Timestamping Schemes for MPEG-2 Systems Layer and Their Effect on Receiver Clock Recovery

Christos Tryfonas and Anujan Varma, *Member, IEEE*

***Abstract*— We propose and analyze several strategies for performing timestamping of an MPEG-2 Transport Stream transmitted over a packet-switched network using the PCR-unaware encapsulation scheme, and analyze their effect on the quality of the recovered clock at the MPEG-2 Systems decoder. When the timestamping scheme is based on a timer with a fixed period, the PCR values in the packet stream may switch polarity deterministically, at a frequency determined by the timer period and the transport rate of the MPEG signal. This, in turn, can degrade the quality of the recovered clock at the receiver beyond acceptable limits. We consider three timestamping schemes for solving this problem: 1) selecting a deterministic timer period to avoid the phase difference in PCR values altogether, 2) fine-tuning the deterministic timer period to maximize the frequency of PCR polarity changes, and 3) selecting the timer period randomly to eliminate the deterministic PCR polarity changes. For the case of deterministic timer period, we derive the frequency of the PCR polarity changes as a function of the timer period and the transport rate, and use it to find ranges of the timer period for acceptable quality of the recovered clock. We also analyze a random timestamping procedure based on a random telegraph process and obtain lower bounds on the rate of PCR polarity changes such that the recovered clock does not violate the PAL/NTSC clock specifications. The analytical results are verified by simulations with both synthetic and actual MPEG-2 Transport Streams sent to a simulation model of an MPEG-2 Systems decoder.**

***Index Terms*— ATM, clock recovery, MPEG-2, MPEG-2 Systems layer, PCR-unaware, set-top box.**

I. Introduction

MPEG-2 is the emerging standard for audio and video compression. By exploiting both spatial and temporal redundancies of the input signal, it achieves compression ratios of up to two orders of magnitude, enabling it to encode a video or audio signal to virtually any desired level of quality. The MPEG-2 Systems layer specifies two mechanisms to multiplex elementary audio, video or private streams to form a *program*, namely the *MPEG-2 program stream* and the *MPEG-2 Transport Stream* formats. In loss-prone environments such as a packet-switched network, the MPEG-2 Transport Stream is the preferred approach for transporting MPEG-2 streams. An MPEG-2 Transport Stream combines one or more programs into a single packet stream. The use of explicit timestamps (called *program clock references* or *PCR's* in MPEG-2 terminology) within the packets facilitates the clock recovery at the decoder end ensures synchronization and continuity of MPEG-2 Transport Streams. For a brief tutorial of the MPEG-2 Systems layer the interested reader is referred to [12].

Several issues must be taken into account while transporting MPEG-2 encoded streams over packet-switched networks. These include the choice of the adaptation layer, method of encapsulation of MPEG-2 packets into network packets, provision of quality-of-service (QoS) in the network to ensure control of delay and jitter, and the design of the decoder. At the decoder end, application-specific requirements such as accuracy and stability determine the approaches that should be taken to recover the system clock [6]. A certain category of applications, such as broadcast TV distribution, uses the reconstructed system clock to directly synthesize a chroma subcarrier for the composite video signal. The system clock, in this case, is used to derive the chroma subcarrier, the pixel clock and the picture rate. The composite video subcarrier in these systems must have at least sufficient accuracy and stability so that any normal television receiver's chroma subcarrier PLL can lock to it, and the chroma signals which are demodulated using the recovered subcarrier do not show any visible chrominance phase artifacts. There are often cases in which the application has to meet NTSC, PAL or SECAM specifications, which are even more stringent. As an example, NTSC requires a subcarrier accuracy of 3 ppm with a maximum long-term drift of 0.1 Hz/s. In such applications, a standard PLL is often used to recover the clock from the PCR timestamps transmitted within the stream.

While clock recovery is critical in set-top box design, it may also be critical in desktop computer environments which display decoded MPEG streams on the local display, and not on an NTSC/PAL monitor. The synchronization constraints are different when the synchronization signals of the display are not generated from the incoming packet stream, but there is still a need to synchronize the display of decoded frames with the frame-buffer's raster clock. Besides, a more stable clock always minimizes the probability of playout buffer underflow or overflow and, therefore, provides uninterrupted and consistent decoding.

The use of applications with stringent clock specifications requires careful design of the decoder since the decoder is responsible for feeding the TV set with a composite signal

Manuscript received December 21, 1998; revised June 30, 1999. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. M. Reha Civanlar. This work was supported by the Advanced Research Projects Agency (ARPA) under Contract F19628-96-C-0038 and the NSF Young Investigator Award MIP-9257103. The OPNET modeling tool used for simulations was donated by MIL-3, Inc.

C. Tryfonas is with Sprint Advanced Technology Laboratories, Burlingame, CA 94010 USA (e-mail: tryfonas@sprintlabs.com).

A. Varma is with the Computer Engineering Department, University of California, Santa Cruz, CA 95064 USA (e-mail: varma@cse.ucsc.edu).

Publisher Item Identifier S 1520-9210(99)07544-6.

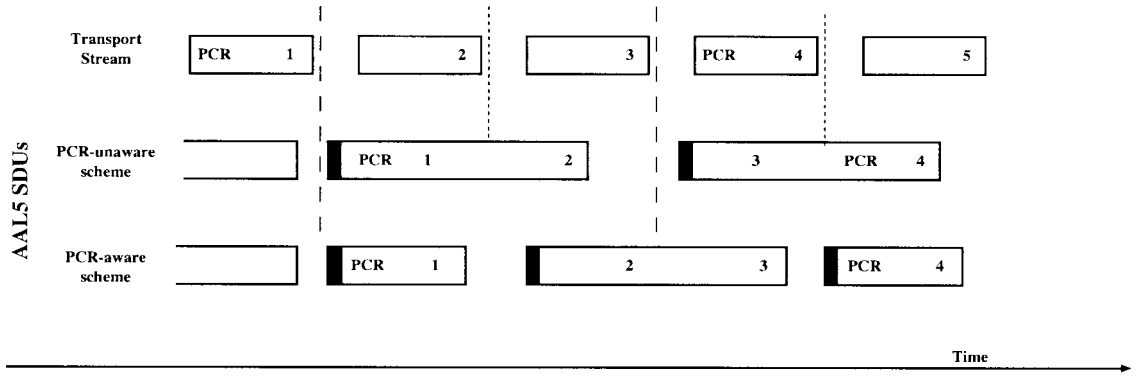

Fig. 1. PCR packing schemes for AAL5 in ATM networks (cut-through operation assumed).

that meets the requirements [2]. The demodulator in the TV set has to extract clock information from this signal for the color subcarrier regeneration process. The frequency requirements for NTSC specify a tolerance of ± 10 Hz (which corresponds to ± 3 ppm) [3]. The central subcarrier frequency is 3.579 545 MHz (or exactly (63/88) $\times$ 5 MHz). The corresponding tolerance for PAL signals is ± 5 Hz around the central subcarrier frequency of 4.433 618 75 MHz. The requirements above define the precision of the oscillators for the modulator and thus, the minimum locking range for the PLL at the receiver end.

The degradation of the reconstructed clock at the receiver is caused primarily by the packet delay variation (jitter). Three different causes contribute to the jitter experienced by an MPEG-2 Transport Stream as seen at the receiving end. The first is the frequency drift between the transmitter and the receiver clocks, which is usually small compared to the other two components. The second component of jitter is due to the packetization at the source, which may displace timestamp values within the stream. Finally, the network may introduce a significant amount of jitter, owing to the variations in queueing delays in the network switches. In this paper, our focus is in the second component, the packetization jitter.

The packet encapsulation procedure is the main cause for the packetization jitter. In the context of asynchronous transfer mode (ATM) networks, two approaches have been proposed for encapsulation of MPEG-2 Transport Streams in ATM adaptation layer 5 (AAL5) packets: the *PCR-aware* and the *PCR-unaware* schemes [1] (these schemes are illustrated in FIg. 1). In the PCR-aware scheme, packetization is done ensuring that when a Transport Stream packet contains a PCR value it will be the last packet encapsulated in an AAL-5 packet. This minimizes the jitter experienced by PCR values during packetization. In the PCR-unaware approach, the sender performs the encapsulation without checking whether PCR values are contained within a transport packet. In the context of this paper, assuming that two transport packets form one AAL5 protocol data unit (PDU), this means that a transport packet containing a PCR value may be either the first or the second packet in the AAL5 PDU. When it is the first packet, we designate the timestamp as falling on an *odd* boundary; similarly, if the PCR value is in the second packet of the AAL5 PDU, we define the timestamp as falling on an *even* boundary. Therefore, the encapsulation procedure may introduce significant jitter on the PCR values. The presence of jitter introduced by the adaptation layer in this case may distort the reconstructed clock at the MPEG-2 audio/video decoder. This, in turn, may degrade the quality when the synchronization signals for display of the video frames on the TV set are generated from the recovered clock.

Several approaches have been proposed in the literature [1], [4], [5], [7], [11], [14] for the design of the MPEG-2 decoder to reduce the effects of jitter and provide acceptable quality for the decoded program. However, the impact of the timestamping process on the clock recovery at the decoder has not been well understood. Insight into some aspects of the timestamping process with a fixed-period timer, when transporting MPEG-2 over AAL5 using the PCR-unaware scheme, was first given by Akyildiz *et al.* [1]. In that work, it was shown that transport packets containing PCR values may switch indices (between odd and even) in a deterministic manner with a period that depends on both the transport rate and the timer period. This behavior was referred to as "pattern switch." This effect can be avoided by forcing all the PCR values to occupy the same phase in the encapsulated packet stream, or by compensating for the phase difference at the receiver.

In this paper, we analyze several strategies for performing timestamping of the outgoing MPEG-2 Transport Stream and analyze their effect on the clock recovery process of the MPEG-2 Systems decoder for applications with stringent clock requirements, assuming that the encapsulation procedure follows the PCR-unaware scheme. We start with an analysis of a timestamping process based on a timer with a fixed period. We derive the deterministic pattern switch frequency as a function of the timer period and the transport rate, and use it to find ranges of the timer period for acceptable quality of the recovered clock. Next, we analyze a random timestamping procedure based on a *random telegraph* process [8] and obtain lower bounds on the rate of change of PCR polarity so that the recovered clock meets the PAL/NTSC specifications. Based on this analysis, we consider three timestamping schemes:

1) setting the deterministic timer period precisely to avoid the phase difference in PCR values altogether;
2) fine-tuning the deterministic timer period to maximize the pattern switch frequency;

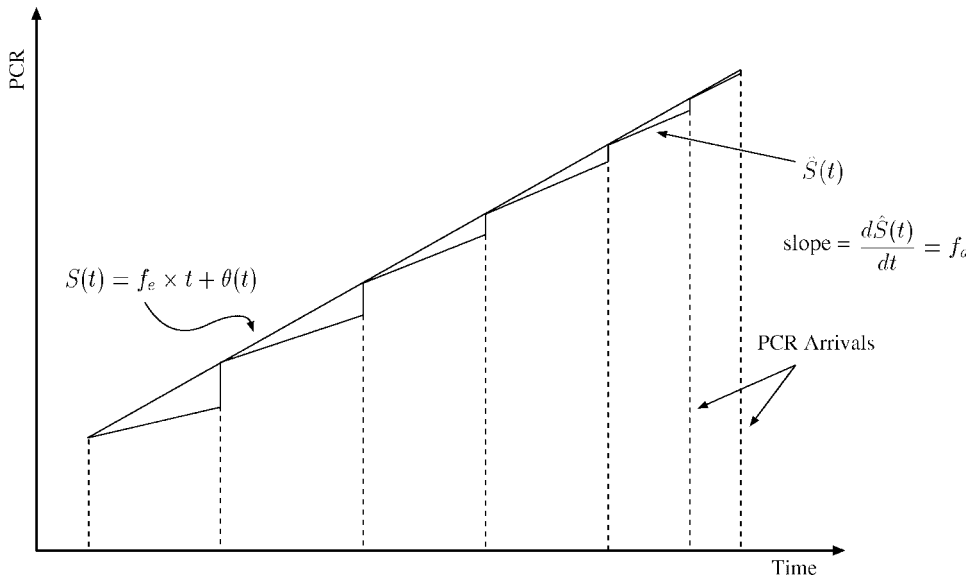

Fig. 2. Actual PCR function and PCR function used in the analysis.

3) using a random distribution for the timer period to eliminate the deterministic pattern switch behavior.

To verify our analytical results, we perform several simulation experiments with both synthetic and actual MPEG-2 Transport Streams sent to an MPEG-2 System decoder model, over both a point-to-point link and a multihop ATM network. Our results show that, for certain ranges of timer periods, the low frequency of the pattern switch results in unacceptable quality of the reconstructed clock. This contradicts with the hypothesis of Akyildiz *et al*. [1], that reducing the frequency of the deterministic pattern switch improves the quality of the recovered clock. For deterministic timer periods, fine-tuning the timer to *maximize* the frequency of pattern switch results in the best quality of the receiver clock.

The rest of this paper is organized as follows. In Section II, we analyze the effects of packetization jitter on the PAL/NTSC clock recovery process due to the PCR-unaware scheme, using a model of a standard PLL in an MPEG-2 Systems decoder under both deterministic and probabilistic timestamping procedures. Based on this analysis, in Section III we propose a number of timestamping schemes for the sender so that the quality of the recovered clock at the decoder is acceptable. In Section IV, we evaluate the schemes by simulation using a standard PLL employed in an MPEG-2 Systems decoder, with input from both synthetic MPEG-2 traces and actual MPEG-2 Transport Stream traces that are sent through a multihop ATM network. Finally, we conclude the paper in Section V with a summary of the results.

II. Analysis

In this section we provide an analysis of the effects of packetization jitter on the MPEG-2 decoder PLL. We first start with the modeling of the actual PLL and derive the transfer function to be used in the analysis. Next, we proceed to characterize the input signal at the PLL resulting from the timestamping and encapsulation schemes at the transmitter. We consider two distinct classes of timestamping schemes. In the first, timestamps are generated by a timer with a deterministic period; and in the second, the timer periods are drawn from a random distribution. In the first case, we derive the pattern switch frequency as a function of the timer period and transport rate of the MPEG-2 stream, which provides the phase of the input signal at the receiver PLL. In the second case, we use a random telegraph process to model the effect of the timestamping process, and use it to derive the variance of the recovered clock. This enables us to derive a lower bound on the required rate of change of PCR polarity in the packet stream to maintain the receiver PLL jitter within the specifications.

A. Derivation of Transfer Function of the Decoder PLL

We follow an approach similar to that of [9] for traditional PLL's. The main difference in our analysis is in the nature of the input signal. In our case, the input signal is a linear function as shown in Fig. 2, whereas in the case of traditional PLL's, the input signal is considered to be a sinusoidal function.

Before continuing with our analysis, we define some notations and assumptions used. Although the PCR's arrive at discrete points in time, we can assume that the incoming PCR's form a continuous-time function $S(t)$ that is updated at the instants when a new PCR value is received. We can model the incoming clock with the function

$$S(t) = f_e \times t + \theta(t) \tag{1}$$

where f_e is the frequency of the encoder sending the MPEG-2 stream and $\theta(t)$ is the incoming clock's phase relative to a designated time origin. As indicated in Fig. 2, there is a small discrepancy when modeling the incoming clock signal. The actual incoming clock signal $\hat{S}(t)$ is a function with discontinuities at the time instants at which PCR values are received, with slope equal to f_d for each of its segments, where f_d is the running frequency of the decoder. For the sake of convenience, however, we use $S(t)$ in place of the actual PCR function $\hat{S}(t)$, since the interval between any two consecutive

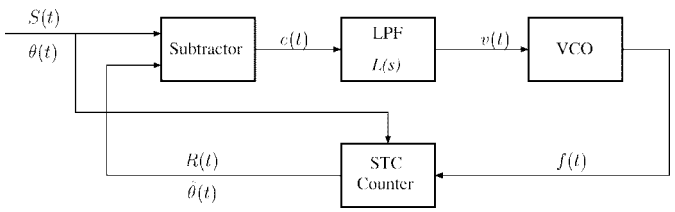

Fig. 3. Equivalent model of the PLL used.

PCR arrivals is bounded by the MPEG-2 standard and equal to at most 0.1 s, which makes the two functions to be very close.

Analogously, the system time clock (STC) corresponds to the function

$$R(t) = f_d \times t + \hat{\theta}(t) \tag{2}$$

where $\hat{\theta}(t)$ is the incoming clock's phase relative to a designated time origin. Therefore, referring to the model of the PLL in Fig. 3, the error term after the subtractor is given by

$$e(t) = S(t) - R(t) = (f_e - f_d)t + (\theta(t) - \hat{\theta}(t)). \tag{3}$$

Without loss of generality, we can assume that $f_e = f_d$. Let us denote this with f_o and insert any frequency difference in the phase terms. We can now work with $\theta(t)$ as being the input to our control system and with $\hat{\theta}(t)$ as being the output of the counter as shown in Fig. 3. Thus, (3) becomes

$$e(t) = \theta(t) - \hat{\theta}(t). \tag{4}$$

The frequency $f(t)$ of the VCO is a function of $v(t)$. The nominal value of this frequency is assumed to be f_0 and when $v(t)$ is applied, it becomes $f_0 + Kv(t)$ where K is the gain factor of the VCO. It is obvious that

$$\frac{dR(t)}{dt} = f_0 + Kv(t). \tag{5}$$

By definition,

$$R(t) = f_0 t + \hat{\theta}(t). \tag{6}$$

Combining (5) and (6) we get

$$\frac{d\hat{\theta}(t)}{dt} = Kv(t). \tag{7}$$

From (4) and (7) we obtain

$$\frac{de(t)}{dt} = \frac{d\theta(t)}{dt} - Kv(t) = \frac{d\theta(t)}{dt} - K\int_0^\infty l(t-u)e(u)\,du. \tag{8}$$

We assume that the Laplace transformations of $e(t)$ and $\theta(t)$ exist and they are $E(s)$ and $\Theta(s)$ respectively, and $L(s)$ is the low-pass filter's transfer function. Equation (8), when transformed to the Laplace domain, becomes

$$sE(s) = s\Theta(s) - KL(s)E(s). \tag{9}$$

We assume that $\hat{\theta}(t)$ has a Laplace transform. Using $E(s) = \Theta(s) - \hat{\Theta}(s)$, where $\hat{\Theta}(s)$ is the Laplace transform of $\hat{\theta}(t)$, and (9) we can now derive the transfer function $H(s)$ of the closed-loop

$$H(s) = \frac{\hat{\Theta}(s)}{\Theta(s)} = \frac{KL(s)}{s + KL(s)}. \tag{10}$$

The Laplace transform $F(s)$ of the recovered frequency function $f(t)$ is given by

$$\begin{aligned} F(s) &= (\Theta(s) - \hat{\Theta}(s))L(s)K \\ &= (1 - H(s))\Theta(s)L(s)K \\ &= K\left(1 - \frac{KL(s)}{s + KL(s)}\right)L(s)\Theta(s) \qquad (11) \\ &= P(s)\Theta(s) \qquad (12) \end{aligned}$$

where $P(s)$ is given by

$$P(s) = K\left(1 - \frac{KL(s)}{s + KL(s)}\right)L(s). \tag{13}$$

In order to model the LPF, we used a first order Butterworth analog LPF with a cutoff frequency of 0.1 Hz. This cutoff frequency was derived from both the experimental results of [12] and the analysis done in [7]. Following the methodology presented in [10] we get

$$L(s) = \frac{0.2\pi}{s - s_1} \tag{14}$$

with s_1 being the single real pole of the filter that has a value of

$$s_1 = 2\pi 0.1 \exp(j\pi) = -0.2\pi. \tag{15}$$

Hence, from (14) and (15) we obtain

$$L(s) = \frac{0.2\pi}{s + 0.2\pi}. \tag{16}$$

B. Possible Input Processes Due to PCR-Unaware Scheme

We now study the effect of possible input processes resulting from the timestamping procedure at the source on the clock recovery process at the decoder. Since we are interested in the tracking performance of the PLL, we can assume that the PLL is locked before the input process is applied as the input function $\theta(t)$ of the PLL.

Under the PCR-unaware scheme, an AAL packet containing two MPEG-2 transport packets may contain a PCR value either in the first or in the second MPEG-2 transport packet. Therefore, a PCR value will suffer either zero or one transport packet delay at the destination. Assuming that the PLL is locked before the input process is applied, the resulting phase difference values at its input $\theta(t)$ will be approximately $\pm f_o/2r$, where f_o is the central frequency in MPEG-2 Systems layer and r is the rate of the MPEG-2 Transport Stream in packets/s. We analyze two cases: a deterministic case in which a timer with a fixed-period is used to perform the timestamping procedure and a probabilistic case in which the PCR values are placed randomly in the MPEG-2 Transport Stream according to a random telegraph process.

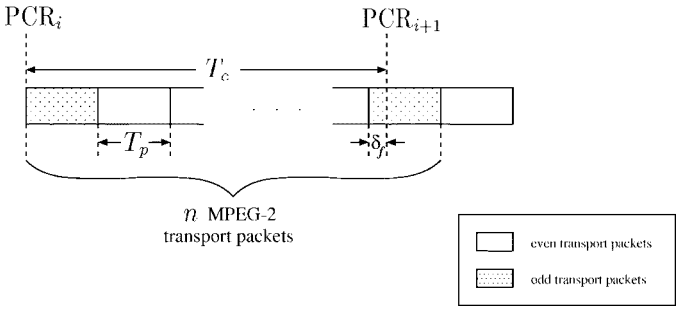

Fig. 4. Illustration of forward drift of PCR values' packet boundaries when $0 \leq \tau < T_p/2$.

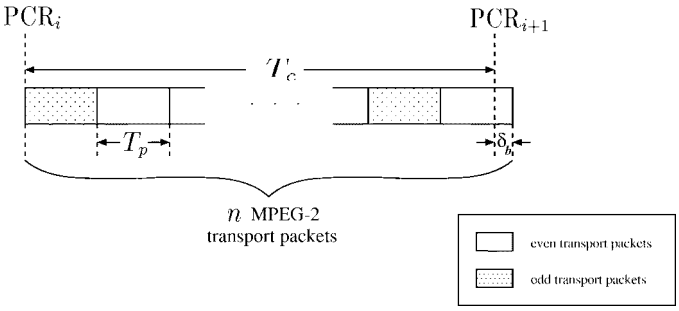

Fig. 5. Illustration of backward drift of PCR values' packet boundaries when $3T_p/2 \leq \tau < 2T_p$.

1) Deterministic Case: When a timer with a constant period is used at the source to timestamp the MPEG-2 packet stream, the positions of the PCR values switch between even and odd boundaries in the AAL packets at a constant frequency. This effect was first observed by Akyildiz *et al.* [1], who termed it "pattern switch." In this section, we systematically derive the pattern switch frequency as a function of the timer period and the transport rate of the MPEG-2 stream.

Let T_p denote the inter-arrival time of MPEG-2 transport packets, and T_c the period of the timer at the transmitter. Since $T_c > T_p$, we can express T_c in terms of T_p as

$$T_c = 2kT_p + \tau \tag{17}$$

where k is a nonnegative integer and $0 \leq \tau < 2T_p$. Since, in general, T_c is not an exact multiple of T_p, the actual time instants at which the PCR values are inserted into the MPEG-2 Transport Stream will drift relative to packet boundaries. More specifically, we need to consider three cases corresponding to different ranges of τ.

Case 1: $0 \leq \tau < T_p/2$. In this case, a forward drift of the resulting packet boundaries of the associated PCR values can be identified as illustrated in Fig. 4. Let n denotes the integer number of transport packets included in a T_c period, that is,

$$n = \left\lceil \frac{T_c}{T_p} \right\rceil. \tag{18}$$

Let δ_f denote the forward drift, given by

$$\delta_f = T_c - (n-1)T_p. \tag{19}$$

The last equation can be easily derived by observing Fig. 4.
From (18) and (19) we obtain

$$\delta_f = T_c - (n-1)T_p = T_c - \left(\left\lceil \frac{T_c}{T_p} \right\rceil - 1 \right) T_p. \tag{20}$$

As becomes evident from (20), the number of continuous PCR packets falling into odd or even positions in the MPEG-2 Transport Stream n_c, is given by

$$\left\lfloor \frac{T_p}{\delta_f} \right\rfloor \leq n_c \leq \left\lceil \frac{T_p}{\delta_f} \right\rceil. \tag{21}$$

Thus, the polarity (even/odd) of timestamp values in the packet stream exhibits a square-wave pattern at the input of MPEG-2 decoder's PLL with a period of $2n_cT_c$ and peak-to-peak amplitude of f_oT_p. Therefore the phase of the input signal at the PLL is given by

$$\theta(t) = f_oT_p \sum_{\lambda=-\infty}^{\infty} \left(u\left(t + \frac{\lambda n_c T_c}{2}\right) - u\left(t - \frac{\lambda n_c T_c}{2}\right)\right) - \frac{f_oT_p}{2} \tag{22}$$

in which $u(t)$ is the unit-step function. If the frequency of the above input signal becomes less than the bandwidth of the PLL, the output of the PLL will follow the pulse with a resulting degradation of the quality of the recovered clock. If we assume that the PLL has a perfect LPF of bandwidth B_L, then the period of $\theta(t)$ should be less than $1/B_L$. That is,

$$2n_cT_c < \frac{1}{B_L}. \tag{23}$$

Case 2: $T_p/2 \leq \tau < 3T_p/2$. In this case, at most two consecutive PCR values may fall into odd- or even-numbered MPEG-2 transport packets. In the specific case that $\tau = T_p$, the PCR values fall in alternate odd- and even-indexed transport packets producing the maximum frequency of changes in timestamp position in the packet stream. The resulting process has high-frequency components which are filtered by the decoder PLL and are unlikely to affect the quality of the recovered clock. This is verified in our simulation experiments of Section IV.

Case 3: $3T_p/2 \leq \tau < 2T_p$. This is similar to the first case, except that the drift of the packet boundaries of the PCR values is in the backward direction (Fig. 5). In this case, let δ_b denote the backward drift, given by

$$\delta_b = nT_p - T_c = \left\lceil \frac{T_c}{T_p} \right\rceil T_p - T_c. \tag{24}$$

Again, the last equation can easily be derived from Fig. 5. Similarly, the number of continuous PCR packets falling into only odd or only even positions in the MPEG-2 Transport Stream n_c, is bounded

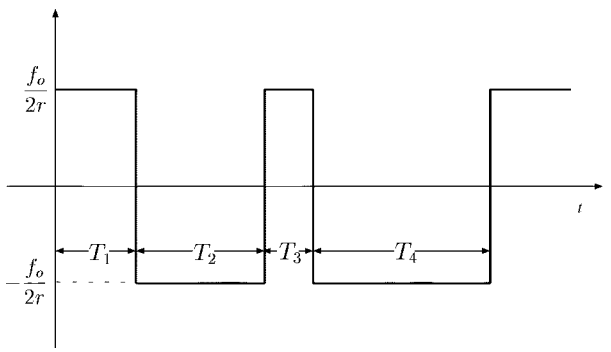

Fig. 6. Sample realization of $T_s(t)$ process. T_i's are independent identically-distributed (i.i.d.) exponential random variables.

by the following inequality

$$\left\lfloor \frac{T_p}{\delta_b} \right\rfloor \leq n_c \leq \left\lceil \frac{T_p}{\delta_b} \right\rceil. \tag{25}$$

The resulting phase at the input of the PLL, in this case also, is a square-wave with a period of $2n_cT_c$ and peak-to-peak amplitude of f_oT_p. Therefore the input function at the PLL is again given by

$$\theta(t) = f_oT_p \sum_{\lambda=-\infty}^{\infty} \left(u\left(t + \frac{\lambda n_c T_c}{2}\right) - u\left(t - \frac{\lambda n_c T_c}{2}\right)\right) - \frac{f_oT_p}{2}, \tag{26}$$

with $u(t)$ being the unit-step function.

2) Probabilistic Case: The generation of MPEG-2 Transport Streams with variable inter-PCR delay can be done by randomizing the timestamping procedure according to some distribution. In the probabilistic case, we assume that the PCR values fall completely in random places in the MPEG-2 Transport Stream. Without loss of generality, we can assume that they have the same probability of being in odd- or even-indexed transport packets (as by using Bernoulli trials). For the sake of convenience, we analyze this behavior by modeling the input phase $\theta(t)$ as a random telegraph process [8].

The objective of our analysis is to obtain the variance or the actual function that describes the recovered clock, i.e., $f(t)$. We derive the variance of the recovered clock in the case that the sequence of $\pm f_o/2r$ values forms a scaled random telegraph process. The random telegraph process $T(t)$ is a random process that assumes values of ± 1, has a mean of zero, and is stationary or cyclostationary. Assuming that initially $T(0) = \pm 1$ with equal probability, $T(t)$ is generated by changing polarity with each occurrence of an event of a Poisson process of rate a. In our analysis, we use a scaled version of the random telegraph process $T(t)$ in which the process gets the values $\pm f_o/2r$. We refer to the scaled version by $T_s(t)$. A sample realization of this process is shown in Fig. 6.

We first obtain the statistic measures of the scaled random telegraph process $T_s(t)$. Since the mean of the random telegraph process is zero, the mean of the scaled version is also zero. Let us now derive the autocorrelation function of $T_s(t)$:

$$\begin{aligned} R_{T_s}(t_1,t_2) &= E[T_s(t_1)T_s(t_2)], \quad \forall \text{ real numbers } t_1, t_2 \\ &= \left(\frac{f_o}{2r}\right)^2 \text{Prob}\left[T_s(t_1) = T_s(t_2)\right] \\ &\quad - \left(\frac{f_o}{2r}\right)^2 \text{Prob}\left[T_s(t_1) \neq T_s(t_2)\right] \\ &= \left(\frac{f_o}{2r}\right)^2 \frac{(1 + e^{-2a|t_2 - t_1|})}{2} \\ &\quad - \left(\frac{f_o}{2r}\right)^2 \frac{(1 - e^{-2a|t_2 - t_1|})}{2} \\ &= \frac{f_o^2}{4r^2} e^{-2a|t_2 - t_1|}. \end{aligned} \tag{27}$$

We also obtain the *power spectral density* (psd) of the input process. The psd function of the input process is given as the Fourier transform of the autocorrelation function $R_{T_s}(k)$. Thus

$$\begin{aligned} S_{T_s}(w) &= \int_{-\infty}^{\infty} R_{T_s}(\tau) e^{-jw\tau} d\tau \\ &= \int_{-\infty}^{\infty} \frac{f_o^2}{4r^2} e^{-2a|\tau|} e^{-jw\tau} d\tau \\ &= \frac{f_o^2}{4r^2}\left(\int_{-\infty}^{0} e^{2a\tau} e^{-jw\tau} d\tau + \int_{0}^{\infty} e^{-2a\tau} e^{-jw\tau} d\tau\right) \\ &= \frac{f_o^2}{4r^2}\left(\frac{1}{2a - jw} + \frac{1}{2a + jw}\right) \\ &= \frac{af_o^2}{r^2(4a^2 + w^2)}. \end{aligned} \tag{28}$$

The psd function of the recovered clock is given by

$$S_f(w) = S_{T_s}(w)|P(w)|^2 \tag{29}$$

where $|P(w)|$ is the magnitude of the Fourier transform of the function defined in (13). We can obtain the Fourier transform of the signal that has Laplace transform of $P(s)$ by substituting s with jw. After substitution, and using (16) we get

$$P(w) = \frac{Kw}{w + \left(\frac{w^2}{0.2\pi} - K\right)j}. \tag{30}$$

The square of the magnitude of $P(w)$ is given by

$$\begin{aligned} |P(w)|^2 &= P(w)P^*(w) \\ &= \left(\frac{Kw}{w + \left(\frac{w^2}{0.2\pi} - K\right)j}\right)\left(\frac{Kw}{w - \left(\frac{w^2}{0.2\pi} - K\right)j}\right) \\ &= \frac{(Kw)^2}{w^2 + \left(\frac{w^2}{0.2\pi} - K\right)^2} \end{aligned} \tag{31}$$

where $P^*(w)$ is the conjugate of $P(w)$.

The variance of the output process is determined by the inverse Fourier transform of $S_f(w)$ at point zero. That is,

$$\begin{aligned} \sigma^2_{f(t)} &= \mathcal{F}\{S_f(w)\}|_0 \\ &= \frac{1}{2\pi}\int_{-\infty}^{\infty} S_f(w) dw \\ &= \frac{1}{2\pi}\int_{-\infty}^{\infty} \left(\frac{af_o^2}{r^2(4a^2 + w^2)}\right)\left(\frac{(Kw)^2}{w^2 + \left(\frac{w^2}{0.2\pi} - K\right)^2}\right) dw \\ &= \frac{af_o^2\pi K^2}{r^2(40a^2 + 4\pi a + 2\pi K)}. \end{aligned} \tag{32}$$

The above equation gives the variance of the clock at the MPEG-2 Systems decoder. The clock of the color subcarrier is derived from this clock using a scaling factor which is different for PAL and NTSC. Since the scaled random telegraph process $T_s(t)$ is bounded, we can assume that the recovered clock deviates from its central frequency by at most $\sigma_{f(t)}$. Applying the requirements for the subcarrier frequency we get

$$\sigma_{f_{\mathrm{PAL}}} = \frac{A_{\mathrm{PAL}} f_o K}{r} \sqrt{\frac{a\pi}{40a^2 + 4\pi a + 2\pi K}} \leq 5 \tag{33}$$

for the recovered PAL subcarrier frequency and

$$\sigma_{f_{\mathrm{NTSC}}} = \frac{A_{\mathrm{NTSC}} f_o K}{r} \sqrt{\frac{a\pi}{40a^2 + 4\pi a + 2\pi K}} \leq 10 \tag{34}$$

for the recovered NTSC subcarrier frequency with A_{PAL} and A_{NTSC} being approximately equal to $\frac{4.43}{27}$ and $\frac{3.58}{27}$ respectively. We can now obtain a lower bound for the allowed rate r in packets/s in the case of PAL and NTSC subcarriers, so that the clock remains within the specifications.

$$r \geq \frac{A_{\mathrm{PAL}} f_o K}{5} \sqrt{\frac{a\pi}{40a^2 + 4\pi a + 2\pi K}} \tag{35}$$

for the PAL case and

$$r \geq \frac{A_{\mathrm{NTSC}} f_o K}{10} \sqrt{\frac{a\pi}{40a^2 + 4\pi a + 2\pi K}} \tag{36}$$

for the NTSC case. Analogously, we can derive a bound on the rate of change of PCR polarity a so that the clock specifications are not violated under a specific transport rate.

$$a \geq \frac{-\pi(4r^2 - X^2) + \sqrt{(4\pi r^2 - \pi X^2)^2 - 320\pi K r^4}}{80r^2} \tag{37}$$

where X is equal to $\frac{A_{\mathrm{NTSC}} f_o K}{10}$ for the NTSC case and $\frac{A_{\mathrm{PAL}} f_o K}{5}$ for the PAL case.

For the sake of convenience, we concentrate on the PAL case only and compute the numerical value for the minimum rate for a typical MPEG-2 decoder PLL. The constant K affects the gain and the stability of the loop. It is also used to scale the input signal to the appropriate levels for the MPEG-2 frequency. More specifically, the design of the VCO takes into account the maximum difference in ticks of a 27 MHz clock when the jitter or the PCR inaccuracy due to remultiplexing operations is at its maximum allowed value, and the limits of the frequency of the decoder. Since, according to MPEG-2 standard [6], the maximum jitter expected is around ± 0.5 ms, the maximum allowable difference is 13 500 ticks. For this maximum difference, the decoder must operate at a frequency within the limits specified in the MPEG-2 standard. That is,

$$\begin{aligned} 27\ \mathrm{MHz} - 810\ \mathrm{Hz} &\leq \text{decoder clock frequency} \\ &\leq 27\ \mathrm{MHz} + 810\ \mathrm{Hz}. \end{aligned} \tag{38}$$

Therefore, the selection of K should be around the value of $\frac{810}{13\,500}$ or 0.06 in order for the decoder to operate correctly, provided that this value results in a stable loop. We can also assume that a is greater or equal to one which corresponds to an underlying Poisson process that has a minimum average rate of one arrival every second. Then, the minimum transport rate for the stream, in order not to have any PAL clock violations, is approximately 19.476 Mbps. A similar result can be found for the NTSC case. Higher values for a produce even lower bounds. For example, a value of 20 gives a bound of approximately 4.970 Mbps. Fig. 7 shows a plot of a versus the MPEG-2 transport rate, where the the shaded region corresponds to the range of a and transport rate for the reconstructed clock to stay within the PAL specifications.

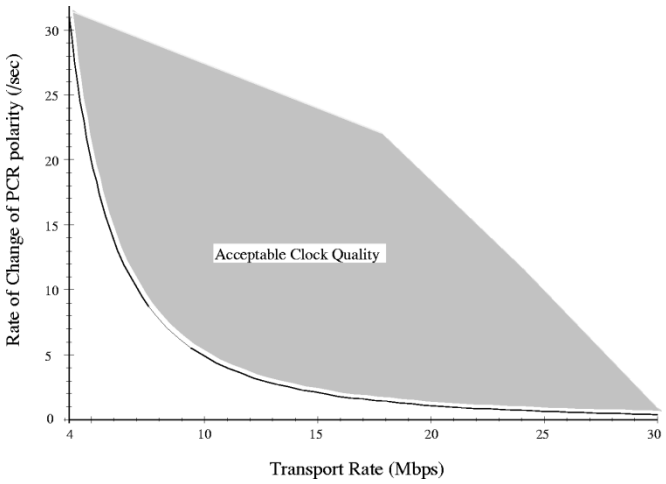

Fig. 7. Minimum rate of change of PCR polarity to maintain acceptable clock quality as a function of the transport rate of the MPEG-2 stream under PAL.

III. Solutions for Providing Acceptable Clock Quality

In the last section, we analyzed and quantified the effect of the timestamping process at the transmitter on the quality of the recovered clock at the receiver. When the timer-period for timestamping is chosen deterministically, the pattern switch behavior may manifest itself as a periodic square-wave signal at the input of the decoder PLL in the MPEG-2 Systems layer. One option to avoid the effect of this pattern switch signal is to eliminate it altogether by forcing all PCR values to occupy the same phase in the AAL packet stream. This would make the receiver clock quality under the PCR-unaware scheme identical to that under the PCR-aware scheme. A second alternative is to maximize the pattern switch frequency by causing the PCR values to switch between odd and even positions in the packet stream at the maximum rate. Finally, a third alternative is to use a random timestamping interval to avoid the deterministic pattern switch behavior. In this section, we discuss the tradeoffs associated with each of these approaches:

The MPEG standard [6] specifies a maximum interval of 0.1 s for transmission of timestamps in the MPEG-2 Transport Stream. Therefore, in all the schemes we consider below, we assume that the timestamping interval T_c is always chosen within this bound.

Scheme 1: Forcing PCR values to stay on one side: The best case in the timestamp process is when the timer period is selected such that the transport rate of the MPEG stream is an exact multiple of the timestamping rate, that is, the ratio T_c/T_p

is an integer. In this case, the PCR values will always fall in either the odd-numbered or the even-numbered transport packets, thus eliminating packetization jitter altogether. Hence, the quality of the recovered clock is similar to that under the PCR-aware case. In practice, however, it is difficult to maintain the timestamping interval precisely as a multiple of the transport period, because of oscillator tolerances and various quantization effects. These effects may cause the PCR timestamp values to switch polarity at a very low frequency in the packet stream, degrading the quality of the recovered clock over the long term. In addition, loss of packets containing PCR values may cause timestamps to change polarity, that is, an odd-indexed PCR packet may become even-indexed or vice-versa. These effects are studied in the simulation experiments of Section IV.

Scheme 2: **Forcing PCR's to change boundary at high frequency:** From the analysis of the previous section, it is clear that the maximum frequency of changes in timestamp position in the packet stream occurs when the timestamping interval T_c satisfies the equality

$$T_c = (2k+1)T_p \tag{39}$$

where T_p is the transport period of the signal and k is any nonnegative integer. If T_c can be chosen precisely to satisfy this equality, the timestamped transport packets will occupy alternate (even/odd) positions in the AAL packet stream. The resulting pattern-switch signal is a square wave with the maximum possible frequency among all possible choices of T_c in the range from $2kT_p$ to $2(k+1)T_p$.

Just as in the previous scheme, it is difficult to set T_c precisely to satisfy (39). However, in this case it is not necessary to maintain T_c precisely. In the light of the analysis in the previous section, if the value of the timer period falls in the interval $[(2k+1/2)T_p, (2k+3/2)T_p]$, the frequency of the resulting pattern-switch pulse is still close to the case when $T_c = (2k+1)T_p$. This allows some tolerance for the clocks. Another significant advantage of this scheme is that random losses of packets containing timestamps are unlikely to affect the quality of the reconstructed clock. These hypotheses are verified in our simulation experiments in the next section.

Scheme 3: **Random setting of timer period:** In this case, the period of the timestamping timer is set to an arbitrary value, chosen randomly. The same timestamping interval is chosen for the entire packet stream, resulting in a deterministic pattern-switch signal at the input of the receiver.

From our analysis of the previous section, the frequency of the pattern switch signal depends on the relative magnitudes of T_c and T_p. Thus, this scheme needs to be used only when the transport rate of the MPEG signal is not known, since a more intelligent choice can be made when T_c is known.

Scheme 4: **Random timer period:** Another alternative when the transport rate is not known is to randomize the timestamping interval, by setting the timer each time to a value drawn from a random distribution. In the previous section, we showed that adequate quality can be maintained for the receiver clock when the timestamping interval is chosen such that the resulting PCR polarity changes in the packet stream exceeds a minimum rate. Although the analysis was based on modeling the PCR polarity changes with a random telegraph process, in practice similar results can be obtained by choosing the timer period from an exponential distribution. Results from our simulation experiments in the next section indicate that an exponentially-distributed timer period results in almost the same quality for the recovered clock as compared to the case when the PCR polarity changes according to the random telegraph process.

Similar to Scheme 2, this solution does not suffer from degradation of clock quality in the presence of random packet losses. Thus, Scheme 4 is useful when the transport rate T_p is not known with adequate precision.

In summary, Scheme 2 is the preferred scheme when the transport rate of the MPEG signal is known precisely, while Scheme 4 may be used when the transport rate is not known. In the next section, we evaluate the four schemes using both synthetic and real MPEG-2 traces to investigate the characteristics of the recovered clock signal at the receiver under various conditions.

IV. Experimental Results

To evaluate the timestamping schemes discussed in the last section, we conducted a number of simulation experiments with both synthetic and real MPEG-2 traces. In this section, we describe the simulation experiments and compare the schemes based on the quality of the recovered clock at the receiver.

Our experiments in this section fall into two sets: the first set is based on a synthetic trace in which the PCR timestamps were generated using the particular timestamping scheme being evaluated, and the second based on actual MPEG-2 traces which were pre-processed by inserting PCR timestamps based on the various timestamping schemes. In all cases, the traces were fed to a simulation model of an MPEG-2 Systems decoder and the behavior of the reconstructed clock was observed. In the second set of experiments, the packet streams were transported over a simulated ATM network and

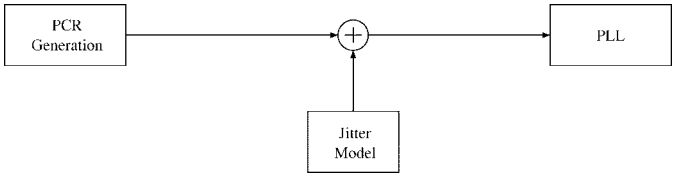

Fig. 8. Simulation model used in the first set of experiments to study the impact of PCR insertion rate on the receiver clock.

the PCR-unaware scheme was used for encapsulation at the Adaptation Layer. Finally, to study how the jitter generated by the network interacts with the packetization jitter introduced by the timestamping process, we conducted simulations in which the MPEG-2 streams were transported over a multihop ATM network model with significant amount of cross-traffic.

A. Experiments with Synthetic Traces

We first describe the models used and then present results from the simulations.

1) PCR Generation and Jitter Models Since our interest is only in studying the effect of the timestamp distribution in the MPEG-2 stream on the behavior of the recovered clock, we constructed a model of the timestamp generation process and used it to drive a model of the MPEG-2 Systems decoder. The PCR insertion frequency was varied based on the specific timestamping scheme simulated, with a lower bound of 10 Hz to meet the MPEG-2 specifications [6]. The transport rate was set to 4 Mbps, which results in an interpacket time T_p of $188*8/(4\times10^6) = 376$ μs. Thus, the maximum packetization jitter under the PCR-unaware scheme is 376 μs.

The model used in the first set of experiments is illustrated in Fig. 8. Two different approaches can be used to model the packetization jitter resulting from the timestamping process: in the first, packets are delayed according to the jitter model before they are delivered to the decoder. This is the traditional way of generating jitter in which the jitter changes the actual arrival times of the packets. Another way of modeling jitter is to assume that packets are delivered with equal spacing, and simulate the jitter by changing the PCR values in the packets once they are received by the decoder. Although the two methods are not equivalent, the convergence time and the steady-state error of the receiver PLL are almost the same in both cases [2]. For the sake of convenience, we chose the second method to model jitter. In our experiments, the jitter is added directly to the incoming PCR values at the decoder, in terms of the number of ticks of a 27-MHz clock.

2) Deterministic Timer Period: We first used a timer with a deterministic period to timestamp the MPEG-2 transport packets, and ran several experiments for various choices of the timer period. The results are summarized in Figs. 9 and 10. Recall that the transport rate is 4 Mbps, corresponding to an interpacket time of $T_p = 376$ μs. We set the timestamping timer period T_c close to or exactly at the transport packet boundaries to examine the three cases considered in Section II.

In the first case, the timer period was chosen very close, but not exactly equal to, a multiple of the interpacket time T_p. The value chosen was $T_c = 50T_p + 10^{-7}$, which corresponds to a timer frequency of $f_c = 1/T_c = 53.1912$ Hz. According to the analysis in Section II, the PCR timestamps in the resulting trace exhibit a forward drift of $\delta_f = 10^{-7}$ s. This, in turn, causes the PCR timestamp positions to switch at intervals of $2n_cT_c$, where $n_c = T_p/\delta_f$, that is, approximately 141.4 s. Thus the packetization delay experienced by transport packets containing PCR values form a square wave as illustrated in Fig. 9.

In the second case, we set the timer period T_c precisely to $(50+1)T_p$, or $f_c = 52.1485$ Hz. According to the analysis of Section II, this maximizes the frequency of PCR polarity changes in the packet stream, for the range of $50T_p \leq T_c \leq 52T_p$. Thus, the frequency of PCR polarity changes is $f_c/2 = 26.0743$ Hz.

In the third case, T_c was set slightly lower than an integer multiple of the interpacket time, by choosing $T_c = 52T_p - 10^{-7}$, or $f_c = 51.1459$ Hz. This results in a backward drift of $\delta_b = 10^{-7}$ s for the timestamp values. The delay distribution resulting from this choice is identical to that shown in Fig. 9.

Assuming the MPEG-2 stream carried NTSC video, Fig. 10 shows the frequency of the recovered NTSC clock at the receiver, in terms of deviation from the ideal 3.58-MHz value. When the low frequency drift is present in the timestamp values, that is, in the first and third cases above, the frequency of the recovered clock does not meet the maximum allowed variation of ±10 Hz in NTSC. In the second case, however, the recovered clock stayed well within the specifications.

For completeness, we also simulated two additional cases: the first was with a timer period of exactly $50T_p$ ($f_c = 53.1914$ Hz), which forces all the PCR timestamp values to stay on one side of the packet boundary. As seen from Fig. 10, this resulted in a perfect clock at the receiver. In the second case, we increased the timestamping rate almost by a factor of 10, but still maintaining the drift in timestamp values. This did not improve the clock behavior in Fig. 10, asserting that the solution to the jitter problem lies in fine-tuning the timestamping frequency, rather than choosing a high frequency of timestamping.

We also conducted several experiments with the synthetic trace where the timestamping intervals were chosen randomly [13]. In summary, the results from these experiments enable us to conclude that a random timestamping process with an adequately large timestamping rate can provide the desired quality for the recovered clock. An interesting observation from these experiments is that, in general, a higher rate of the exponential distribution (smaller timer period) results in better quality of the reconstructed clock, which may not be true in the case when the timer period is chosen deterministically. The reason is that the random timestamp distributions resulting from the former avoids the low-frequency patterns caused by the latter.

B. Simulations with Real MPEG-2 Traces in an ATM Network Simulation Model

To further validate our results, we performed a number of simulation experiments with real MPEG-2 Transport Stream traces. To simulate the effect of the timestamp distribution on the clock recovery process, the PCR values were inserted

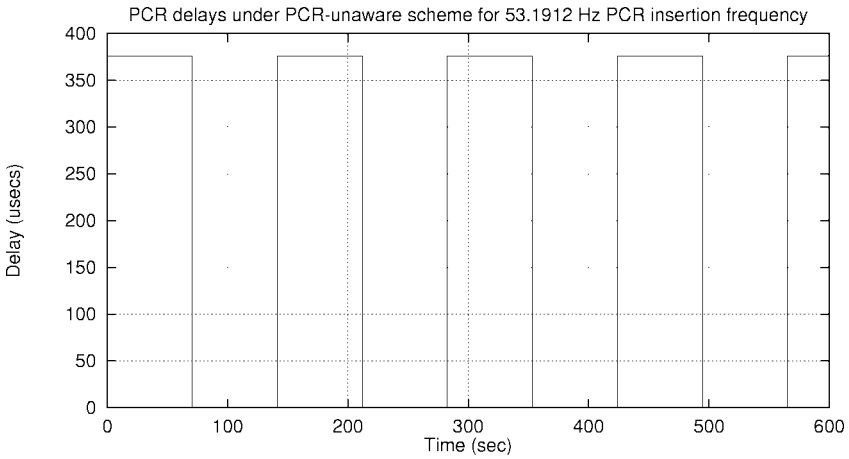

Fig. 9. Delays of transport packets containing PCR values due to packetization jitter with a PCR insertion rate of 53.1912 Hz.

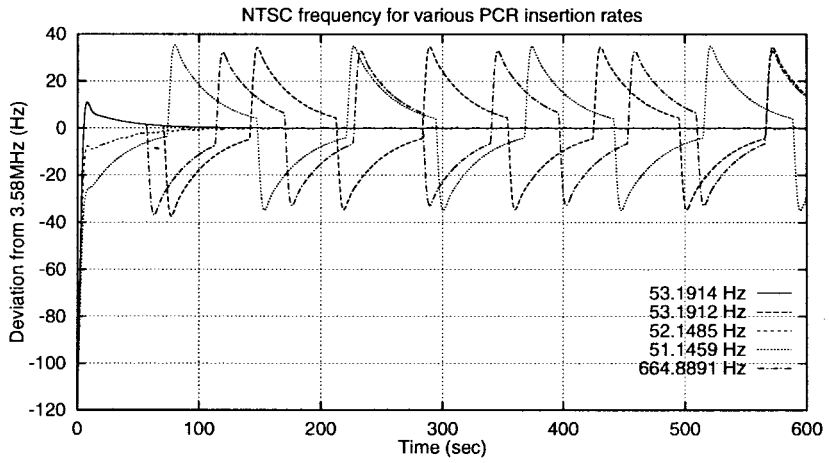

Fig. 10. NTSC color subcarrier generation frequency for varying PCR insertion frequencies. Transport rate is 4 Mbps.

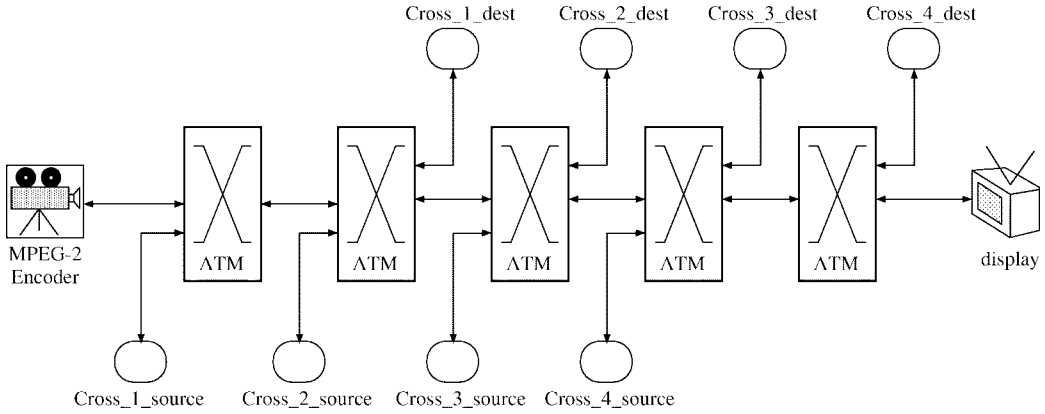

Fig. 11. Network topology used in the experiment.

in the traces using a timer with tunable frequency. Three traces were used throughout the experiments; in the first one, all the transport packets were inserted into even-numbered packets (referred to as "One Side" in the experiments). In the second trace, the PCR values were added so as to introduce a small forward drift when encapsulated using the PCR-unaware scheme (referred to as "Forward Drift" in the experiments). Finally, in the third trace, the transport packets containing PCR timestamps fall into alternating odd and even indices (referred to as "Max. Freq." in the experiments). The oscillator used for the timer was simulated with a 2 ppm frequency variation, to model the effect of clock drift. The transport rate of each trace was 6.5 Mbps. Each trace contained a single PAL video signal.

To model the influence of network-induced jitter, the trace was sent through an ATM network model with a large number of ON-OFF sources whose ON and OFF periods were dis-

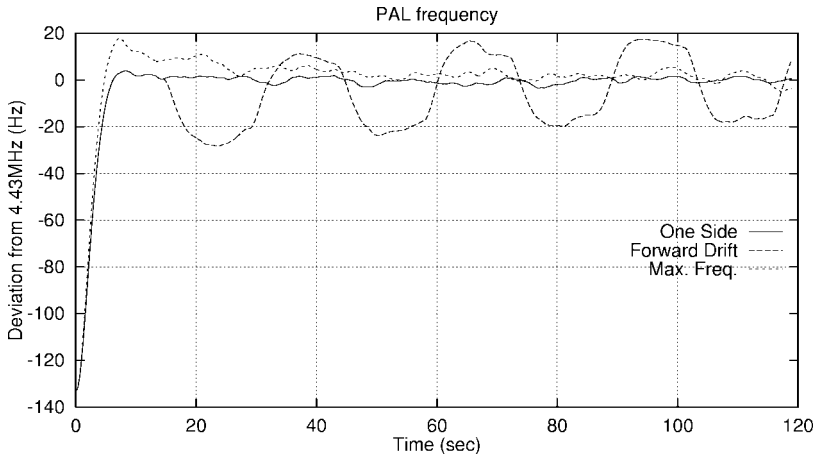

Fig. 12. PAL color subcarrier generation frequency for varying PCR insertion frequencies for the experiment.

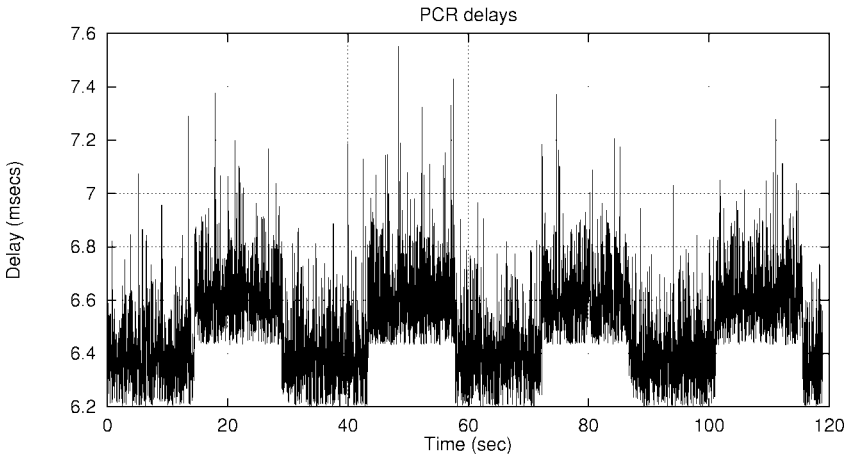

Fig. 13. PCR delays of "Forward Drift" trace for the experiment.

tributed exponentially. The network topology used is shown in Fig. 11. It consists of five cascaded ATM switches. The switch nodes are nonblocking, output-buffered crossbar switches. The MPEG-2 Transport Stream traces are sent through all the cascaded switches to the display device at the other end. At each hop of the network, the end-to-end MPEG stream shares the network link with cross traffic generated by a set of cell sources. All the cross-connections are between nodes that are connected to adjacent ATM switches. The propagation delay for each network link was set to 1 ms.

1) Experiment with Moderate Cross-Traffic: This experiment was performed to test whether jitter introduced by the network can mask the effect of the square wave pattern generated by the packetization jitter. In this experiment, 30 ON–OFF sources from each cross-traffic node were multiplexed with the MPEG-2 stream at each network link. The overall load on any downstream output link of the ATM switches was set to approximately 60%, resulting in a 10-Mbps aggregate rate of the ON–OFF sources per hop, or 0.334 Mbps load per source.

As seen in Fig. 12, the recovered clock is still affected by the square-wave pattern of the packetization jitter when the "Forward Drift" trace was used, resulting in unacceptable quality of the recovered clock. The recovered PAL color subcarrier generation frequencies for the other traces are also shown in Fig. 12. The "One Side" and "Max. Freq." traces produce similar variations in the recovered clock which stays within the acceptable levels when the PLL becomes locked.

The plot of the delays experienced by transport packets containing PCR's for the "Forward Drift" case contains the square shape due to the forward drift described in Section II and is not masked by the jitter induced by the network. This is illustrated in Fig. 13. However, the PCR delays in the "Max. Freq." case are spread out and the two delay zones due to the PCR-unaware scheme can no longer be distinguished clearly (Fig. 14). The maximum jitter observed was approximately 1.5 ms. Finally, in the "One Side" case, the PCR delays have similar shape as in the "Max. Freq." case but with a lower average value (Fig. 15). This results in the slightly better quality of the reconstructed clock in the "One Side" case compared to the "Max. Freq" case, although in both cases the PAL specifications were not violated (Fig. 12).

The interested reader is referred to [13] for two additional experiments with no cross-traffic: The first one studies the behavior of the clock recovery process at the decoder due to packetization jitter only, resulting from the use of different schemes to timestamp the MPEG-2 trace. In that experiment,

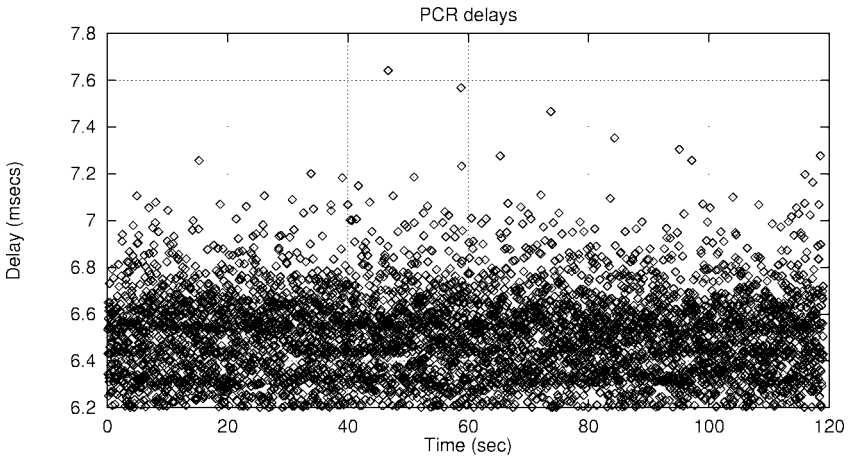

Fig. 14. PCR delays of "Max. Freq." trace for the experiment.

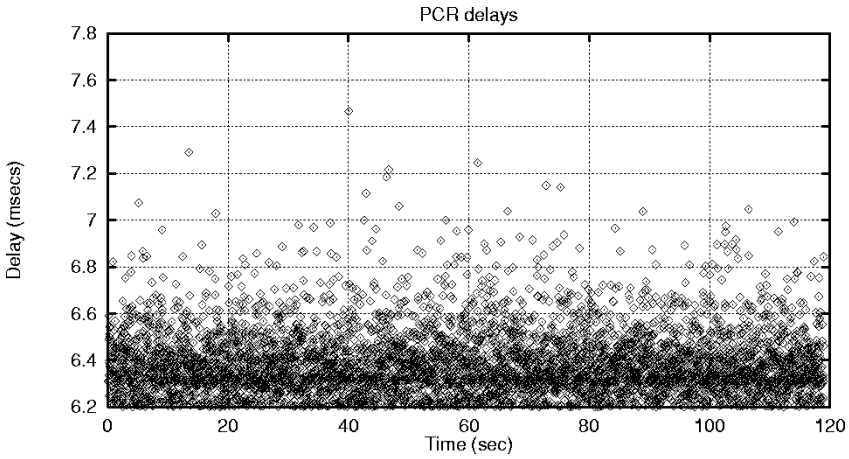

Fig. 15. PCR delays of "One Side" trace for the experiment.

we verified that the "Forward Drift" trace results in the worst performance of the recovered clock compared to the other two traces. The second experiment deals with the effect of random losses on the recovered clock for all three traces ("One Side", "Forward Drift", and "Max. Freq" traces). The results showed that the "Max. Freq." case is the only one in which the reconstructed clock is not influenced by the losses in transport packets containing PCR values. The other two cases ("One Side" and "Forward Drift") resulted in major disturbances in the recovered clock.

V. Conclusion

Our primary contribution in this paper is in providing guidelines for selecting the timestamping interval for transmission of PCR timestamps in a packetized MPEG-2 Transport Stream. Based on a systematic analysis of the jitter introduced by the timestamping process at the receiver, we identified three approaches for setting the timer used to derive the timestamps. In the first approach, the timer period is set precisely so that the transport rate of the MPEG stream is an exact multiple of the timestamping rate. This completely eliminates packetization jitter, but is difficult to implement in practice because of the precision required in the timer setting. In addition, loss of packets carrying timestamp values can cause the PCR values in the packet stream to switch position, affecting the quality of the recovered clock.

The second approach is to fine-tune the timer period to maximize the frequency of changes in PCR polarity. To maximize the frequency, the timestamping interval must ideally be set to $T_c/(2k+1)$, where k is any nonnegative integer and T_c the inverse of the transport rate in packets per second. This causes consecutive PCR values in the packet stream to alternate in polarity. This scheme has the advantage that, even when the timer cannot be set precisely to $T_c/(2k+1)$, the frequency of PCR polarity changes in the packet stream is still close to ideal. In addition, the scheme is robust in the presence of packet losses. Hence, this is the preferred scheme when the timestamps are generated with a fixed period.

When the transport rate of the MPEG-2 stream is not known and/or when a deterministic timer period is not practical, generating timestamping intervals randomly (with a certain minimum rate) can still provide adequate quality for the recovered clock. The quality of the receiver clock in this case depends on the process of PCR polarity changes, which, in turn is dependent on the distribution of the timestamping interval. By modeling the PCR polarity changes with a random telegraph process, we were able to derive lower bounds on the rate of PCR polarity changes to achieve the desired clock

quality. Although derivation of a corresponding bound for the timestamping interval is not straightforward, the former provides us with a guideline for setting the random timestamping interval.

The solutions in this paper apply in cases where the recovered clock must meet stringent requirements on its stability, as in the case when the recovered clock is used to generate the color subcarrier at the receiver. However, it must be pointed out that the stability of the reconstructed clock is not as critical in some environments. An example is in low-end systems such as set-top boxes for connection to home TV sets without component inputs. In some cases, although clock recovery is performed for the central 27-MHz clock, the set-top boxes use independent oscillators for the NTSC or PAL subcarrier, relaxing the requirements significantly. A second example is in high-end systems where the reconstructed composite signal must conform to full specification either to comply with FCC broadcasting specifications or because it will be further processed in a TV production environment. In this case, an elastic buffer may be used to hide the effect of frequency drifts between the recovered clock and the local clock by dropping or repeating entire frames when necessary. With such an approach, the stability of the recovered 27-MHz clock becomes less important.

ACKNOWLEDGMENT

The authors would like to thank Prof. Yiannis Kontoyiannis at Purdue University for his valuable input on the analysis, and the anonymous reviewers for their helpful comments.

REFERENCES

[1] I. F. Akyildiz, S. Hrastr, H. Uzunalioglu, and W. Yen, "Comparison and evaluation of packing schemes for MPEG-2 over ATM using AAL5," in *Proc. ICC'96*, June 1996, vol. 3, pp. 1411–1415.

[2] G. F. Andreotti, G. Michieletto, L. Mori, and A. Profumo, "Clock recovery and reconstruction of PAL pictures for MPEG coded streams transported over ATM networks," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 5, pp. 508–514, Dec. 1995.

[3] CCIR 624, *Specifications for the Color TV Signal*, Tab. II, Item 2.11.

[4] P. Hodgins and E. Itakura, "The issues of transportation of MPEG over ATM," ATM Forum, ATM94-0570, July 1994.

[5] P. Hodgins and E. Itakura, "VBR MPEG-2 over AAL5," ATM Forum, ATM94-1052, Dec. 1994.

[6] International Organization for Standardization, *Information Technology—Generic Coding of Moving Pictures and Associated Audio: Systems, Recommendation H.222.0, ISO/IEC 13818-1*, Ist Edition, Apr. 1996.

[7] Y. Kaiser, "Synchronization and dejittering of a TV decoder in ATM networks," in *Proc. Packet Video Workshop'93*, 1993, vol. 1.

[8] A. L. Garcia, *Probability and Random Processes for Electrical Engineering*, 2nd ed. Reading, MA: Addison-Wesley, May 1994.

[9] H. Meyr and G. Ascheid, *Synchronization in Digital Communications*, vol. 1, Wiley Series in Telecommunications. New York: Wiley, 1990.

[10] J. Proakis and D. G. Manolakis, *Introduction to Digital Signal Processing*. New York: Macmillan, 1988.

[11] R. P. Singh, S.-H. Lee, and C.-K. Kim, "Jitter and clock recovery for periodic traffic in broadband packet networks," *IEEE Trans. Commun.*, vol. 42, pp. 2189–2196, May 1994.

[12] C. Tryfonas, "MPEG-2 transport over ATM networks," M.S. thesis, Univ. California, Santa Cruz, Sept. 1996 [Online.] Available http://www.cse.ucsc.edu/research/hsnlab.

[13] C. Tryfonas and A. Varma, "Timestamping schemes for MPEG-2 Systems layer and their effect on receiver clock recovery," Tech. Rep. UCSC-CRL-98-2, Dept. Comput. Eng., Univ. California, Santa Cruz, May 1998.

[14] C. Tryfonas and A. Varma, "A restamping approach to clock recovery in MPEG-2 Systems layer," in *Proc. ICC'99*, June 1999.

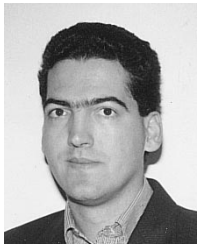

Christos Tryfonas received the Diploma in computer engineering and information sciences from the University of Patras, Greece, in 1993, and the M.S. and Ph.D. degrees in computer engineering from the University of California, Santa Cruz, in 1996 and 1999, respectively.

He is a Principal Member of Technical Staff, Sprint Advanced Technology Laboratories, Burlingame, CA. His research interests lie in the areas of multimedia networking, congestion control, quality-of-service (QoS), and traffic management in broad-band networks.

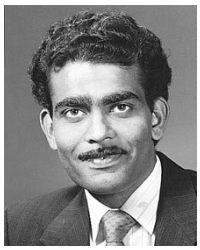

Anujan Varma (M'86) received the Ph.D. degree in computer engineering from the University of Southern California, Los Angeles, in 1986.

He was with the IBM Thomas J. Watson Research Center, Yorktown Heights, NY, until 1991. he is now a Professor of computer engineering at the University of California, Santa Cruz. His research interests lie in the areas of packet switching, traffic management in broad-band networks, and optical networks. He has published more than 70 papers and holds six patents. He has consulted extensively for the networking and semiconductor industries.

Dr. Varma has received several awards, including the National Science Foundation Young Investigator Award, IEEE Darlington Award, and a teaching innovation award from the University of California.