

COMMON TIME REFERENCE FOR INTERACTIVE MULTIMEDIA APPLICATIONS

Mario Baldi and Yoram Ofek
Synchrodyne Networks, Inc.
75 Maiden Lane
New York, NY 10038

ABSTRACT

A delay of about 100 ms gives human communicators the feeling of live interaction. Since in a global network the propagation delay alone is about 100 ms, every other delay component, such as processing and queuing, should be kept as short as possible.

Moreover, the deployment of new high bandwidth multimedia applications will boost network traffic and consequently the demand for very high capacity transmission technologies, such as Wavelength Division Multiplexing (WDM). Networks will suffer (i) electronic switching bottlenecks among high-speed links and (ii) communications link bottlenecks between high capacity core technologies and low speed access technologies.

This paper addresses the design of interactive systems for applications such as toll quality telephony, videotelephony and videoconferencing, highlighting the benefits brought by the availability of *global common time reference* derived from GPS (Global Positioning System). Common time reference is essential to keep the user perceived delay within the 100 ms bound while avoiding the two above mentioned bottlenecks. The proposed solution can be applied to both IP and ATM networks, does not require changes to any of the existing protocols, and enables traffic aggregation in the core of the network—thus not requiring nodes to keep state information on microflows—while providing a guaranteed quality service to individual applications.

1. INTRODUCTION

The deployment of new high bandwidth multimedia applications will boost network traffic and consequently the deployment of very high capacity transmission technologies, such as Wavelength Division Multiplexing (WDM). On the other side, since multimedia services will have to be widely available, various “low speed” access technologies, such as wireless, DSL, and cable modem will be deployed. In this scenario networks will suffer (i) *electronic switching* bottlenecks and (ii) *communications link* bottlenecks, the latter being created by the bandwidth mismatch between high capacity core technologies and low speed access technologies.

Many interactive applications, such as, telephony, videotelephony and videoconferencing require at the receiver *continuous playing* of samples captured at the sender. Continuous playing implies a *constant delay* service to be provided at the application layer, i.e., where samples are acquired and played. Since some of the end-to-end delay components can vary during a session, specific action is required at the receiver to keep constant the end-to-end delay between the application layers, thus enabling continuous playing. Before samples are played, delay variations should be “smoothed out” by buffering the samples that have experienced (in the network and in the decoder) a delay shorter than the maximum. This introduces a *resynchronization delay* component that is typically the time spent in a *replay buffer* [4]. On exiting the

replay buffer, all the samples have experienced the same delay since the time they were acquired at the sender side. Such an overall delay is equal to or larger than the *delay bound* the system can guarantee.

This paper shows that the end-to-end delay bound can be minimized by a system with a *global common time reference*, independently of the underlying packet technology (e.g., ATM or IP) and the session rate. Resynchronization components can be kept small, e.g., 125 μ s, as well as the queuing delay in each switch, e.g., 250 μ s. Global time is used in two ways:

1. To implement time-driven priority (TDP) forwarding of packets, which
 - a. guarantees a maximum per-hop queuing delay below one millisecond, independent of the flow rate and the network load, also in bandwidth mismatch points;
 - b. enables the implementation of efficient packet switch architectures based on low complexity switching fabrics that do not require speedup with respect to input line capacity. This increases the scalability of switches and eliminates the electronic switching bottleneck.
2. To synchronize the acquisition of samples at the sender (i.e., video capture card) and their continuous playing at the receiver (i.e., video display) with one another and with the TDP forwarding.

Global time eliminates communications link bottlenecks since it completely avoids congestion, also in bandwidth mismatch points. Moreover, this paper describes how global time is beneficial in the implementation of highly scalable packet switches whose design is based on simple switching fabrics without any speedup with respect to input link capacity.

Section 2 describes how synchronous switches use the common time reference to implement Time-Driven Priority (TDP) forwarding and how such synchronous switches can be deployed in any network—one significant specific case being the Internet—together with traditional asynchronous switches. Moreover, the creation of Synchronous Virtual Pipes (SVPs) enables traffic aggregation while providing service guarantees to individual packet flows. Section 3 discusses the benefits stemming from the application of time-driven priority forwarding to voice traffic. The deployment of the same technology for the transmission of video is addressed in Sections 4 and 5. Section 6 describes the advantages of using the global time in the video capture card and the video display card. Scalability properties of synchronous packet switches are highlighted in Section 7. Finally, conclusions are drawn in Section 8.

2. GLOBAL TIME AND PERIODIC FORWARDING: TIME-DRIVEN PRIORITY

All packet switches are synchronized to a global common clock with a basic time period that is called *time frame* (TF). The TF

duration is derived from the UTC (coordinated universal time) second received from the GPS (Global Positioning System). For example, by dividing the UTC second by 8,000, the duration of each time frame is $T_f = 125 \mu s$; however, the time frame duration can be set (longer or shorter) as needed. Note that different links can have different time frame duration; for example, $12.5 \mu s$ for high capacity links or $500 \mu s$ for low capacity links. The accuracy of the GPS time signal can be very high; however, the time-driven priority forwarding scheme can operate correctly with time accuracy of about half a TF, e.g., $60 \mu s$.

TFs are grouped into a *time cycle*; Figure 1 shows an example of a time cycle that contains 100 TFs, i.e., there are 80 time cycles in a UTC second. Time cycles are further organised in *super cycles*, each one typically equal to one UTC second. Since a super cycle is equal to one UTC second the insertion/deletion of leap second is possible without affecting existing call schedules. The described timing structure is useful to perform resource reservation in order to provide guaranteed services.




Figure 1: Global Common Time Reference

Thus, all switches around the globe have an identical time structure that is collectively called a *Common Time Reference* (CTR) and can be used to coordinate the acquisition of samples at the sender with the playing of them at the receiver, as discussed the subsequent sections. Moreover, the CTR enables the implementation of *Time-Driven Priority* (TDP) [1][2] for periodically forwarding real-time packets, for example inside IP and ATM networks. Real-time packets are *periodically granted the highest priority*, while “best effort” traffic is sent with lower priority. In order to guarantee that a packet can be forwarded during a predefined TF, the right to transmit the corresponding amount of bits on links during the TFs is reserved beforehand by applications. *Periodic* forwarding indicates that the reservation, and hence the forwarding pattern, repeats itself in every time cycle and in every super cycle. TDP guarantees that the end-to-end delay jitter is less than one TF and that reserved real-time traffic is transferred from the sender to one or more receivers with no loss due to congestion.




Figure 2: Synchronous Virtual Pipe (SVP) and Access Bandwidth Brokers (ABBs).

TF reservations for each micro-flow should be made in each node in order to guarantee constant delay and no loss due to congestion. Even though no per-micro-flow information is needed to forward packets, the processing load introduced by signaling can affect scalability in the network core. Maximum scalability can be

obtained by reserving TFs for semi-permanent Synchronous Virtual Pipes (SVPs). Packets traveling through the core are carried within an SVP. In this scenario, when an application attempts to reserve network resources for its packet micro-flow, its signaling message is not forwarded to core switches. Rather, as shown in Figure 2, it is processed by the Access Bandwidth Broker (ABB) at the ingress of the SVP the micro-flow is to traverse. If the available TFs—among those reserved for the SVP—are compatible with the application’s request (i.e., an end-to-end schedule is feasible), the ABB assigns the corresponding TFs to the new micro-flow and forwards the signaling request to the ABB at the other end of the SVP. Even though intermediate switches traversed by the SVP are not aware of the TFs reserved to each micro-flow, they can forward packets according to the TDP forwarding principles since a reservation had been placed for the SVP. As a result, packets experience the same guaranteed quality service as if resources had been reserved to their micro-flow, but only the ABB is aware of the TF assigned to the micro-flow.




Figure 3: Interoperation between TDP Networks and Asynchronous Networks

In case a micro-flow gets to the device at the ingress of the SVP through an asynchronous network, as in the configuration shown in Figure 3, the border device buffers each packet until the proper TF—among those reserved to the SVP—assigned to the micro-flow. Notice that SVPs can be set up over multiple synchronous subnetworks interconnected by asynchronous ones. Only the border device at the boundary with an asynchronous network, possibly co-located with the ABB, needs to be aware of the TF assigned to the micro-flow.

3. CONSTANT BIT RATE SERVICES: TOLL QUALITY TELEPHONY

Assuming constant size packets, each phone call consists of a packet flow at a constant rate. If TDP is deployed and, for example, 100 samples are transmitted in each packet, resources should be reserved for the transmission of one packet in each time cycle, as shown in Figure 4. Packets containing samples will be delivered to the destination with no loss due to congestion and within a time depending primarily on propagation delay.




Figure 4: Periodic Forwarding of Voice Packets

The queuing delay introduced by TDP is independent of the packet size, session rate, and the amount of resources reserved to it. This is unlike other asynchronous scheduling algorithms, such as, weighted fair queuing, in which the queuing delay is proportional to the packet size and inversely proportional to the rate. Such asynchronous algorithms possibly require overallocation of resources to meet the end-to-end delay requirement [3]; in some cases where propagation delay is large, voice compression may lead to a resource allocation larger than without compression. Paradoxically, it may be the case that where voice compression would be most beneficial (e.g., trans-continental phone calls) it becomes useless.

On one hand, since TDP does not require overallocation, the bandwidth reduction stemming from compression can provide full benefit. Moreover, the low queuing delay bound in the network allows the sender to:

1. spend more time performing particularly effective compression, and
2. gather samples in larger packets, which introduces a longer packetization delay, in order to reduce the percentage overhead stemming from the transmission of fixed size headers (40 bytes in IP networks).

On the other hand, since TDP allows full link utilization with real-time traffic, compression can be avoided in end-systems that have low processing power and must be kept simple.

TDP is beneficial to echo cancellation in two ways:

1. When the propagation delay is below 50 ms, echo cancellation can be avoided, and
2. When echo cancellation is needed, the small jitter (bounded by 125 μ s) simplifies it.

Finally, since unused reserved capacity can be exploited to transmit “best effort” traffic, the overall transfer capacity of the network benefits from phone sources performing silence suppression.

4. PERIODIC BURSTY SERVICES: VIDEOTELEPHONY AND VIDEOCONFERENCING

Videotelephony and videoconferencing, like telephony, rely on continuous playing at the receiver of samples acquired at a fixed rate at the sender. Video frames have two main differences from voice samples.

1. The sampling rate is usually lower, from a few to 30 frames per second.
2. The amount of bits required to encode each video frame sample is much larger, at least a few kilobits.

When circuit switching is used to transfer video frames, the encoder is operated in such a way that it produces a constant bit rate flow. This is required in order to fully utilize the channel allocated to the session. Consequently, the transmission delay of a single video frame is the time between two successive video frames. This is because the transmission of the current video frame should continue, in a constant rate, until the next video frame is ready. For example, if the sampling rate is ten video frames a second, the transmission delay alone is 100 ms, which is unacceptable in interactive applications.

The elimination of such a long transmission delay is achieved by transmitting the captured video frame as a short burst. Packet switching allows burst transmission of video frames in packets, i.e., as shown in Figure 5, a video frame is captured, put into a packet, and then transmitted as a burst into the network. Therefore, the only way to transmit video frames for interactive applications with minimum delay is over packet-switched network.




Figure 5: Periodic Bursty Transmission of a Video Stream

Since video frames are captured periodically, in order to minimize the delay bound, *periodic resource allocation* with *periodic transmission synchronized with their capture* are required. TDP is the only known way to satisfy those requirements, while guaranteeing absence of loss with minimum delay bound.

It is worth noting that even though all the video frames are not encoded with exactly the same amount of bits, the capacity reserved on the links is not wasted since it is used to forward “best effort” (i.e., non-reserved) traffic. Avoidance of loss due to congestion is guaranteed to all the video frames, provided that the amount of bits encoding them does not exceed the reservation.

5. COMPLEX PERIODICITY: MPEG VIDEO

Some video encoding schemes, like MPEG, encode frames with significantly different amounts of bits in a *periodic* fashion. MPEG encodes pictures in one of two different ways¹: intra-frame coding, which produces *I-frames*, and predictive coding, which produces *P-frames* that are typically from 2 to 4 times smaller than I-frames. The slower moving is a scene, like in a videoconference, the smaller the amount of bits produced for each P-frame.

It may be inefficient to transfer such a compressed video stream over a constant bit rate channel, e.g., the one provided by a circuit switched network (see [4] for a detailed discussion). If the encoder is operated in such a way that it produces a constant bit rate flow, it can introduce a delay up to the time between two successive I-frames.




Figure 6: Complex Periodicity Scheduling of MPEG Video Stream

Complex periodicity scheduling allows MPEG video frames to be transmitted as soon as they are encoded, analogously to what is described in Section 4 for fixed size video frames. TDP together with global time facilitates the realization of complex periodicity

¹ Actually, a third type of encoding, called bi-directional predictive coding exists, but since it introduces a delay of multiple video frame periods, it is not usable in this context given the 100 ms end-to-end delay bound requirement.

scheduling, which provides deterministic quality of service guarantees to variable bit rate traffic. In complex periodicity scheduling the amount of transmission capacity reserved on the links traversed by a session varies in a repetitive manner. This is what is needed to transmit an MPEG encoded video stream as shown in Figure 6. A related study [5] showed that an MPEG encoder can be successfully implemented in a way that encoded video frame size never violates the resource allocation.

6. SYNCHRONIZING CAPTURE AND DISPLAYING

In a videophone or videoconference call the delay perceived by the user is the time from when a picture is captured by the sender camera until when it is displayed on the receiver monitor. The global time can reduce the delay perceived by users if both the capture card and the display use it, and both are synchronized with the TDP scheduling in the network.

Synchronizing the frame grabber to the network The encoder must produce the encoded frame right before the TF during which capacity has been reserved for transmission on the outgoing link. Otherwise, a *shaping delay* is introduced at the sender. Synchronizing the capture card with the network minimizes the user perceived delay.

Synchronizing the video card and the display to the network A dual situation exists on the receiver side. The receiver decodes each picture and stores it into the video buffer that is periodically scanned by the video card to refresh the image on the screen. If the decoder does not write the next picture in the video buffer just in time for the video card to scan it, a *presentation delay* is introduced. This can be avoided if the video card and the display are synchronized with the network.

7. SCALABILITY OF SYNCHRONOUS SWITCHES

The deployment of time to control packet forwarding has a twofold impact on switch scalability.

TDP enables to control traffic patterns across each switch, i.e., the maximum number of packets that during each TF are to be moved from every input to the same output. This can be leveraged of in the switch design: optimal input-output switching is obtained with low (x2 or lower) speed up in the switching fabric, or even without speed up at all. Instead, asynchronous switches require high speed up in order to achieve high throughput.

A non-blocking fabric allows any possible input-output connection *at any time*; a simpler (blocking) fabric allows only a limited number of simultaneous input-output connections. With TDP, packet arrival can be controlled in a way that incompatible input-output connections are not required during the same TF, thus avoiding unfeasible switching configurations. Thus, thanks to the increased flexibility introduced by the time dimension, it is possible to design synchronous packet switches based on blocking switching fabrics that achieve a throughput comparable to that of asynchronous switches with non-blocking fabrics. Or, in other words, given the state of the art aggregate switching fabric capacity, synchronous switches can accommodate more ports, each at a higher capacity, than asynchronous switches.

8. CONCLUSIONS

This work shows how global time can be used to minimize the end-to-end delay for applications that require at the receiver continuous playing of samples captured at the sender.

When a TDP network is deployed to carry voice calls, compression can be fully benefited to reduce the amount of link capacity used by each call. This is not the case with other asynchronous queuing schemes that possibly require overallocation to satisfy end-to-end delay requirements.

When dealing with videotelephony, encoded video frames are transmitted in bursts of packets with controlled delay and no loss. The delay perceived by users is lower than the one obtained by carrying video calls over a circuit switching network which requires delay to be introduced for smoothing out the burstiness of the video source.

Global time and TDP forwarding offer the only solution for the transmission of video frames also when they are encoded with a highly variable amount of bits, such as with MPEG. Each video frame can be transmitted in a burst of packets as soon as it is encoded with no shaping delay and no loss due to congestion. Through complex periodicity scheduling, resource reservation is fitted to the size of encoded video frames thus leading to efficient resource utilization.

The paper also shows that global time can be beneficial when its deployment is not ubiquitous, as it will be the case especially in the initial transition phase when synchronous switches will start to be installed on backbones and will coexist with legacy asynchronous ones. Moreover, the concept of synchronous virtual pipe enables flow aggregation, which is particularly useful in the core of global networks in order to increase scalability.

Finally, global time is beneficial to increase the scalability of packet switches. Switch scalability is essential to satisfy the increasing demand for electronic switching capacity that will be boosted by new bandwidth-consuming multimedia applications. Moreover, when very high capacity backbones will be built to satisfy the above demand, the congestion free operation guaranteed by TDP will avoid potential bottlenecks at the boundary with lower speed access networks.

REFERENCES

- [1] C-S. Li, Y. Ofek, A. Segall and K. Sohraby, "Pseudo-Isochronous Cell Switching in ATM Networks," Computer Networks and ISDN Systems, No. 30, 1998.
- [2] C-S Li, Y. Ofek and M. Yung, "Time-driven Priority Flow Control for Real-time Heterogeneous Internetworking," IEEE INFOCOM'96, Apr. 1996.
- [3] M. Baldi, D. Bergamasco, F. Risso, "On the Efficiency of Packet Telephony," 7th IFIP International Conference on Telecommunication Systems, Nashville, Tennessee, USA, March 1999.
- [4] M. Baldi and Y. Ofek, "End-to-end Delay Analysis of Videoconferencing over Packet Switched Networks," in IEEE INFOCOM 1998, Apr. 1998.
- [5] M. Baldi, Y. Ofek, "End-to-end Delay of Videoconferencing over Packet Switched Networks," RC 20669 (91480), IBM - T. J. Watson Research Center, Yorktown Heights, NY, USA, Dec. 1996.