line drawn through the initial endowment with a slope of $-(p^1/p^2)$ is the *budget line*. User 1's budget set (denoted by $B_1(p)$) consists of all the non-negative vectors below and to the left of the budget line. The area on the other side of the budget line is the budget set for user 2 (denoted by $B_2(p)$). Any total allocation of the two goods on the budget line will be affordable at price system p to both the users simultaneously.

Given $c_i = (c_i^1, c_i^2)$, user i can calculate his utility $U_i(c_i)$, which is a measure of "goodness" or satisfaction with the consumption vector c_i. The locus of all x_i yielding the same utility $U_i(x_i) = u_i$ is called an *indifference curve* of user i. The family of all indifference curves is the collection of level sets of the utility function $U_i(x_i)$, as shown in Fig. 4. As we move away from his origin, the utility associated with successive indifference curves for user i increases, as more of both goods increases utility. In Fig. 4, considering the indifference curves for user 1, the utility for user 1 that is associated with the indifference curve that passes through $c1$ is higher than the utility associated with

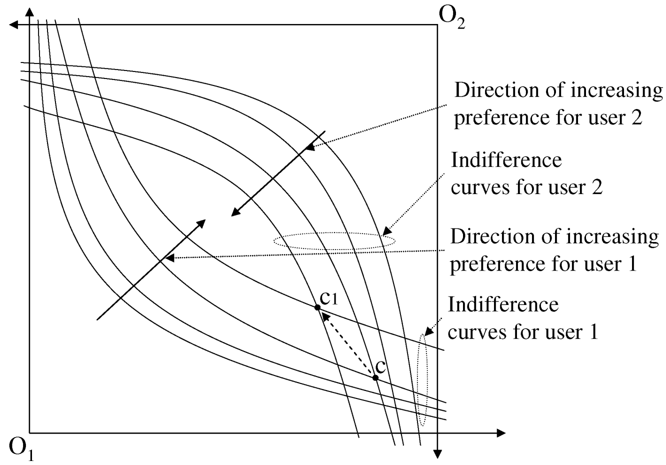

Fig. 4. Preferences in the Edgeworth box.

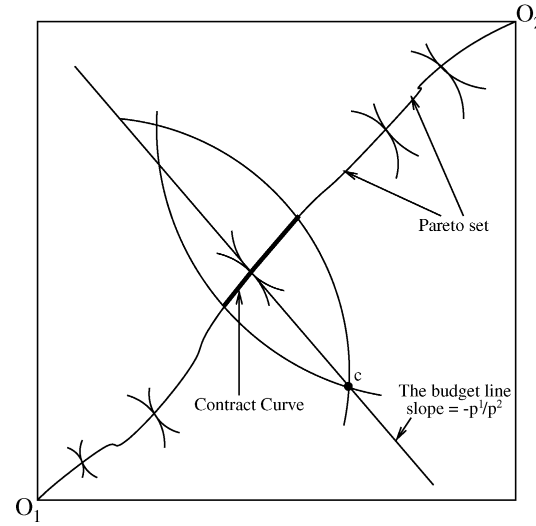

Fig. 5. Pareto set and the contract curve in the Edgeworth box.

the indifference curve through c. The same is true for user 2. For the purpose of explaining the Edgeworth box, we assume that these curves are convex. For the problem of allocating bitrate among multiple video streams, we will discuss the validity of this assumption in the next section.

Consider the family of indifference curves for each user, consisting of all indifference curves through every allocation. Under our assumptions of convexity and smoothness, each indifference curve for a user will be tangential to one indifference curve of the other user at some point. The points where the indifference curves for both users are tangential to each other are *Pareto optimal* allocations. At these allocations, it is not possible to increase the utility of one user without decreasing the utility of the other user. The set of all Pareto optimal allocations is the *Pareto set*. It is a curve that connects all the Pareto optimal allocation points in the Edgeworth box from one origin to the other origin. The part of the Pareto set where both users do at least as well as at their initial endowments is called the *contract curve* (Fig. 5). The contract curve lies between the indifference curves for both the users passing through the initial endowment. Free and unfettered bargaining between the users might be expected will result in some point on the contract curve as these are the only points at which both users do at least as well as at their initial endowments and for which there is no alternative further trade that can make both users better off [11]. This is shown in Fig. 5.

Given a price vector p, a user demands his most preferred allocation in his budget set. The most preferred point is the point where the budget line is tangential to one of his preference curves. As p is varied, the budget line pivots around the initial endowment point c, and the quantity demanded by a user will be a set of points where different budget lines are tangent to different preference curves. The locus of the user i's optimal choice given current price (which defines current wealth) is known as the *offer curve* (OC_i) and is shown in Fig. 6. The offer curve always passes through the endowment point because, for any p, the initial endowment is affordable for the user and the tangent to the indifference curve through the endowment point defines a price p at which the endowment is the most preferred allocation in the corresponding budget set.

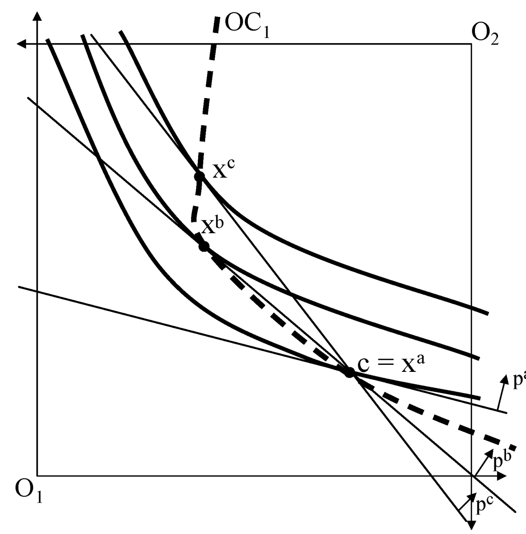

Fig. 6. Offer curve for user 1. Budget line p^a, p^b, and p^c are tangential to various indifference curves at x^a, x^b, and x^c, respectively.

A *competitive equilibrium* for an Edgeworth box economy is a price vector p^* and an allocation $x^* = \{x_1^*, x_2^*\}$ such that

$$U_i\left(x_i^*\right) \geq U_i\left(x_i'\right) \quad \forall x_i' \in B_i(p^*),\ \forall i = 1, 2 \tag{2}$$

and

$$\sum_{i=1}^{2} x_i^{j*} = c^j\ \forall j = 1, 2. \tag{3}$$

At an equilibrium, each user i's demanded bundle at price vector p^* is x_i^* and each user's net demand for a good is exactly matched by the other's net supply. The intersection of a budget line and contract curve, where the budget line is also tangential to the indifference curve for both the users on the contract curve, defines a competitive equilibrium. At this point, the offer curves for both the users intersect at a point on the contract curve (Fig. 7). At an equilibrium point, both users are better off compared to their initial endowment. Under our assumptions, at least one competitive equilibrium will exist for every initial

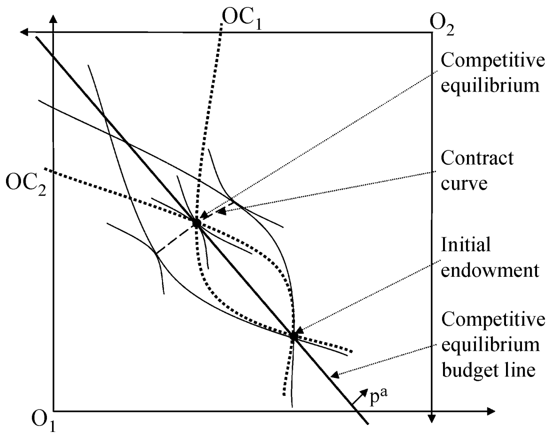

Fig. 7. Competitive equilibrium allocation. The offer curves for both the users intersect at the competitive equilibrium allocation. One of the indifference curves for both the users is tangential to a budget line at this allocation.

endowment allocation. More details about the Edgeworth box and competitive equilibrium can be found in [11].

III. Competitive Equilibrium Approach for Video Multiplexing

In this section, we explain the competitive equilibrium approach for bitrate allocation among various video streams. Suppose there are N video users. The video stream of each user is divided into T time-slots (TS). In this work, we consider one TS to be one GOP, but a TS can be larger or smaller than a GOP. We will use the terms GOP and TS interchangeably. We assume that the videos are synchronized at the GOP level. Such synchronization can be achieved by a small amount of buffering of the input videos at the expense of a small amount of delay. The problem of synchronization does not exist when a TS is of a frame size.

We extend the concept of the Edgeworth box from two users to N users and from two goods to T TS. The quantity of each good is represented by the number of bits available in each TS. A user i in TS t is initially endowed with c_i^t bits. Therefore, the total number of bits available in TS t is

$$c^t = \sum_{i=1}^{N} c_i^t. \tag{4}$$

The users compete to receive bits from the pool of c^t bits in TS t. Let $U_i(x_i^1, x_i^2, \ldots, x_i^T)$ be the utility for user i. Therefore, the optimization problem for user i is given by

$$\max_{\{x_i^t\}} U_i\left(x_i^1, x_i^2, \ldots, x_i^T\right) s.t. \sum_{t=1}^{T} p^t.x_i^t = \sum_{t=1}^{T} p^t.c_i^t \tag{5}$$

and the constraint over all users is

$$\sum_{i=1}^{N} x_i^t = c^t \quad \forall t = 1 \; to \; T \tag{6}$$

where $(p^1, p^2, \ldots, p^T)$ are the competitive equilibrium prices. The solution $\{x_i^{t*}\}$ obtained from (5) and (6) is the competitive equilibrium bitrate allocation.

The above method would work for archived videos where the utility function over all TS is available in advance. However, for real-time applications, a limited amount of video is available and it is generally not possible to achieve a global competitive equilibrium solution. Thus, we reduce the problem to a sequence of problems, each of which solves for a competitive equilibrium for the current TS and a representative of all future TS. For video streams, we define the utility to be the negative of MSE. We generate the RD curve for each TS by calculating the MSE at different bitrates. Note that the complexity of generating the RD curve can be reduced by using the method described in [13].

Suppose that, in a two user system, user 1 and user 2 each has an initial endowment of 500 bits in TS 1 and in TS 2. Therefore, a total of 1000 bits are available in each TS. If the RD curves for the two users are such that giving 600 bits to user 1 in TS 1 and 400 in TS 2 (and vice versa for user 2) produces a more favorable total MSE than the equal initial endowment, then the Edgeworth box approach would favor this allocation over the initial one.

While trading across TS is the basic idea behind our approach, often adjacent TS have similar RD curves. Therefore, little benefit can be gained by trading bits between adjacent TS for any two users. One would like to trade between the current encoding TS and some other TS widely separated in time. But, since the specific RD curve for a distant TS is typically not known in a real-time application, we consider trades between the current encoding TS and an *expected* or *representative* RD curve for all the future TS.

Specifically, for our method, in each TS for each user, the central controller will reoptimize the decision for the current and all future TS using estimated values for future RD curves. Assuming future TS are identical in expectation (the future environment is perceived as stationary), then each TS's decision problem is just an optimization problem with two decisions only—the allocation x_i^t for the current TS and $\bar{x}_i^t$, the common allocation for each of the remaining $(T-t)$ TS. We start at the first TS and sequentially process each TS in the same manner. Using the estimated utility for the future TS, the optimization problem for user i becomes

$$\max_{x_i^t, \bar{x}_i^t} U_i\left(x_i^t, \bar{x}_i^t\right)$$
$$s.t. \quad p^t.x_i^t + (T-t).\bar{p}^t.\bar{x}_i^t = p^t.c_i^t + (T-t).\bar{p}^t.\bar{c}_i^t \tag{7}$$

where p^t is the equilibrium price for the current TS and $\bar{p}^t$ is the estimated equilibrium price of the remaining TS. $\bar{c}_i^t$ is the average initial endowment for user i for TS $t+1$ to T.

As given in [14], the RD curve for user i in TS t is fitted by

$$D_i^t\left(R_i^t\right) = a_i^t + \frac{b_i^t}{R_i^t + d_i^t} \tag{8}$$

where R_i^t is the number of bits and D_i^t is the MSE distortion for TS t in video stream i. We use the unconstrained nonlinear minimization approach to find a_i^t, b_i^t, and d_i^t, the coefficients for generating this curve-fitting model. Other curve-fitting models are available in the literature [13]. Note that the model used in (8) is convex. The convexity of the RD curves is an empirical observation. Frequently in the previous literature, convexity is either empirically observed or is assumed to hold. All the videos investigated in this paper exhibit this property. Were it not to be the case, the computed competitive equilibrium solution in the Edgeworth box would not necessarily be an efficient solution for bitrate allocation.

Using (8), the utility function is represented by

$$U_i\left(x_i^t, \bar{x}_i^t\right) = -\left(a_i^t + \frac{b_i^t}{x_i^t + d_i^t}\right) - (T-t).\left(\bar{a}_i^t + \frac{\bar{b}_i^t}{\bar{x}_i^t + \bar{d}_i^t}\right) \tag{9}$$

that is, the negative sum of the MSE in TS t and the estimated weighted MSE for the remaining $(T-t)$ TS. Then the indifference curve through the initial endowment at TS t can be derived as

$$-\left(a_i^t + \frac{b_i^t}{x_i^t + d_i^t}\right) - (T-t).\left(\bar{a}_i^t + \frac{\bar{b}_i^t}{\bar{x}_i^t + \bar{d}_i^t}\right) = -\left(a_i^t + \frac{b_i^t}{c_i^t + d_i^t}\right) - (T-t).\left(\bar{a}_i^t + \frac{\bar{b}_i^t}{\bar{c}_i^t + \bar{d}_i^t}\right) \tag{10}$$

for different combinations of x_i^t and $\bar{x}_i^t$. Since the RD curves for both current TS and average for future TS are convex; these indifference curves are also convex in nature.

A competitive equilibrium is found by solving

$$\max_{x_i^t, \bar{x}_i^t} -\left(a_i^t + \frac{b_i^t}{x_i^t + d_i^t}\right) - (T-t).\left(\bar{a}_i^t + \frac{\bar{b}_i^t}{\bar{x}_i^t + \bar{d}_i^t}\right)$$
$$s.t.\ p^t.x_i^t + (T-t).\bar{p}^t.\bar{x}_i^t = p^t.c_i^t + (T-t).\bar{p}^t.\bar{c}_i^t,\ \forall i = 1\ to\ N. \tag{11}$$

The Lagrangian expression for user i is

$$L_i = -\left(a_i^t + \frac{b_i^t}{x_i^t + d_i^t}\right) - (T-t).\left(\bar{a}_i^t + \frac{\bar{b}_i^t}{\bar{x}_i^t + \bar{d}_i^t}\right) + \lambda_i\left(p^t.c_i^t + (T-t).\bar{p}^t.\bar{c}_i^t - p^t.x_i^t - (T-t).\bar{p}^t.\bar{x}_i^t\right). \tag{12}$$

The constraints on the total available bits in each TS are given by

$$\sum_{i=1}^{N} x_i^t = c^t \tag{13}$$

$$\sum_{i=1}^{N} \bar{x}_i^t = \bar{c}^t. \tag{14}$$

By differentiating L_i with respect to x_i^1, x_i^2, and λ_i, equating the results to 0 and solving for x_i, and substituting in (13), we get

$$\sum_{i=1}^{N} \sqrt{\frac{b_i^t}{p^t}}.\left(\frac{p^t.\left(c_i^t + d_i^t\right) + (T-t).\bar{p}^t.\left(\bar{c}_i^t + \bar{d}_i^t\right)}{\sqrt{p^t.b_i^t} + (T-t).\sqrt{\bar{p}^t.\bar{b}_i^t}}\right) - \sum_{i=1}^{N} d_i^t = c^t. \tag{15}$$

To determine the competitive equilibrium, we need to find the equilibrium prices, p^t and $\bar{p}^t$ that solve (15). Since the solution of (15) is homogeneous of degree 0 in prices, we only need to find an equilibrium price ratio $p^t/\bar{p}^t$. Therefore, without loss of generality, we may take $\bar{p}^t = 1$ and solve (15) numerically for p^t. With p^t, we find x_i^t and $\bar{x}_i^t$ which comprise a competitive equilibrium.

To predict the future average RD function, we consider several alternatives that differ in the information assumed to be held by the user at the time he makes the forecast. In all cases, the user will use the competitive equilibrium approach to calculate the allocated bitrate for the current TS with respect to the forecasted average RD function for the future. Suppose we have information about the average RD function for a video user i over *all* TS (1 to T). Then we can use such information to calculate the bitrate demand at a competitive equilibrium for a user in the current TS by trading the bits with the average RD function for all the TS. We call this method of bitrate allocation **ALL_TS** which assumes knowledge of the average RD function over all the TS for a user. The average RD curve in ALL_TS is approximated by averaging the individual coefficients (a, b, and d) separately for a user over all TS. The coefficients generated by actually averaging the RD curves for all the TS of a video are extremely close to the average coefficients.

A user is assumed to always know the actual RD function for the current TS t and all past TS (1 to $t-1$). Given the average RD function for all TS, he can calculate the average RD function for the *remaining* TS (**REM_TS**). With the information on the current RD function and average RD function for the remaining TS, the central controller uses (15) to calculate the competitive equilibrium price and bitrate allocation for the current TS. The average RD curve in REM_TS is approximated by averaging the individual coefficients (a, b, and d) separately for a user over all the remaining TS. Both ALL_TS and REM_TS are *ex ante* approximation models where we assume some information about the future video in advance.

Suppose a user has no knowledge about future TS (*ex post*) (as is the real-time case) but assumes that the video is a stationary process at the GOP level. Future TS properties (for example, complexity) can be estimated by looking at the past. We calculate the bitrate demand for the current TS using the average of all *previous* TS (**PRE_TS**) as the estimate of future TS. This method would be expected to work well for long videos but may not work for short videos if the previous TS are very different from the future TS. If averaged over a sufficiently long interval, the complexity for most video streams can be assumed to be almost stationary. The assumption about stationarity is important for PRE_TS to work well. The average RD curve in PRE_TS is approximated by averaging the individual coefficients (a, b, and d) separately for a user over all the past TS.

To approximate an upper bound on video quality improvement, we consider a method in which each user has full information about his RD curves in all TS, and proposes to divide the bits among all the TS based on his relative complexity. This can only be done for archived videos where the coefficients of the RD curves are calculated offline. Each video stream uses this criterion for bit allocation among its TS independently. Since the total number of bits in each TS is given by the initial endowment, we normalize the number of bits allocated to each video stream by the total available bits for a TS (**FUL_TS**). Note that, for this method, each user attempts to allocate bits across TS but does not trade with other users. FUL_TS is, of course, not the real upper bound. The real upper bound would be given by computing the competitive equilibrium for all TS simultaneously [(5) and (6)], an extremely large computational problem if there are many users and many TS.

To compare the improvement in video quality of competitive equilibrium bitrate allocation using the various multiplexing

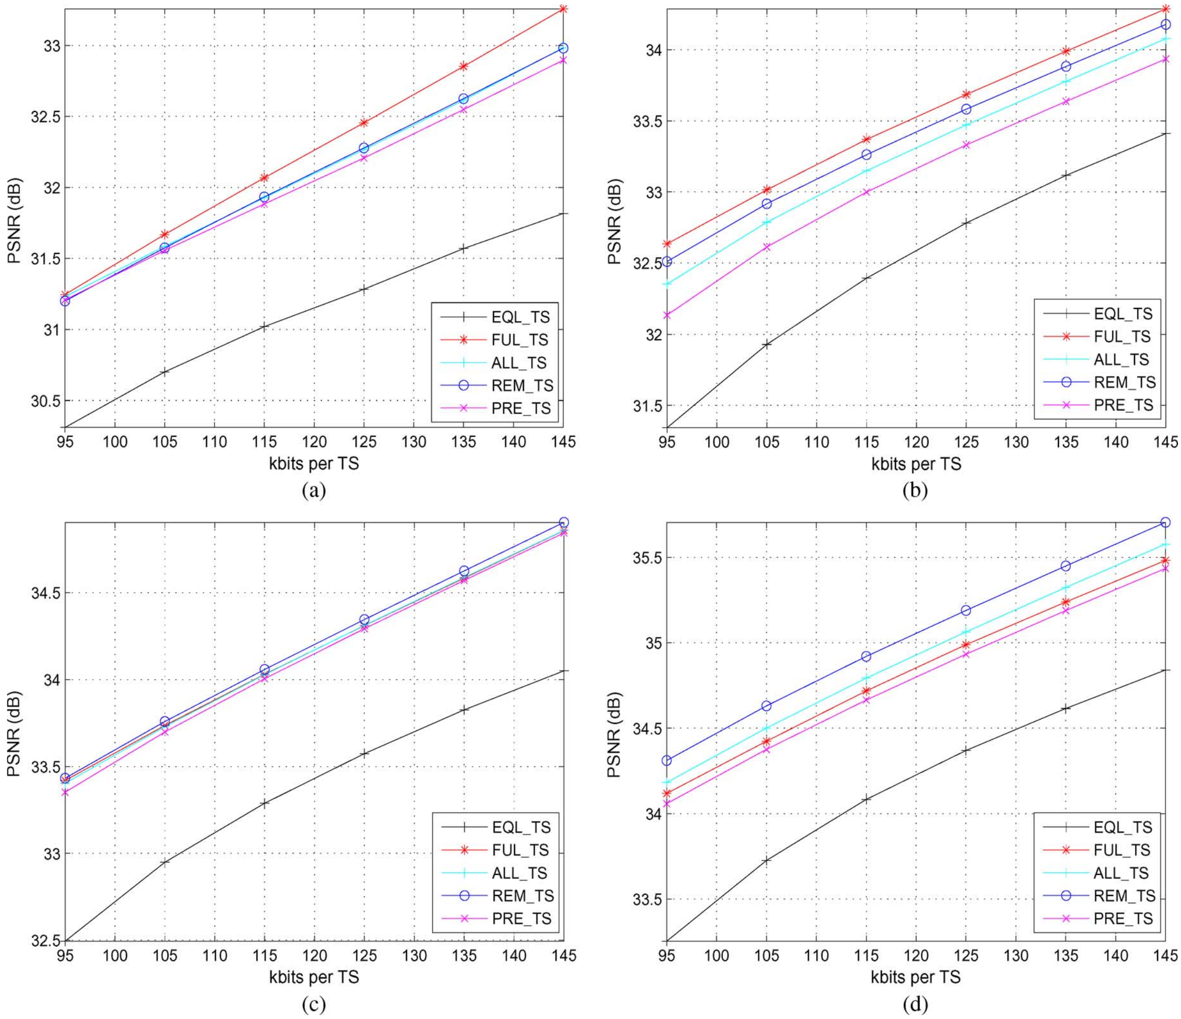

Fig. 8. PSNR variation with bitrate for four multiplexed video streams. (a) g8 video stream; (b) g9 video stream; (c) g10 video stream; (d) g11 video stream.

schemes, we consider the equal bitrate allocation for each user in every TS (**EQL_TS**). Here, each user in every TS receives an equal number of bits to encode his video. Note, for a TS of GOP length, the rate control algorithms used in conjunction with most of the video standards strive to achieve equal bitrate allocation for all GOPs, similar to EQL_TS. We use equal bitrate allocation as an initial endowment for our competitive equilibrium allocation.

IV. Results

The simulation was performed using the baseline profile of H.264/AVC [15] reference software JM 11.0 [16]. The GOP size is 15 frames (I-P-P-P). The frames inside a TS are encoded using H.264 rate control [17]. The test video sequences were taken from travel documentaries at a resolution of 352 $\times$ 240 pixels (SIF) and at 30 frames per second. We chose 12 test sequences (denoted by g1 to g12) and each sequence was 250 seconds (7500 frames) in length. The coding parameters such as resolution, GOP size or structure can be changed for any appropriate application as our multiplexing method is independent of such parameters. We considered a lossless channel for transmitting multiple video streams.

Each video sequence contained various types of scenes with varying camera motions such as zooming and panning. The high motion scenes included dancing, bike racing, and a vegetable market. The low motion scenes showed buildings, maps, sculptures, scenery, etc. Other types of scenes included a flying airplane, showing flowers, people talking, farming, cooking, children playing chess, etc. The videos also had scenes with varying spatial content such as a crowded market, bird's eye view of a city, sky, still water, etc. Each video sequence contained many types of scenes and motions.

Fig. 8 shows the results of multiplexing four video streams. The five curves in each plot represent the various bitrate allocation methods for multiplexing video streams as described previously. Each plot shows PSNR versus bitrate (ranging from 95–145 kbits per TS (190–290 kbps) per user). We calculate the MSE of each frame and average across all frames of a video, then convert it to PSNR. The performance of EQL_TS is worst in all the videos. This is the method used in most video standards for GOP level rate control. For archived videos, the RD curves for all TS are available and we see that FUL_TS performs better than the other methods for most of the videos. The PSNR gain over EQL_TS varies from 0.62–0.87 dB for g11 to 0.94–1.44 dB for g8. However, this method cannot be used for real-time video multiplexing. The bitrate allocation of EQL_TS is considered to be the initial endowment for the competitive equilibrium bitrate allocation methods (ALL_TS, REM_TS, and PRE_TS). If we consider that a user knows the average RD curve for all the TS, then the competitive equilibrium bitrate allocation method, ALL_TS, is used to improve the video quality of all the video streams individually. We found that ALL_TS performs 0.67–1.00 dB for g9 to 0.88–1.17 dB for g8 better than EQL_TS. We see that even a small amount of information such as the average RD curve for the entire video is useful to improve the quality of all the videos.

If a user knows the average RD curve for future TS, this information can be used to improve the video quality, as shown by REM_TS. This method finds a competitive equilibrium point for the current TS when compared to its average RD of the remaining TS. This method improves the quality of each video stream from 0.77–0.94 dB for g10 to 0.87–1.17 dB for g8 over the EQL_TS method. As this allocation method uses the knowledge of the average RD curve for all the future TS, in general, its performance is slightly better than the ALL_TS method.

Finally, we assume that we have no prior knowledge about the video and we estimate the future RD curves by looking at the previous TS. Again we compute the competitive equilibrium for the current TS and the estimated average RD information for future TS based on the average of the previous TS (PRE_TS curve in the figure). This method improves the PSNR from 0.56–0.81 dB for g11 to 0.86–1.08 dB for g8 over EQL_TS. All the competitive equilibrium bitrate allocation methods improve the quality of all the video streams. We see that, even with absolutely no knowledge about the future video RD characteristics in PRE_TS, we are able to improve the quality of all the video streams by calculating a competitive equilibrium bitrate allocation. The PSNR improvement over EQL_TS using this allocation method is in the vicinity of other competitive equilibrium bitrate allocation methods described in this paper.

Our competitive equilibrium bitrate allocation method aims at improving the video quality of each user. If we calculate the MSE averaged across all the users, then our method may perform worse than the methods for minimizing the MSE across all the videos [1], [2], [4]. For example, consider the case for the four videos used in Fig. 8 at 95 kbits per TS per user. If we maximize the quality averaged over all the videos at each TS using MINAVE from [1] applied to time-domain RD curves, then we achieve 32.91 dB as the average PSNR. On the other hand, the average video quality for the ALL_TS, REM_TS, and PRE_TS methods is 32.65, 32.71, and 32.55 dB, respectively. Clearly, our multiplexing method does not minimize the average distortion. For improving the video quality for each user individually, we incur some performance penalty when compared to the method that explicitly has as its goal the minimization of average distortion over all videos.

Estimates of MSE for future TS by our various estimation methods are shown in Fig. 9 for the g11 video sequence at a bitrate of 100 kbits per TS. On the x-axis is the TS index and on the y-axis is the actual MSE and three types of estimated MSE for the future TS. We encode each TS at the given bitrate and this is shown as the "Actual" curve in the figure. When the RD curve is averaged over all the TS (ALL_AVG), then the MSE remains constant. When the RD curve is averaged over the remaining video at any TS then the calculated MSE is the same as ALL_AVG initially, then deviates, and finally converges with the actual MSE of the last TS. When the average RD curve is estimated from the past TS (PRE_AVG), then the calculated MSE starts with the actual MSE at the beginning and then converges to the ALL_AVG at the end. We compare the actual MSE variation at any TS with the MSE of these averages. The important conclusion that can be derived from this figure is the low variation of all the types of averaging methods compared to the variation in the actual MSE at any TS. When a user calculates his bitrate requirement for the current TS, then the user actually calculates the usefulness of the bits currently compared to the future. If the video is less complex in the current TS than the average of the future, then the user demands fewer than the average bits for the current TS in anticipation that he will receive more bits in the future when his complexity is expected to be higher. By trading the bits across time, a user is able to improve his video quality. At any given TS, the demands from some users are less than their average allocation of bits while other users' demands are greater. At a competitive equilibrium allocation for the current TS and the average of their future TS, the expected video quality of all the users is improved.

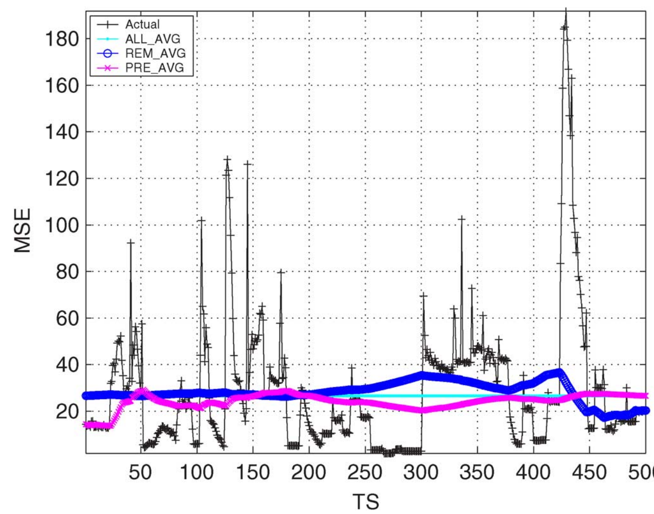

Fig. 9. Actual MSE and estimates of MSE for future TS by our various estimation methods for g11 video sequence at 100 kbits per TS.

Similarly, Fig. 10 shows the results for six video streams that are multiplexed together using the methods described above. The PSNR of all the six videos improves for a wide range of bitrates. For the competitive equilibrium bitrate allocation methods in these six videos, g8 and g9, in general, produce most PSNR improvement. The improvement is as high as 1.5 dB. The least PSNR improvement is seen for g11 but it is still in the range of 0.60–1.08 dB above EQL_TS. A similar result is shown in Table I when ten video streams are multiplexed together at an average bitrate of 100 kbits per TS per user. The table shows the PSNR improvement over EQL_TS by various multiplexing methods.

We note that the largest PSNR gain is achieved by finding the competitive equilibrium when there is a lot of fluctuation in the video motion, for example g9. Conversely, the PSNR gain

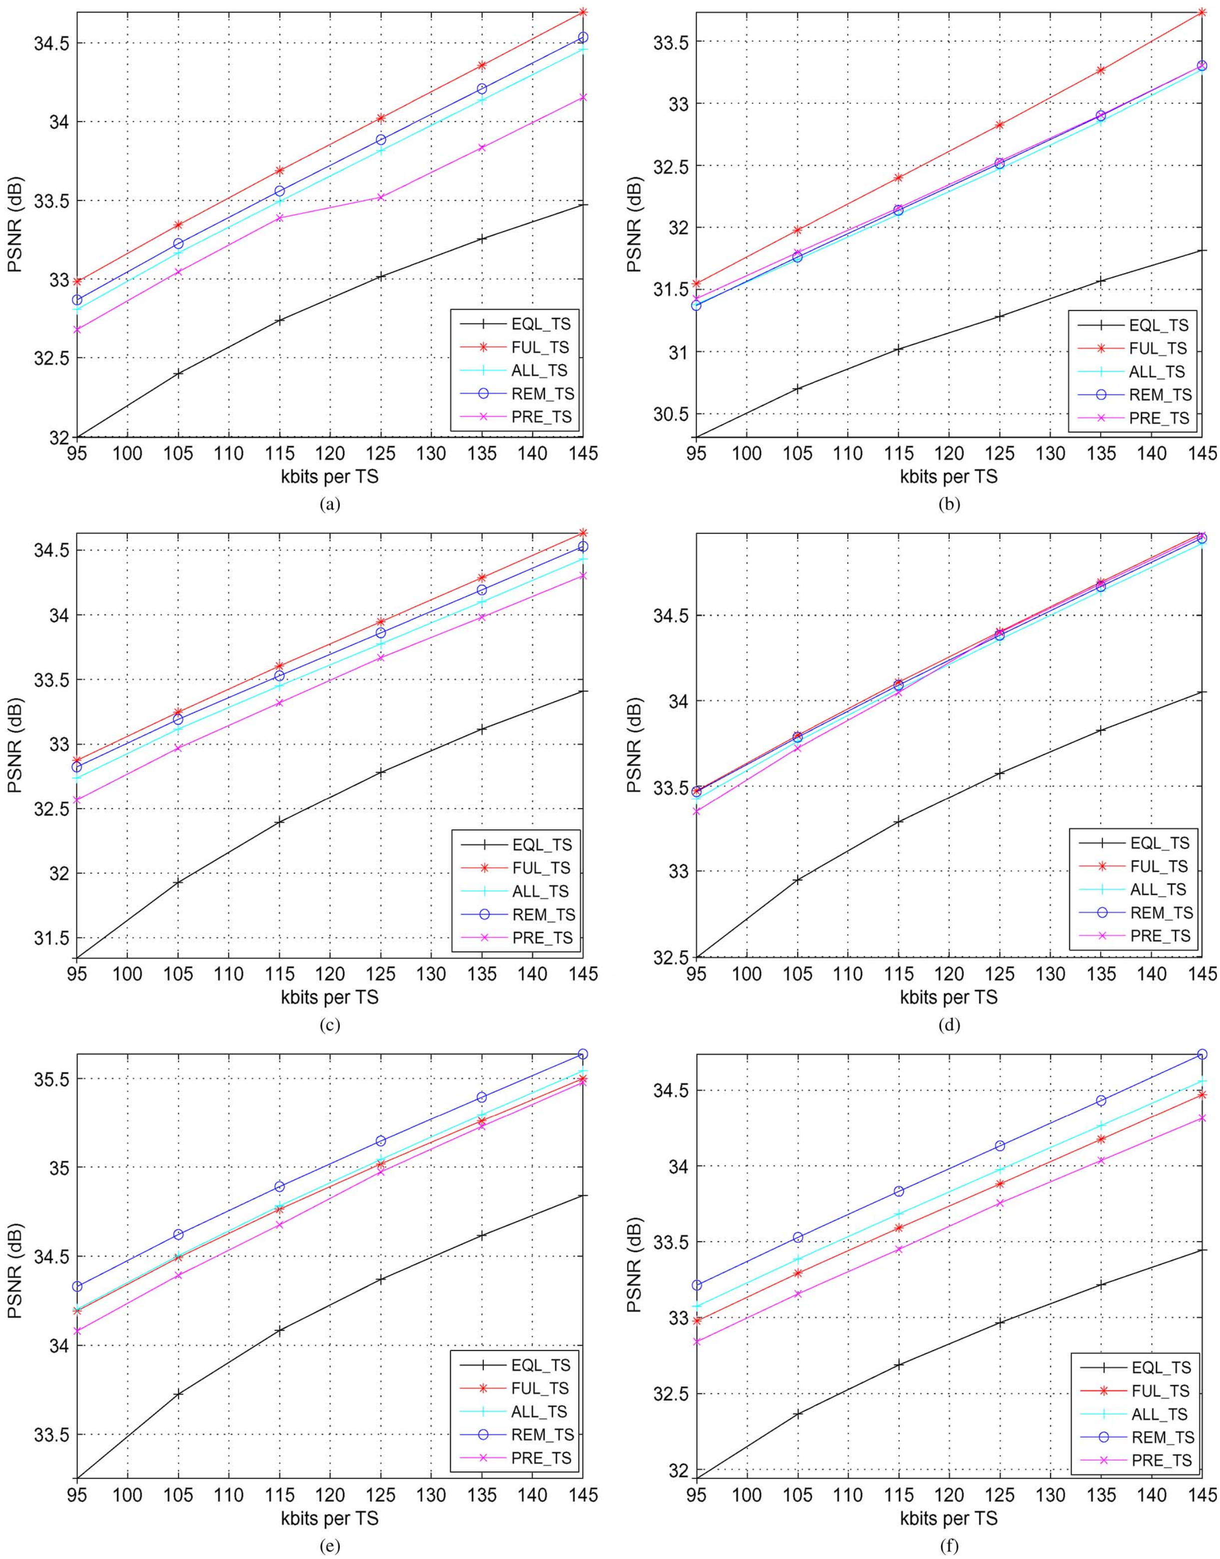

Fig. 10. PSNR variation with bitrate for six multiplexed video streams. (a) g7 video stream; (b) g8 video stream; (c) g9 video stream; (d) g10 video stream; (e) g11 video stream; (f) g12 video stream.

is low if the motion fluctuation in a video stream is low, for example g1. Most of the video streams have significant motion fluctuation and scene changes, so multiplexing them by computing the competitive equilibrium should improve their quality.

The performance of PRE_TS depends on the accuracy of the estimation of future TS from past TS. Suppose we have a video whose complexity is monotonically decreasing. Therefore, it would be desirable if the bitrate demanded in the current TS would be higher than the bitrate demanded in any future TS.

TABLE I
PSNR (db) OF EQL_TS AND IMPROVEMENT OVER EQL_TS BY USING VARIOUS MULTIPLEXING METHODS WHEN TEN VIDEO STREAMS ARE MULTIPLEXED TOGETHER AT AN AVERAGE BITRATE OF 100 kbits PER TS PER USER

	g1	g2	g5	g6	g7	g8	g9	g10	g11	g12
EQL_TS	34.18	35.91	34.45	34.15	32.21	30.51	31.63	32.72	33.47	32.14
	Improvement over EQL_TS									
FUL_TS	0.35	1.77	0.49	0.74	1.18	1.42	1.52	0.98	0.91	1.06
ALL_TS	0.48	0.76	0.62	0.75	1.08	1.17	1.45	0.96	0.99	1.08
REM_TS	0.48	0.87	0.65	0.87	1.09	1.14	1.41	0.93	1.06	1.29
PRE_TS	0.32	0.77	0.44	0.63	0.97	1.03	1.39	0.91	0.85	0.87

Using PRE_TS, the competitive equilibrium bitrate allocation will compare the RD function for the current TS and estimated RD function for the future TS. However, since the estimated RD curve for future TS is based on the past TS, the bitrate allocation for current TS at competitive equilibrium for this video will be less than the average allocation for the future TS. In the REM_TS method, the allocation for the current TS will be more than the average allocation for the future since fewer bits are actually required for the future TS. In such a case, REM_TS and FUL_TS will perform much better than EQL_TS but it is possible that EQL_TS might perform better than PRE_TS. However, in real video examples, we rarely encounter such pathological cases, and our multiplexing methods were found to improve the quality of all of the video streams studied.

As can be seen from Figs. 8 and 10, all the video streams gain from the multiplexing process. The multiplexing method using the competitive equilibrium borrows bits from a low motion TS of a video and gives these bits to another video in the same TS with the expectation of getting back later when the need arises. Thus, the multiplexing method exchanges bits between video streams as well as across the TS. This leads to another observation that the quality fluctuation for each video stream is slightly reduced compared to EQL_TS because the high motion TS get more bits than the low motion TS instead of getting the same number of bits for all TS. Fig. 11 shows the PSNR fluctuation for g9 for all the multiplexing methods. The EQL_TS method has the highest PSNR fluctuation (20. 31–44.83 dB) and FUL_TS has the lowest (24.29–42.50 dB). Among the competitive equilibrium bitrate allocation methods, REM_TS has minimum PSNR fluctuation (22.15–42.50 dB) compared to ALL_TS or PRE_TS. The quality fluctuation can further be reduced by imposing a constraint on maximum and minimum video quality. However, any method of reducing the quality variance comes at the cost of reduction in overall quality as can be seen in [1]. In our paper, reduction in the quality variance is achieved by trading the bits across time. Further work can be done in reducing the quality variance for a video while maintaining the overall video quality. Perceptually, we see a huge improvement in the subjective quality by our multiplexing method compared to EQL_TS.

Depending on the prices at every TS, the videos might receive unequal total numbers of bits in the multiplexing process, but all the videos benefit from these multiplexing methods. By changing the encoding technique inside a GOP (e.g., using multiple reference frame prediction or using hierarchical B-frames), along with these multiplexing methods, the overall video quality can be expected to further improve. The PSNR gains are negligible if the videos have similar complexity at every TS. In such a case, the bitrate requirements for all the videos are similar for the current TS and future TS and so very little trading will take place. Similarly, if each individual video has nearly identical complexity in each one of its TS compared to any other TS, then, even if the videos differ hugely in complexity compared to one another, the competitive equilibrium bitrate allocation would result in negligible PSNR improvement over EQL_TS. This is because the bitrate requirement for each video at the current TS is nearly the same as that for the future TS; no user would be willing to trade bits for the current TS with respect to the future TS. Then, all the videos receive almost the same number of bits at each TS. Thus, the competitive equilibrium bitrate allocation would result in negligible PSNR improvement.

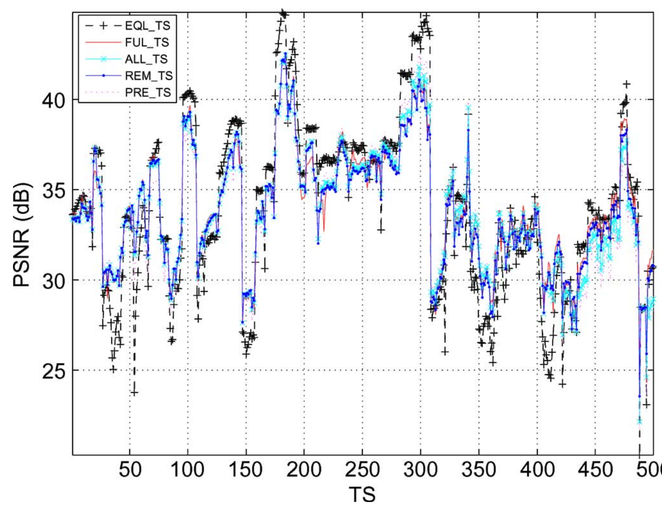

Fig. 11. Variation of PSNR with TS for g9 video at 100 kbits per TS.

The utility function for this paper was defined in terms of MSE. However, the proposed competitive equilibrium bitrate allocation method can be applied to any other utility function provided that it is convex. The interpretation would vary since users would have to communicate their individual subjective utility information instead of objective RD information about their videos.

A. Video Users With Variable Start and End Times

The results presented above are for N users who are simultaneously transmitting their videos. All the users are present during all TS. This is the same condition used in previous work on video multiplexing [1]–[6], [8]. We now consider video streams with different start and end times, so there are different numbers of users involved in the multiplexing at different times.

Let t_i and T_i be the start and end times for user i. The utility function for user i at TS t such that $t_i \leq t \leq T_i$ is given by

$$U_i\left(x_i^t, \bar{x}_i^t\right) = -\left(a_i^t + \frac{b_i^t}{x_i^t + d_i^t}\right) - (T_i - t).\left(\bar{a}_i^t + \frac{\bar{b}_i^t}{\bar{x}_i^t + \bar{d}_i^t}\right). \tag{16}$$

A competitive equilibrium at TS t is found by solving

$$\begin{aligned} &\max_{x_i^t, \bar{x}_i^t} -\left(a_i^t + \frac{b_i^t}{x_i^t + d_i^t}\right) - (T_i - t).\left(\bar{a}_i^t + \frac{\bar{b}_i^t}{\bar{x}_i^t + \bar{d}_i^t}\right) \\ &s.t. \quad p^t.x_i^t + (T_i - t).\bar{p}^t.\bar{x}_i^t = p^t.c_i^t + (T_i - t).\bar{p}^t.\bar{c}_i^t \\ &\qquad \forall i = 1 \; to \; N, i : t_i \leq t \leq T_i. \end{aligned} \tag{17}$$

The constraint on the total available bits in TS t is given by

$$\sum_{\substack{i=1 \\ i:t_i \leq t \leq T_i}}^{N} x_i^t = c^t. \tag{18}$$

The equilibrium prices can be found by solving

$$\sum_{\substack{i=1 \\ i:t_i \leq t \leq T_i}}^{N} \sqrt{\frac{b_i^t}{p^t}}.\left(\frac{p^t.\left(c_i^t + d_i^t\right) + (T_i - t).\bar{p}^t.\left(\bar{c}_i^t + \bar{d}_i^t\right)}{\sqrt{p^t.b_i^t} + (T_i - t).\sqrt{\bar{p}^t.\bar{b}_i^t}}\right) - \sum_{\substack{i=1 \\ i:t_i \leq t \leq T_i}}^{N} d_i^t = c^t. \tag{19}$$

Thus, the competitive equilibrium bitrate allocation can be calculated using these prices.

There are two scenarios for the system involving users with different start and end times. In the first, we consider that each user is offered some fixed dedicated bitrate but then chooses to add their allocation into the communal pool and join a competitive equilibrium allocation process. The total bitrate at any TS, therefore, scales up with the number of users in that TS. In the second scenario, users also start and end their video transmission at different times, but the shared channel has a bitrate that may be constant or variable over time, but in any case does not scale with the number of users. These scenarios are briefly discussed below.

Suppose we have a system where a user is assigned a specific bitrate when he enters. There are two options available for such users: (a) the user can transmit his video at his given bitrate and (b) the user can collaborate with other users, adding his allocation into the communal pool and achieving a competitive equilibrium bitrate allocation, with the expectation that he will be better off in terms of video quality by doing so. The overall bitrate depends on the number of users present at any TS and their initial endowments. Our competitive equilibrium bitrate allocation method was easily extended to such systems to improve the quality of all the users. Using the simulation results for such systems, it was found that the quality improvement using the competitive equilibrium bitrate allocation depends on the amount of time that the users overlap with each other. As one would expect, the case of large overlaps among large numbers of users produces higher quality improvement. As the number of such users increases in the system for the competitive equilibrium allocation, the performance improvement for collaborating (compared to retaining one's own equal initial endowment) increases since more trading takes place between the users for mutual advantage. The results are largely the same as the case of same start and end times. Current users of the system know that any additional users who join in will bring their own equal allocation with them for the communal pool, so there is, on the average, a slight improvement when new users join (because of the advantages of having more people involved in trades), and there is a slight disadvantage when people leave.

We now consider the second case, where the shared channel has a bitrate that does not depend on the number of users present. This may be a Constant Bitrate (CBR) channel or a Variable Bitrate (VBR) channel. In either case, users enter and leave the system at different times, and the average bitrate at any TS depends on both the total bitrate at that TS and on the number of users present. Therefore, the initial endowment to each user varies, and is performed at that TS. In such a case, the competitive equilibrium bitrate allocation depends on the number of users present at any TS for trading and the total available bitrate at that TS. If there are many users present, since the total bitrate does not scale up with the numbers of users, everyone's quality will be worse on the average compared to if there were fewer users, regardless of whether equal allocation or competitive equilibrium allocation is used. However, as in the previous case, the improvement for collaborating compared to equal allocation generally increases with the number of users present at any TS. However, if there are many users present, the quality of most of them (unless they have a very low complexity TS) is lower than their expected quality for the future, and users would then be unwilling to give away bits in that TS in exchange for future bits. In such a case, little or no trading will take place between the users. Table II shows the results for multiplexing 10 video streams with different start and end times for a shared constant bitrate of 500 kbits per TS. The start and end times for each user are given in Fig. 12. At some TS, there are as many as six video streams, and at other times, there are as few as three. To estimate the average video quality for future TS, we assume that the users know the average bitrate. At any TS, the competitive equilibrium is achieved for those users who are present at that TS. Even with different start and end times, all users improve their video quality compared to EQL_TS, as shown in Table II. Fig. 12 and Table II provide one example of differing start and end times. The results, in general, depend on the distributions of start times and end times, as well as on what users know about these distributions when they forecast their future bitrate demands.

We note also that our competitive equilibrium model for bitrate allocation is intended to be an approximation to a "large" system in which fluctuations in the number of users on the shared channel at any one time will be "close" to some average number. In that case, variable starting and ending times present no problem because the number of bits allocated to each user initially can be considered to be constant per TS. For all these different scenarios, the computations we make are exactly the same—each user is trading off bits between the current TS and an average of all future TS (which may be different for different

TABLE II
PSNR (db) IMPROVEMENT OVER EQL_TS BY USING VARIOUS MULTIPLEXING METHODS WHEN TEN VIDEO STREAMS WITH DIFFERENT START AND END TIMES ARE MULTIPLEXED TOGETHER AT A CONSTANT BITRATE OF 500 kbits PER TS

	g1	g2	g5	g6	g7	g8	g9	g10	g11	g12
FUL_TS	0.16	0.96	0.39	0.83	1.14	1.13	1.46	1.04	0.98	0.92
ALL_TS	0.49	0.22	0.52	0.94	0.86	0.88	1.32	1.05	1.18	1.01
REM_TS	0.46	0.08	0.63	1.07	0.76	0.45	1.28	1.03	1.20	1.25
PRE_TS	0.34	0.58	0.36	0.79	0.79	1.04	1.20	1.07	0.95	0.73

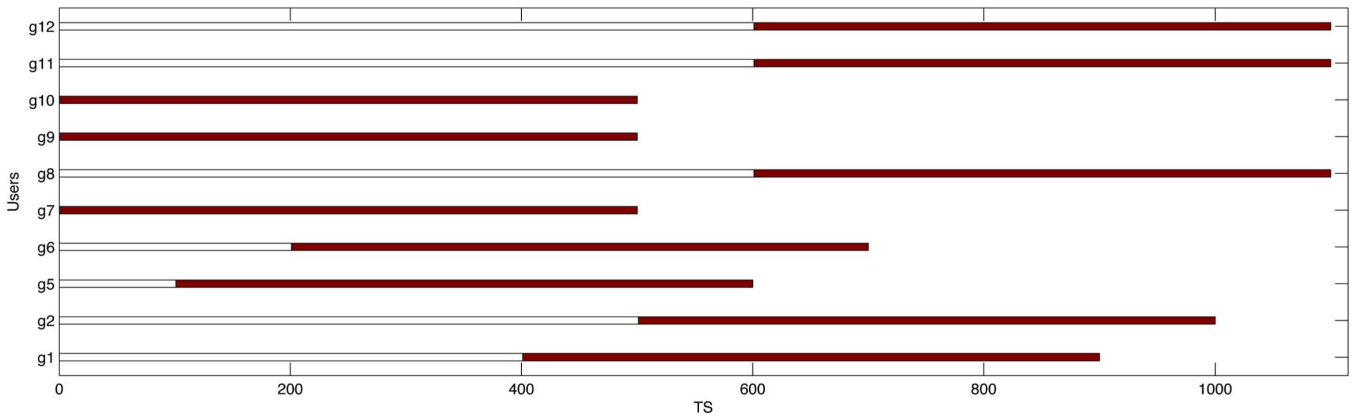

Fig. 12. Start and end times of ten video users.

users, depending on the length of their video and when they started).

V. CONCLUSION

We discussed various methods for multiplexing video streams for improving the quality of each individual video. Note that because this technique is applicable to both ad hoc networks that employ cluster heads and cellular architectures, it is relevant to both military and commercial scenarios. We considered the competitive equilibrium approach for allocating bitrate among various video streams. Using an Edgeworth box solution, we graphically showed the process of bitrate allocation at a competitive equilibrium. A central controller collects rate-distortion information from all the users. The central controller performs the competitive equilibrium calculation for a bitrate allocation to all the users simultaneously. The final bitrate allocation is a Pareto optimal solution and all the users do at least as well as they would with an individual allocation. The bitrate allocation information is sent to the video users, and the users use this information to encode their video streams. We proposed three different methods for estimating the future RD information for a video stream. The estimation of future RD information was used to trade bits for the current TS with respect to the expected bitrate requirement in the future. All the estimation methods work well for the competitive equilibrium allocation. The results show PSNR improvement for all the video streams. Comparing the decoded videos after multiplexing, we found that the subjective quality was improved by using the competitive equilibrium bitrate allocation when compared with EQL_TS. Typically, the video quality improvement is clearly visible in high motion parts of a video stream where more bits are allocated in the competitive equilibrium bitrate allocation methods. Generally, the PSNR improvement depends on the accuracy of estimating the RD information for future TS. The PSNR improvement is greater for the videos with higher motion fluctuation, even though their estimation of the future TS is almost stationary over time. Such videos have varying demand for the current TS with respect to the almost constant demand for the future TS, and so, are willing to trade away bits for now in anticipation of gaining bits at some future TS or vice versa.

ACKNOWLEDGMENT

The authors would like to thank Prof. M. Tagliasacchi from Politecnico di Milano for providing the data of [1] to compare the results. They would also like to thank the anonymous reviewers for their many constructive comments.

REFERENCES

[1] M. Tagliasacchi, G. Valenzise, and S. Tubaro, "Minimum variance optimal rate allocation for multiplexed H.264/AVC bitstreams," *IEEE Trans. Image Process.*, vol. 17, no. 7, pp. 1129–1143, Jul. 2008.

[2] L. Wang and A. Vincent, "Bit allocation and constraints for joint coding of multiple video programs," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 9, no. 6, pp. 949–959, Sep. 1999.

[3] G. Su and M. Wu, "Efficient bandwidth resource allocation for low-delay multiuser video streaming," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 15, no. 9, pp. 1124–1137, Sep. 2005.

[4] M. Tiwari, T. Groves, and P. Cosman, "Multiplexing video streams using dual-frame video coding," in *Proceedings IEEE Int. Conf. Acoustics, Speech, Signal Processing*, Mar. 2008, pp. 693–696.

[5] A. Fattahi, F. Fu, M. van der Schaar, and F. Paganini, "Mechanism-based resource allocation for multimedia transmission over spectrum agile wireless networks," *IEEE J. Sel. Areas Commun.*, vol. 25, no. 3, pp. 601–612, Apr. 2007.

[6] H. Park and M. van der Schaar, "Bargaining strategies for networked multimedia resource management," *IEEE Trans. Signal Process.*, vol. 55, no. 7, pp. 3496–3511, Jul. 2007.

[7] F. Fu and M. van der Schaar, "Noncollaborative resource management for wireless multimedia applications using mechanism design," *IEEE Trans. Multimedia*, vol. 9, no. 4, pp. 851–868, Jun. 2007.

[8] F. Fu, T. Stoenescuand, and M. van der Schaar, "A pricing mechanism for resource allocation in wireless multimedia applications," *IEEE J. Sel. Areas Commun.*, vol. 1, no. 2, pp. 264–279, Aug. 2007.

[9] M. Tiwari, T. Groves, and P. Cosman, “Bitrate allocation for multiple video streams at competitive equilibria,” in *Proc. IEEE Asilomar Conf. Signals, Systems, Computers*, pp. 2248–2252.
[10] J. Marschak and R. Radner, *Economic Theory of Teams*. New Haven, CT: Yale Univ. Press, 1972.
[11] A. Mas-Colell, M. Whinston, and J. Green, *Microeconomic Theory*. Oxford, U.K.: Oxford Univ. Press, 1995.
[12] F. Edgeworth, *Mathematical Psychics: An Essay on the Application of Mathematics to the Moral Sciences*. London, U.K.: C. Kegan Paul, 1881.
[13] L. Lin and A. Ortega, “Bit-rate control using piecewise approximated rate-distortion characteristics,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 8, no. 4, pp. 446–459, Aug. 1998.
[14] K. Stuhlmüller, N. Färber, M. Link, and B. Girod, “Analysis of video transmission over lossy channels,” *IEEE J. Sel. Areas Commun.*, vol. 18, no. 6, pp. 1012–1032, Jun. 2000.
[15] T. Wiegand, G. Sullivan, G. Bjontegaard, and A. Luthra, “Overview of the H.264/AVC video coding standard,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 13, no. 7, pp. 560–576, Jul. 2003.
[16] H.264/AVC Ref. Software [Online]. Available: http://iphome.hhi.de/suehring/tml
[17] K.-P. Lim, G. Sullivan, and T. Wiegand, “Text description of joint model reference encoding methods and decoding concealment methods,” JVT of ISO/IEC MPEG and ITU-T VCEG, JVT-K049, Mar. 2004.

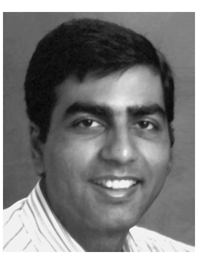

Mayank Tiwari (S’03) received the B.E. degree in electronics and communication engineering from the Indian Institute of Technology, Roorkee, India, in 1999, and the M.S. degree in electrical engineering (communication and signal processing) from the Arizona State University in 2004. He is currently pursuing the Ph.D. degree in electrical and computer engineering (signal and image processing) at the University of California at San Diego, La Jolla.

From June 1999 to May 2001, he was with Hughes Software Systems in Gurgaon, India, working on the design and development of video and speech codecs. From June 2001 to August 2002, he was with Cybernetics Infotech Inc. in Rockville, Maryland, working on the design and development of low bitrate speech codecs. His research interests include image and video compression, and mobile multimedia communication.

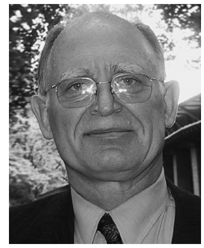

Theodore Groves received his B.A. degree from Harvard University, Cambridge, MA, and the Ph.D. degree in economics from the University of California, Berkeley, in 1970.

Prior to coming to the University of California at San Diego (UCSD), La Jolla, as a Professor of economics in 1979, he was a faculty member at the University of Wisconsin, Madison, Northwestern University’s Kellogg School of Management, and Stanford University, Stanford, CA. He was a founder of mechanism design theory and the discoverer of the “Groves Mechanism” for eliciting truthful information in an incentive-compatible manner. He and coauthor J. Ledyard also developed the first general solution to the “free rider problem” of public goods. He has also studied the Chinese economy’s transition to a market economy, optimal policies for minimizing the occurrence of oil spills, and, currently, improved methods for the multiplexing of videos over wireless transmission channels. He is the Director of the Center for Environmental Economics in the Department of Economics at UCSD and is involved in research on water pricing, consumer responses to smart-meter technology for electrical energy consumption, and numerous projects for managing marine resources and the protection of endangered species.

Dr. Groves is an elected Fellow of the Econometric Society and the American Academy of Arts and Sciences.

Pamela C. Cosman (S’88–M’93–SM’00–F’08) received the B.S. degree with Honors in electrical engineering from the California Institute of Technology, Pasadena, in 1987, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1989 and 1993, respectively.

She was an NSF postdoctoral fellow at Stanford University and a Visiting Professor at the University of Minnesota during 1993–1995. In 1995, she joined the faculty of the Department of Electrical and Computer Engineering, University of California at San Diego, La Jolla, where she is currently a Professor. She was the Director of the Center for Wireless Communications from 2006 to 2008. Her research interests are in the areas of image and video compression and processing and wireless communications.

Dr. Cosman is the recipient of the ECE Departmental Graduate Teaching Award (1996), a Career Award from the National Science Foundation (1996–1999), a Powell Faculty Fellowship (1997–1998), and a Globecom 2008 Best Paper Award. She was a guest editor of the June 2000 special issue of the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS on “Error-resilient image and video coding,” and was the Technical Program Chair of the 1998 Information Theory Workshop in San Diego. She was an associate editor of the IEEE COMMUNICATIONS LETTERS (1998–2001), and an associate editor of the IEEE SIGNAL PROCESSING LETTERS (2001–2005). She was the Editor-in-Chief (2006–2009) as well as Senior Editor (2003–2005, 2010–present) of the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS. She is a member of Tau Beta Pi and Sigma Xi.