

Uncoordinated Real-Time Video Transmission in Wireless Multi-Code CDMA Systems: An SMPT-Based Approach (Extended Version) *

Frank Fitzek[†] Martin Reisslein[‡] Adam Wolisz

July 2002

Abstract

We consider the real-time transmission of encoded video from distributed, uncoordinated wireless terminals to a central base station in a multi-code CDMA system. Our approach is to employ the recently proposed Simultaneous MAC Packet Transmission (SMPT) approach at the data link layer (in conjunction with UDP at the transport layer). We consider the real-time transmission of both video encoded in an open-loop (i.e., without rate control) and video encoded in a closed loop (i.e., with rate control). We conduct extensive simulations and study quantitatively the trade-off between the video quality, the transmission delay (and jitter), and the number of supported video streams (capacity). We find that the simple to deploy SMPT approach achieves significantly higher video quality and smaller delays than the conventional sequential transmission approach, while ensuring a high capacity. In typical scenarios, with SMPT the probability of in-time video frame delivery is more than twice as large as with sequential transmission (for given delay bounds). Our results provide guidelines for the design and dimensioning of cellular wireless systems as well as ad-hoc wireless systems.

Keywords: Multi-code CDMA; Rate Control; Real-Time; Simultaneous MAC Packet Transmission; Uplink Transmission; Video.

1 Introduction

Video traffic is expected to account for a large fraction of the traffic in future wireless networks. Generally, the transport of video over wireless links is a very challenging problem. This is due to (1) the stringent playout deadlines of the video frames (as well as the variability of the frame sizes for open-loop encodings), and (2) the unreliability of the wireless links. For streaming services (e.g., the web-based streaming of prerecorded video clips) which do not require real-time interactions, these challenges can be overcome by relaxing the timing constraints with receiver side buffering and taking advantage of multi-user diversity. Multi-user diversity is based on the observation that in a typical wireless system with multiple users, at any point in time, some users are experiencing favorable transmission conditions on their wireless links, while others are experiencing adverse transmission conditions. (This is due to the

*Please direct correspondence to M. Reisslein.

[†]F. Fitzek and A. Wolisz are with the Telecommunication Networks (TKN) group, Dept. of Electrical Eng., Technical University Berlin, Sekr. FT5-2, Einsteinufer 25, 10587 Berlin, Germany, (e-mail: {fitzek, wolisz}@ee.tu-berlin.de).

[‡]M. Reisslein is with the Dept. of Electrical Eng., Telecommunications Research Center, Arizona State University, Goldwater Center, MC 7206, Tempe, AZ 85287-7206, USA, phone: (480)965-8593, fax: (480)965-8325, (e-mail: reisslein@asu.edu).

typically location-dependent, time-varying, and bursty errors on the wireless links.) The basic idea of multi-user diversity is to transmit at any point in time only over the links that are currently experiencing favorable transmission conditions. If a link is currently experiencing adverse transmission conditions (which would require high power levels or several re-transmissions), then no transmissions are scheduled for this link and the receiver continues video playback from its buffer. (See, for instance, [7] for a detailed study of such a system.) Applications that require real-time interactions (e.g., video conferencing, tele-medicine and games), however, allow only for very limited relaxation of the timing constraints (i.e., very limited receiver side buffering). In addition, taking advantage of multi-user diversity typically requires the central scheduling of packet transmissions. This central scheduling requires an established signalling and coordination structure, which, however, is not available in the emerging ad-hoc wireless networks. In summary, the *real-time* video transmission by *uncoordinated* wireless terminals is especially challenging.

In this paper we develop and evaluate simple, yet quite effective techniques for the real-time video transmission by distributed, uncoordinated wireless terminals in a common interference environment (e.g., local cluster in an ad-hoc network or cell in a cellular network). Our approach is to employ Simultaneous MAC Packet Transmission (SMPT) techniques, which have been recently developed and evaluated for the transmission of data traffic (e.g., FTP traffic). The basic idea of SMPT is to transmit multiple packets (on multiple CDMA codes) in parallel to make up for packets lost due to errors on the wireless channel. While SMPT techniques have been studied extensively for the transmission of data traffic without any fixed deadlines [5, 4], they have not yet been studied in the context of continuous media (such as video) with strict timing constraints (except for the initial study [9] which focused on the stabilization of the TCP throughput for video, see Section 1.1 for details).

In this paper we study the real-time transmission (with UDP as the transport protocol) of video encoded with rate control as well as of video encoded without rate control. Video encoded with rate control has typically only moderate bit rate variations on typically small time scales. On the other hand, encoding video without rate control — which requires less complexity (and saves cost and energy at the wireless device) — results in video traffic with large bit rate variations over many different time scales (including long time scales) [6]. In addition, video encoded without rate control typically exhibits long range dependence (or self-similarity). Consequently, we find that different variations of the SMPT technique are appropriate for the different types of encoded video. The so-called *slow-healing* SMPT mechanism, which resorts to transmitting multiple MAC packets in parallel only after suffering loss on the wireless link, is suited for the relatively smooth video encoded with rate control. On the other hand, the so-called *fast-start* and *slow-start* SMPT mechanisms, which transmit from the outset on multiple parallel codes, are suited for the bursty video encoded without rate control.

Throughout our study we pay close attention to implementation aspects. Our focus is on techniques that are simple to deploy in practical systems and give tangible performance improvements, rather than techniques that are optimized for performance at the expense of increased complexity. The presented SMPT mechanisms operate exclusively at the data link layer and do not require any higher layer information. Also, the SMPT mechanisms do not require any coordination among the wireless terminals in the cluster of an ad-hoc network or the cell of a cellular network. The SMPT mechanism running in a given wireless terminal schedules the MAC packet transmissions completely independently from the other wireless terminals in the cluster (or cell). The wireless terminals “feel” each other only through the interference level generated by their transmissions. By varying the number of used CDMA codes and monitoring the success of its own transmissions, the SMPT mechanism in a given wireless terminal probes the capacity of the cluster (cell). In typical scenarios the SMPT mechanism more than doubles the probability of in-time video frame delivery

This paper is organized as follows. In the following section we give an overview of related work on

video transmission in wireless environments. In Section 2 we describe the real-time transmission of video using the SMPT mechanisms. We first review the SMPT technique and then adapt it to the transmission of video traffic. We also define the performance metrics used in our performance evaluations and outline the simulation scenarios. In Section 3 we quantitatively evaluate the transmission of video encoded with rate control. Section 4 studies the transmission of video encoded without rate control. We summarize our conclusions in Section 5.

1.1 Related Work

The area of video transmission in wireless environments has attracted a great deal of attention recently and a large body of literature on the topic has emerged. Several schemes have been proposed for improving the video quality by employing adaptive video coding schemes, see for instance [2, 10, 11, 13, 20]. Our work is orthogonal to these adaptive video encoding schemes in that we adapt the scheduling of the transmissions of the MAC packets carrying the video (instead of the encoding of the video).

A hybrid scheme employing forward error correction (FEC) and automatic repeat request (ARQ) is proposed in [14]. The ARQ component in [14] does not transmit multiple packets simultaneously in response to lost packets. Thus, our SMPT approach is orthogonal to [14] in that it may be used as a refined ARQ component in the hybrid scheme of [14]. We note that an adaptive error recovery scheme for wireless video transmission is developed in [18]. An enhanced UDP transport protocol for wireless video transmission is developed in [24].

In [9] we studied the streaming of video using TCP as the transport protocol. This paper differs from [9] in two important aspects. First, [9] focuses on video streaming with delays on the order of one second. In this paper, on the other hand, we focus on real-time transmission with delays on the order of less than a few hundred milliseconds, which allow of real-time communication, as needed for interactive applications, such as video conferences, tele-medicine or games. Secondly, [9] studied a “reliable” video service, that temporarily pauses the video playout when the client’s consumption exceeded the supply of video information (and does not drop any video frames). This paper, on the other hand, considers a real-time (albeit lossy) video service, that trades of tight delay bounds for some small loss (whereas [9] guarantees no loss at the expense of large delay and play-back pauses).

In [7] we developed a prefetching scheme for the streaming of video in wireless environments. This prefetching scheme may be employed for downlink streaming as well as uplink streaming, and does also support real-time transmission when the tolerable delays are of the order of a few video frame periods. The fundamental difference between the scheme studied in this paper and the prefetching scheme [7] is that [7] requires centralized scheduling, whereas the scheme studied in this paper is for the uncoordinated transmission by distributed wireless terminals.

We note that video transmission in multi-code CDMA systems is also studied in [1]. The scheme proposed in [1] is similar to ours in that multiple codes are used in parallel to accommodate the variable sized video frames (of a given video). The main difference between [1] and our scheme is that [1] requires a significant amount of coordination among the videos being transmitted (for instance, the video streams are aligned such that a (typically large) Intracoded (I)-frame of one video stream does not coincide with the I-frame of another video stream). Our scheme on the other hand does not require any coordination among the ongoing video flows, and is thus well suited for wireless networks with little or no coordination among the wireless terminals, such as ad-hoc networks.

We finally note that mechanisms for the transmission of scalable video (i.e., video encoded into multiple layers) over wireless links are discussed in [12, 15, 23]. These mechanisms strive to schedule the individual layers so as to maximize the overall video quality. Throughout this paper we are focusing on video that is encoded into a single-layer (i.e., non-scalable).

2 Real-Time Video Transmission with SMPT

In this section we discuss the transmission of video traffic using the Simultaneous MAC Packet Transmission (SMPT) approach. At the sending wireless terminal the video is encoded into video frames. Each video frame is immediately passed down the networking protocol stack through the UDP transport protocol and the IP network protocol (each of which appends its header to the video frame). For simplicity, we assume that each video frame is encapsulated into a single transport layer segment. (Thus, we use the terms “segment” and “video frame” interchangeably in our discussion of the lower layer mechanisms.) At the data link layer, the segment (video frame plus UDP and IP protocol headers) is partitioned into fixed size Link-layer Packet Data Units (LPDUs). (Padding is used to fill the last LPDU for a given video frame.) With the standard *sequential* transmission approach, the LPDUs are transmitted in send-and-wait fashion using one single CDMA code. (Throughout, a slotted timing structure is assumed, where one CDMA code provides sufficient transmission capacity to transmit one LPDU in one slot of duration τ_{slot} .) An LPDU successfully received by the receiving terminal is immediately acknowledged, and the next LPDU is transmitted in the subsequent slot. (We assume that the acknowledgment for a successful LPDU is returned and processed before the next LPDU is sent in the next slot; this is feasible with typical hardware configurations of wireless communication systems [22].) If an LPDU is lost on the wireless link, the sender retransmits the lost LPDU until it is received successfully, and then moves on to the next LPDU (see [8] for details).

2.1 Slow-Healing SMPT for Video Encoded with Rate Control

We now briefly review the SMPT approach (referring the interested reader to [8] for a more detailed discussion) and describe how to transmit video encoded with rate control with the SMPT approach. The basic idea of SMPT is to transmit multiple LPDUs in parallel using multiple CDMA codes (one for each LPDU) when an LPDU was lost on the wireless link. Suppose an LPDU is not successfully received (and hence not acknowledged). In the most basic SMPT scheme, in the next slot the sending terminal transmits the lost packet *and* the subsequent LPDU (which would have been transmitted in that slot, had there not been a link error) on two CDMA codes. If these LPDUs are successfully received and acknowledged the sender returns to sending one packet using one CDMA code. Otherwise (i.e., if the packets are not successful) the terminal sends three LPDUs (the two unsuccessful LPDUs plus the LPDU next in line) using three CDMA codes. This process continues until the LPDUs are successful or the terminal has “ramped up” to using a pre-specified maximum number R of CDMA codes (where typically $R = 8$, due to the physical limitations of the radio front ends of practical wireless devices; for low-cost devices typically $R = 3$). We note that modern wireless systems, such as IS-95 (Rev. B) and UMTS, allow for the deployment of SMPT, as these systems provide the capability to adapt the transmission rates by varying the number of used codes. (We also remark that conceptually the transmission rate could be adapted by varying the spreading gain or a combination of varying the number of codes and the spreading gain. However, to fix ideas for our study, we consider a system that varies the number of codes and keeps the spreading gain fixed.)

To save precious wireless transmission resources (and energy) as well as to reduce the created interference, the “ramping” mechanism of SMPT can be combined with a link probing mechanism. The basic idea of link probing (first proposed in [25]), is to probe the link after an LPDU was not acknowledged. In our SMPT context, the sending terminal retransmits the lost LPDU (as a link probe) using only one single CDMA code until this probing LPDU is acknowledged. The terminal then starts to build the SMPT ramp to clear the backlog that has accumulated during the probing. With the *Slow-Healing* variation of SMPT the terminal ramps up by using two CDMA codes in the slot right after the probing

LPDU was received and acknowledged successfully, three codes in the subsequent slot, and so on, until all R codes are used. If at any point an LPDU is not acknowledged the terminal returns to probing and rebuilds the ramp after a successful probing LPDU. This probing refinement is especially effective with the highly-correlated, bursty errors of a typical wireless link (usually modeled as a two-state Markov Chain, see Section 2.4 and [19] for details) and is employed throughout the remainder of this study.

In this study on video transmission using SMPT we employ the outlined slow-healing SMPT approach for video encoded with rate control. The video frames are partitioned into LPDUs and buffered in a data link buffer of size L_{Queue} . This buffer is used to smooth out short-time scale variation of the rate-controlled video (and holds the LPDUs that are backed up during link probing). The motivation for using slow-healing SMPT for rate-controlled video is that rate-controlled video typically has only moderate variations on relatively short time-scales around a pre-specified target bit rate [6]. Slow-healing SMPT, which strives to stabilize the throughput of the wireless link is therefore well suited for this type of video traffic. An important advantage of slow-healing SMPT is that it generally has very good interference properties [4]; it has a small variance of the code usage in a wireless cell, reducing the demands on the power control, and has a small inter-cell interference to neighboring cells.

2.2 Fast/Slow-Start SMPT for Video Encoded Without Rate Control

For video encoded without rate control, slow-healing SMPT is not well suited. This is because video encoded without rate control exhibits a highly variable bit rate over a wide range of time scales (including long time scales). A moderately sized buffer at the link layer which is ideally drained at a constant rate is therefore not very efficient in serving this type of video traffic. Instead, the link layer needs to adapt to the varying bit rate of the video traffic. To achieve this we employ the more aggressive *Slow-Start* and *Fast-Start* SMPT mechanisms. With fast-start SMPT the wireless terminal transmits R LPDUs in parallel using R CDMA codes as long as there are backlogged LPDUs in the queue and all LPDUs are acknowledged. When an LPDU is not acknowledged the terminal begins to probe the channel with one LPDU per slot. Once the probing LPDU is acknowledged, the terminal resumes to transmit with R codes. With slow-start SMPT, on the other hand, the terminal starts building a ramp whenever more than one LPDU is in the queue (even when the backlog is not due to previous link errors). The terminal increases the number of used codes by one for each slot in which all LPDUs are acknowledged up to the maximum of R codes. When an LPDU is not acknowledged, the terminal starts probing the channel with one LPDU. Once the probing LPDU is acknowledged, the terminal resumes building the ramp (provided there are backlogged LPDUs in the queue). By striving to work off any backlog in the link layer buffer, the start SMPT mechanisms adapt to the varying bit rates. We note however, that these start mechanisms generally lead to a higher variability of the total number of used codes in a cell of a wireless cellular system, thus putting more burden on the cell's power control. Also, the interference seen in neighboring cells is larger, compared to the less aggressive slow-healing approach.

2.3 Performance Metrics

Continuous media applications (such as video) have stringent timing constraints. For the case of video, the receiver has to decode and display a new frame with every frame period (typically 40 msec or multiples thereof, see [6]). If a video frame is not completely received by its playout deadline, the receiver loses (a part or all of) the frame. The loss may result in jerky motions or artifacts in the video, which reduce the perceived video quality.

For the purpose of our study we express the timing constraints in the delay and jitter bounds standardized in RFC 1193 [3], which we briefly review here for convenience. Suppose at time t_0 a video

frame is generated and instantaneously passed through the transport and network layer down to the data link layer. Suppose the video frame (plus higher layer protocol headers) is partitioned into N LPDUs. With sequential transmission the minimal delay of the video frame is clearly

$$D_{\min} = N \cdot \tau_{slot}.$$

This minimal delay is attained when there are no errors on the wireless link, i.e., when no LPDU needs to be re-transmitted. However, typically there are some wireless link errors that result in some LPDU(s) being dropped, and subsequently re-transmitted. Suppose that the transport layer at the receiver accepts only segments that arrive with a delay no larger than

$$D_{\max} = \min\{\tau_{\text{delay}}, D_{\min} + \tau_{\text{jitter}}\}, \quad (1)$$

where τ_{delay} denotes the deterministic delay bound, and τ_{jitter} denotes the deterministic delay-jitter bound defined in [3]. We define D_{\max} as the *transmission window* within which a segment has to be transmitted (delivered) from the sender to the receiver in order to be considered successful. Suppose that the segment arrives at time t_2 at the transport layer at the receiver. The delay D of a successful segment satisfies

$$D = t_2 - t_0 \leq \tau_{\text{delay}}. \quad (2)$$

The jitter J of the segment is

$$J = t_2 - (t_0 + N \cdot \tau_{slot}) \leq \tau_{\text{jitter}}. \quad (3)$$

In our quantitative study we consider the following metrics:

- The *jitter* J as defined in (3). In our simulations we record the jitter of all successful segments (video frames) and report the resulting average jitter.
- The *probability of successful video frame (segment) delivery* for a given delay constraint τ_{delay} .
- The *goodput* (in bit/sec) is obtained by summing the sizes of the link layer payloads of all successfully transmitted LPDUs (in bit) and dividing this sum by the duration of a given video flow. (The link layer payload consists of the video frame plus UDP and IP headers plus padding, but does not include FEC and link layer header. Due to the UDP and IP protocol headers as well as the padding the goodput may be larger than the average bit rate of the video streams.)

Additionally, we consider the *capacity* which we define as the number of simultaneous video flows that can be supported in a given wireless cell while meeting specific delay constraints and goodput requirements.

2.4 Simulation Scenario

To evaluate the transmission of video using SMPT mechanisms we have developed comprehensive simulation programs based on `ptolemy` [17] and `ns-2` [16]. Because of space constraints we give here only a brief overview of the simulation programs and refer the interested reader to [8] for details. In our simulations we consider a scenario where multiple distributed wireless (and possibly mobile) terminals transmit video (exactly one stream per terminal) to a central base station. We chose this uplink transmission scenario within a wireless cell to fix ideas. Our SMPT mechanisms do not rely on the cellular structure; in fact with our SMPT mechanisms each wireless terminal schedules its LPDU transmissions without any knowledge of the other terminals' activities. (The wireless terminals only "feel" each other

through the shared interference environment (within the cell.) The simulation programs simulate the entire network protocol stack at each sending terminal and the corresponding protocol stacks at the base station (which serves as a receiver for the video streams).

For the simulation of the video traffic we use frame size traces of 25 videos encoded with rate control (for the scenario with rate control) and frame size traces of 25 videos encoded without rate control (for the scenario without rate control). The sets of frame size traces are described in more detail at the beginnings of Sections 3 and 4. In each scenario, for each of the ongoing video streams we randomly pick one out of 25 (respective) frame size traces as well as a random phase into the selected traces.

Each video frame is encapsulated into a single UDP transport layer segment and passed down through the IP network layer to the link layer. At the link layer the video frame (plus UDP and IP headers) is partitioned into LPDUs of 128 bytes each. (80 bytes of link layer payload (video frame, UDP and IP protocol headers, padding) + 47 bytes of FEC + 1 byte of link layer header.) The LPDUs are placed into the data link buffer of default size $L_{Queue} = 100$ LPDUs = 12.8 kByte.

At the physical layer, pseudo-noise spreading sequences are used. Each wireless links (consisting of up to R parallel code channels) is modeled using the two-state (“good” and “bad”) Markov Chain (Gilbert–Elliot) model with typical settings for the transition probabilities between the two states [19, 11, 8]. In the “bad” state all LPDUs sent over the link are dropped with probability one. In the “good” state an improved Gaussian approximation is used to evaluate the bit error probability on the link as a function of the total number of used pseudo-noise codes (i.e., interference level) in the cell. The LPDU drop probability is then calculated from the bit error probability assuming BCH(1023, 640, 41) forward error correction for each LPDU. The slot length is fixed at $\tau_{slot} = 10$ msec. The bit rate for one CDMA channel (in the wireless terminal to base station direction) is 64 kbps.

All simulations are run until the 99% confidence interval of the metrics of interest is less than 1% of the corresponding sample mean.

3 Simulation Results for Video Encoded With Rate Control

For the simulation of the transmission of video encoded with rate control we use the 25 frame size traces of video encoded in H.263 with a target bit rate of 64 kbps available from [6]. These H.263 encodings are characteristic of video encoded with rate control. The video traffic has only small to moderate variations around the target bit rate.

The spreading gain is set to 16 throughout this section.

3.1 Impact of the Number of Wireless Terminals

In Figure 1 we plot the goodput as a function of the number of wireless terminals (which is equivalent to the number of ongoing video streams). We give the goodput for the sequential transmission mode and the slow-healing SMPT approach (with a maximum number of $R = 3$ parallel CDMA codes for a given terminal). The corresponding jitter results are presented in Figure 2. In this experiment we do not impose any delay bound τ_{delay} or jitter bound τ_{jitter} . Thus at the receiver, there are no video frames discarded due to violated playback deadlines. Loss occurs only when the LPDUs carrying parts of a video frame find the link layer buffer full. We observe from the figures that with SMPT up to nine video streams can be supported with a goodput of 66 kbit/s and an average jitter J smaller than 10 ms. The sequential transmission mode achieves a goodput of 63 kbit/s for up to twelve video streams. However, the average jitter of the sequential transmission mode is roughly one second throughout. This is unacceptable for real-time video transmission, especially for interactive applications, such as video

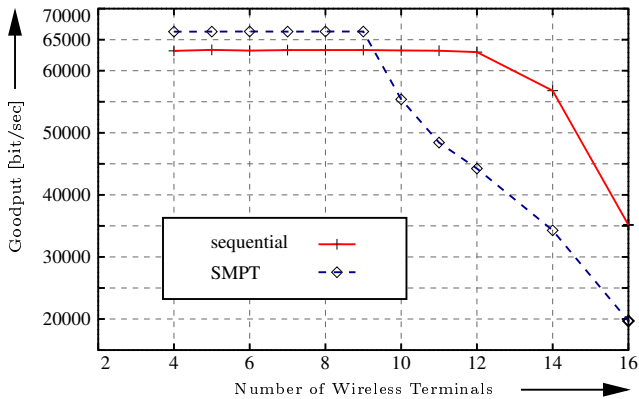


Figure 1: Goodput as a function of number of wireless terminals ($L_{Queue} = 100$ LPDUs).

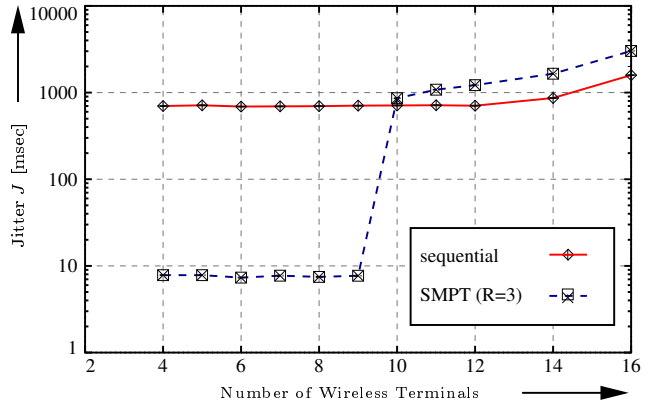


Figure 2: Average jitter J as a function of number of wireless terminals ($L_{Queue} = 100$ LPDUs).

conferences. We define the *operational phase* (capacity) as the range of video streams over which the QoS provided is stable. For the scenario shown in Figures 1 and 2 the operational phase for the sequential transmission mode is 12 video streams and for SMPT it is nine video streams. We observe that the jitter increases steeply when the capacity of the SMPT system is exceeded, i.e., when the number of video streams increases from nine to ten (an 11% increase in the cell load). To explain this, note that SMPT stabilizes the throughput by using more codes when LPDUs get backlogged in the link layer buffers of the wireless terminals. Once the traffic load is increased beyond the system capacity, the interference level increases significantly, leading to more dropped LPDUs and thus more backlog. In response, SMPT uses more codes, trying to clear the backlog. However, by using more codes the interference level is further increased. This in turn leads to more dropped LPDUs and more backlog, which SMPT is not able to clear when the system is loaded beyond its capacity. As we observe from Figure 2, when the SMPT is loaded beyond its capacity it gives about the same jitter performance as sequential transmission.

We thus observe that there exists a trade-off between the range of the operational phase (capacity) and QoS provided in terms of goodput and jitter. For a smaller operational phase the QoS provided by SMPT is much better than the QoS provided by sequential transmission. (We note that the average jitter studied here is only a first, coarse assessment of the performance. Even with an average jitter below a certain threshold τ_{jitter} , video frames may miss their deadline if the variability of the jitter is large. We study therefore the performance with respect to a fixed delay constraint in the next section and the jitter distribution in detail in Section 3.3.)

3.2 Impact of Delay Constraint τ_{delay}

In Figures 3 and 4 we plot the probability of successfully delivering a video frame as a function of the number of ongoing video streams for the sequential transmission mode and slow-healing SMPT. We give the probability of successful video frame delivery for the delay bounds $\tau_{delay} = 50, 100, 150, 200,$ and 250 msec. We note that the link layer at the sending wireless terminal is not aware of these delay bounds. The link layer simply tries to transmit the LPDUs in its buffer; it is not aware of the fact that the LPDUs carry video frames with playout deadlines. Thus, our approach preserves the isolation of the layers of the networking protocol stack and allows for the deployment of our mechanisms in low-cost wireless terminals. At the receiver's transport layer only the video frames meeting the delay bound are passed up to the application. We observe from Figure 3 that only a very small fraction (less than one

percent) of the video frames is transmitted successfully with the sequential transmission mode for the chosen delay bounds. As expected, larger delay constraints result in a larger probability of successful video frame delivery. However, even for a delay bound of $\tau_{\text{delay}} = 250$ msec, less than one percent of video frames are transmitted successfully. (We note that this very poor performance is due to the relatively large default size of the link layer buffer of $L_{\text{Queue}} = 100$ LPDUs. The link layer is not aware of the video frame deadlines and transmits all the LPDUs in the buffer, even though they may carry video frames that have already missed their deadline. The larger the buffer the more delay the LPDUs may experience in the buffer. We study the impact of the buffer size in detail in Section 3.3.)

We observe from Figure 4 that with the slow-healing SMPT mechanism the probability of sending a video frame successfully is very high. For nine and less wireless terminals (each sending one video stream) the success probability for a delay constraint of $\tau_{\text{delay}} = 250$ msec is 98%. Smaller delay constraints τ_{delay} decrease the success probability, but even for $\tau_{\text{delay}} = 150$ msec, the success probability is still around 60%. We conclude that for its operational phase up to nine ongoing video streams, the slow-healing SMPT mechanism achieves high probabilities of successful video frame transmission. With ten wireless terminals the SMPT system is loaded beyond its capacity and we observe a sharp drop-off in the success probability.

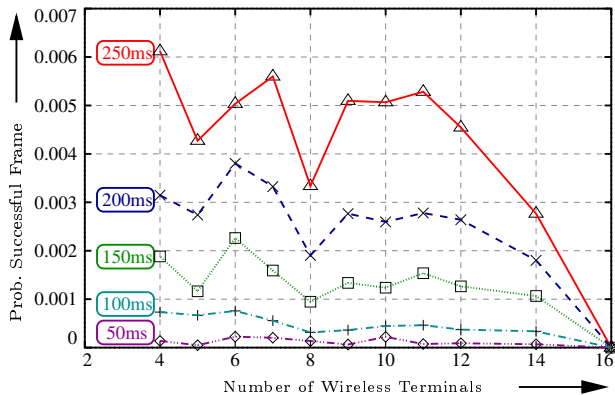


Figure 3: Probability of successfully delivering a video frame for delay constraints $\tau_{\text{delay}} = 50, 100, 150, 200,$ and 250 msec for the sequential transmission mode ($L_{\text{Queue}} = 100$ LPDUs, fixed).

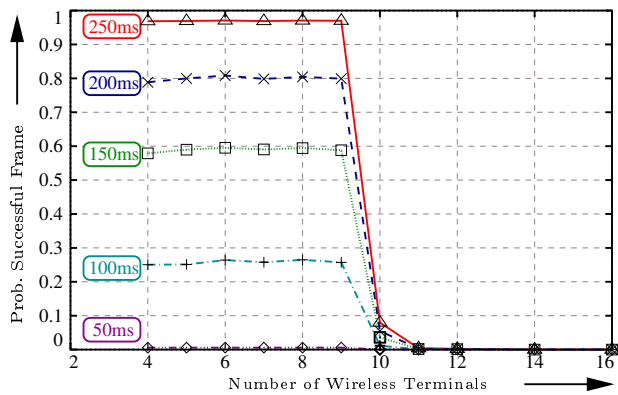


Figure 4: Probability of successfully delivering a video frame for delay constraints $\tau_{\text{delay}} = 50, 100, 150, 200,$ and 250 msec for the slow-healing SMPT mechanism ($L_{\text{Queue}} = 100$ LPDUs, fixed).

3.3 Impact of Link Layer Buffer Size L_{Queue}

For the simulation results reported so far, the buffer at the link layer was set to $L_{\text{Queue}} = 100$ LPDUs = 12.8 kByte. We now vary the size of the link layer buffer L_{Queue} . It is well known that smaller buffers reduce the jitter. However, smaller buffers also result in larger loss (and hence a smaller probability of successful video frame transmission), and consequently in a smaller goodput. We now investigate this trade-off quantitatively. Figures 5 and 6 give the goodput and the jitter as functions of the link layer buffer size L_{Queue} . (We do not impose a delay bound τ_{delay} or jitter bound τ_{jitter} in this experiment.) We investigate the situation where nine and ten wireless terminal are transmitting rate-controlled video simultaneously using either the sequential transmission mode or slow-healing SMPT. We observe that for the sequential transmission mode (i) the cell load (either nine or ten WTs) has no significant impact, and (ii) as expected, both the jitter and the goodput decrease as the buffer size is decreased from the




Figure 5: Goodput as a function of the link layer buffer size L_{Queue} .

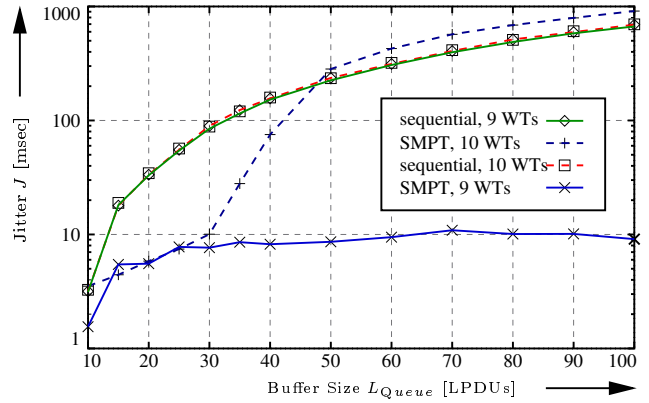


Figure 6: Average jitter J as a function of the link layer buffer size L_{Queue} .

default value of 100 LPDUs to smaller values. For a buffer size of $L_{Queue} = 30$ LPDUs, for instance, the average jitter is $J = 100$ msec and the goodput is 56 kbit/s. The implication of this experiment is that the performance of the sequential transmission mode is very sensitive to the link layer buffer size. For a relatively small buffer of 30 LPDUs the sequential transmission mode reaches its maximum goodput. Further increases in the buffer do not affect the goodput, but increase the jitter dramatically (and thus lead to a small probability of successful video frame delivery, as seen in the previous section).

For buffer sizes between 5 and 30 LPDUs there is no significant difference between the performance of SMPT for nine or ten wireless terminals. Note that the operational phase for SMPT includes nine wireless terminals with a buffer size of 100 LPDUs. For nine terminals the average jitter J is always below 10 ms and the goodput reaches 66 kbit/s for large buffers. For ten wireless terminals supported with SMPT the jitter is below 10 msec for small buffers (30 LPDUs or less); for larger buffers the jitter increases dramatically. The good news from this experiment is that within its operational range of up to nine ongoing video streams, SMPT is relatively insensitive to the buffer size (as long as the buffer has a certain minimum size, of 40 LPDUs in our setting); this simplifies the configuration of the mechanism in practice.

As noted above, real-time video transmission requires that the video frames are delivered within a tight delay bound. To achieve a tight delay bound the buffers should be small. In Figures 7, 8, 9, and 10 we plot the probability masses of the jitter J (in msec) as a function of the video frame size (in byte). We consider slow-healing SMPT and the sequential transmission mode in this experiment. The link layer buffer is set to $L_{Queue} = 20$ LPDUs and $L_{Queue} = 40$ LPDUs, respectively. There are nine wireless terminals — sending one video stream each — in the cell. Only the jitter values of successful video frames (i.e., video frames which had none of their LPDUs dropped due to a full link layer buffer) are considered. All figures have the bi-modal frame distribution — which is typical for video encoded with H.263 with rate control (see [6] for details) — in common. We observe that SMPT gives significantly smaller jitter J than the sequential transmission mode. For SMPT and a link layer buffer of $L_{Queue} = 20$ LPDUs, almost all of the probability mass is located at jitter values less than 30 msec; for a link layer buffer of $L_{Queue} = 40$ LPDUs, most of the probability mass is located at jitter values less than 75 msec. On the other hand, with the sequential transmission mode, video frames are very likely to experience a jitter of up to 150 msec for a link layer buffer of $L_{Queue} = 20$ LPDUs; with a link layer buffer of $L_{Queue} = 40$ LPDUs, jitter values in the range between 200 and 400 msec are very likely (especially for small video frames).




Figure 7: Probability masses for video frame jitter J as a function of the video frame size. (Sequential transmission, $L_{Queue} = 20$ LPDUs, 9 video streams.)




Figure 8: Probability masses for video frame jitter J as a function of the video frame size. (slow-healing SMPT, $L_{Queue} = 20$ LPDUs, 9 video streams.)

Upon closer inspection, we observe from Figures 7 and 9 that with the sequential transmission mode, smaller video frames are more likely to experience large jitter values, compared to large video frames. The reason for this is that large video frames fit only in an almost empty link layer queue and therefore typically are not delayed by retransmissions of preceding video frames. Small video frames, on the other hand, fit almost always into the queue and therefore are more likely to be delayed by retransmissions of preceding frames. Note also that frame loss is more likely for larger video frames than for smaller video frames for the sequential transmission mode. From Figures 8 and 10 we observe that for SMPT the smaller video frames are only slightly more likely to experience larger jitter values, compared to the larger video frames. In brief, this is due to the dynamics of the ramping process (see [8] for details).

Besides the simulations with the frame size traces we have also conducted experiments with the actual video streams to subjectively evaluate the video quality achieved by the studied transmission schemes. Because of space constraints we give here only a few illustrative results and refer the interested reader to [8] for more results. Figures 11 and 12 give snapshots for the video *Aladdin*¹ after transmission using the sequential transmission mode and slow-healing SMPT. The picture obtained with an error-free transmission is given for comparison. In the considered scenario there are nine ongoing video streams, the link layer buffer is $L_{Queue} = 40$ LPDUs, and the delay bound is $d_{delay} = 100$ msec. We observe that for both pictures the SMPT approach gives significantly better quality than the sequential transmission approach. In fact the SMPT picture is almost as good as the picture that would be obtained with an error-free transmission; it is only noticeable that in the SMPT picture the contours are somewhat “washed out” and not as sharp as in the error-free picture. The picture obtained with sequential transmission, on the other hand, is severely degraded in quality; it has several obvious artifacts (also, the left two-thirds of the picture in Figure 12 lag behind the error-free video picture). The degradations in the pictures are caused by video frames that missed their deadline and could therefore not be decoded. With each missed video frame the decoder misses the update of a number of macro-blocks which then

¹©DISNEY

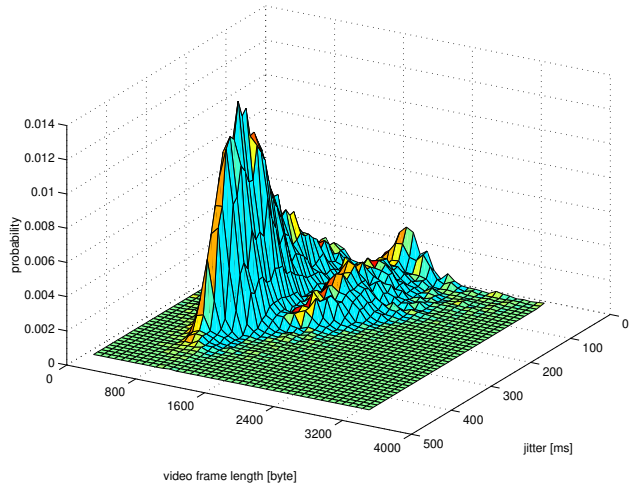


Figure 9: Probability masses for video frame jitter J as a function of the video frame size. (Sequential transmission, $L_{Queue} = 40$ LPDUs, 9 video streams.)




Figure 10: Probability masses for video frame jitter J as a function of the video frame size. (slow-healing SMPT, $L_{Queue} = 40$ LPDUs, 9 video streams.)

results in artifacts and washed out contours in future frames that are predictive encoded with respect to the missing frame.

We note that these pictures are only indented to give a rough impression of how effective the simple SMPT technique is. We did not consider any refinements, such as error concealment techniques [21] which can decode partial video frames and thus improve the quality further at the expense of added complexity.



Figure 11: Comparison I of video sequence *Aladdin* (©DISNEY) after sequential (left), SMPT (middle), and error-free (right) transmission.




Figure 12: Comparison II of video sequence *Aladdin* (©DISNEY) after sequential (left), SMPT (middle), and the error-free (right) transmission.

4 Simulation Results for Video Encoded Without Rate Control

In this section we study the uncoordinated real-time transmission of video encoded without rate control (i.e., in an open loop). For these simulations we used the frame size traces of 25 videos encoded in MPEG-4 with the fixed quantization parameters 10 for I-frames, 14 for P-frames, and 18 for B-frames (available from [6]). This open-loop encoding avoids the complexity introduced by rate control. (Also, it results in a constant video quality at the encoder output, whereas the video quality at the encoder output is slightly variable when rate control is employed). However, the traffic produced by open-loop encoding is highly variable. Not only are the individual video streams highly variable in the sizes of their video frames, but also the different video streams differ significantly in their frame size statistics, e.g., the video streams vary in their mean bit rate (see [6] for the exact statistical properties of the 25 video traces used). These variabilities pose a particular challenge for the network transport. We demonstrate that the simple to deploy slow- and fast-healing SMPT mechanisms are able to transport this highly variable traffic efficiently in real-time over the wireless links.

Throughout this section the spreading gain is set to a default value of 32 (whereas it was 16 throughout the preceding section). The reason for this larger spreading gain setting is that CDMA systems generally achieve better statistical multiplexing for larger spreading gains (especially for video traffic, see [8] for a detailed study). In the previous section the goal of the slow-healing SMPT transmission mechanism was to stabilize the throughput of the wireless link to the target bit rate of the rate-controlled video encodings; which have only small variations and hence require only a small amount of statistical multiplexing. For the highly variable open-loop encoded video considered in this section, a significant amount of statistical multiplexing is required for efficient transport.

4.1 Impact of the Number of Wireless Terminals

We first investigate the impact of the number of wireless terminals in the cell on the goodput and the average jitter J . We consider three transmission approaches: (1) sequential with a bit rate of 64 kbps (i.e., a spreading gain of 32) and a bit rate of 128 kbps (i.e., a spreading gain of 16), (2) slow-start SMPT (with up to $R = 8$ parallel code channels), and (3) fast-start SMPT (with up to $R = 8$ parallel code channels). In Figure 13 we plot the goodput as a function of the number of wireless terminals for the different transmission approaches. We observe that the sequential transmission mode with 64 kbps

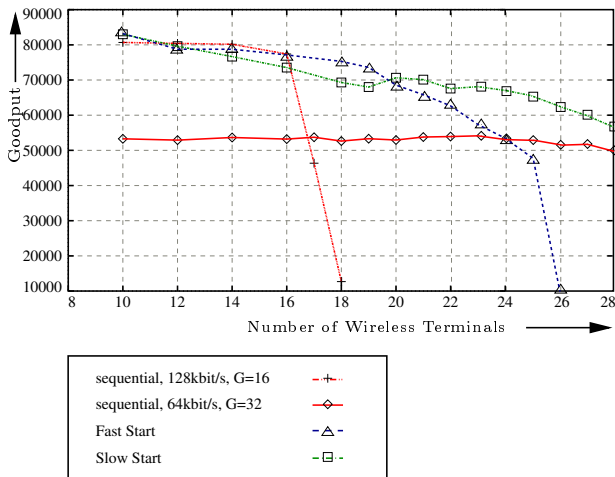


Figure 13: Goodput as a function of the number of wireless terminals ($L_{Queue} = 100$ LPDUs).

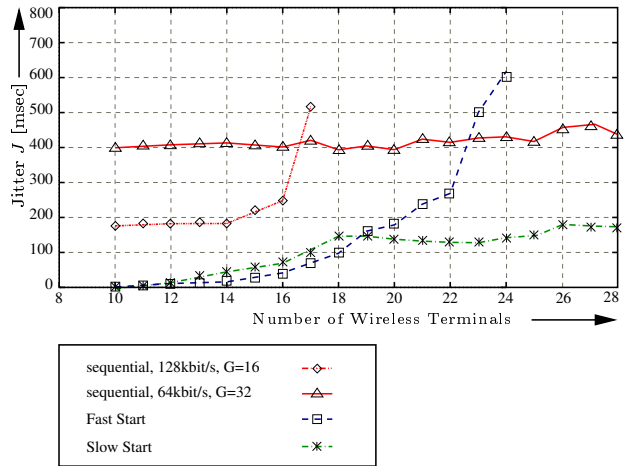


Figure 14: Average jitter J as a function of the number of wireless terminals ($L_{Queue} = 100$ LPDUs).

has a stable goodput of 53 kbps over the entire considered range of the number of wireless terminals (each transmitting one video stream). The sequential transmission mode with 128 kbps, which effectively transmits always on two CDMA code channels and produces therefore significantly larger interference, achieves a higher goodput of approximately 80 kbps up to 16 ongoing video streams. Up to 20 ongoing video streams the fast-start SMPT approach achieves the highest goodput. For more video streams the goodput drops slowly to the level of the sequential transmission mode with 64 kbps. The slow-start SMPT approach gives a slightly smaller goodput than the fast-start SMPT approach for less than 20 video streams, but for more video streams it achieves the largest goodput.

In Figure 14 we plot the average jitter J as a function of the number of wireless terminals. Sequential transmission with 64 kbps gives an average jitter J of 400 msec. This jitter is too large for real-time video transmission, considering that for video conferencing the jitter constraint is typically $\tau_{jitter} = 150$ msec. By doubling the bit rate a smaller jitter (which is still above the threshold for video conferencing) is achieved for up to 16 ongoing video streams. Only the SMPT mechanisms achieve jitter values that are acceptable for real-time communication. For up to 18 wireless terminals in the cell the fast-start SMPT mechanism gives a slightly smaller jitter than the slow-start SMPT mechanism. With a larger number of ongoing video streams in the cell, the fast-start SMPT mechanism becomes unstable. This is because the fast-start SMPT mechanism always uses all R CDMA code channels without taking the interference level in the cell (governed by the other terminals' activities) into consideration. The slow-start SMPT mechanism, on the other hand, gives an average jitter below 150 msec for up to 24 ongoing video streams. The slow-start SMPT mechanism gives stable performance as the cell load increases since it probes out the capacity in the cell by slowly increasing the number of used CDMA codes.

4.2 Impact of Delay Bound τ_{delay}

In Figures 15, 16, and 17 we plot the probability of successfully delivering a video frame as a function of the number of wireless terminals for the sequential transmission with 64 kbps, the fast-start SMPT mechanism, and the slow-start SMPT mechanism. We plot the probability of successful video frame delivery (i.e., the probability that a video frame is delivered to the receiver with a delay value smaller

than the delay bound τ_{delay}) for $\tau_{\text{delay}} = 50, 100, 150, 200,$ and 250 msec. We note that the delay bound is not known at the sender's link layer. The sender simply tries to send the video frames as fast as possible. We observe from Figure 15 that for the sequential transmission the probability of sending a video frame successfully is constant over the entire considered range of the number of wireless terminals in the cell. The success probability depends only on the delay constraint. For a delay constraint of $\tau_{\text{delay}} = 150$ msec, the probability of in-time delivery of a video frame is around 30%. We observe from Figure 16 that

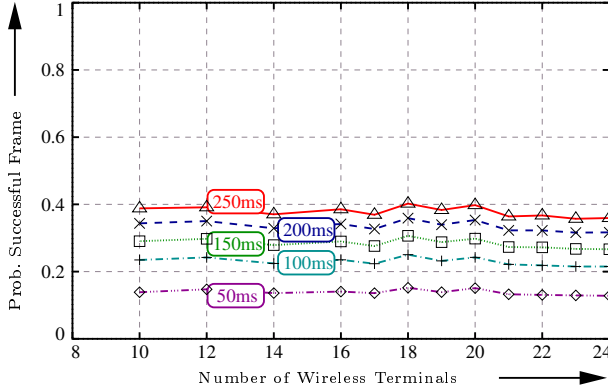


Figure 15: Probability of successfully delivering a video frame with delay constraints of $\tau_{\text{delay}} = 50, 100, 150, 200,$ and 250 msec for the sequential transmission mode ($L_{\text{Queue}} = 100$ LPDUs).

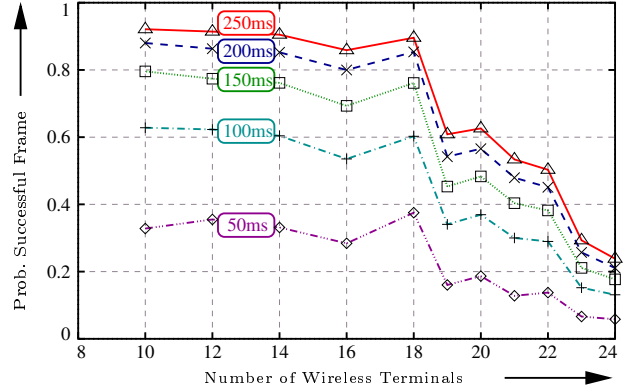


Figure 16: Probability of successfully delivering a video frame with delay constraints of $\tau_{\text{delay}} = 50, 100, 150, 200,$ and 250 msec for the fast-start SMPT mechanism ($L_{\text{Queue}} = 100$ LPDUs).

for the fast-start SMPT approach the probability of sending a video frame successfully is over 70% if the number of ongoing video streams is less than 19 and the delay constraint τ_{delay} is 150 msec or larger. We

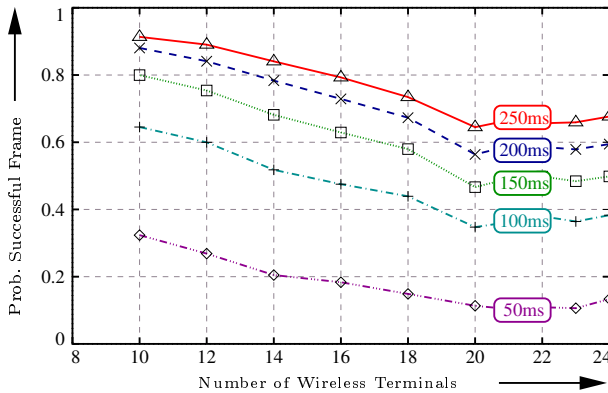


Figure 17: Probability of successfully delivering a video frame with delay constraints of $\tau_{\text{delay}} = 50, 100, 150, 200,$ and 250 msec for the slow-start SMPT mechanism ($L_{\text{Queue}} = 100$ LPDUs).

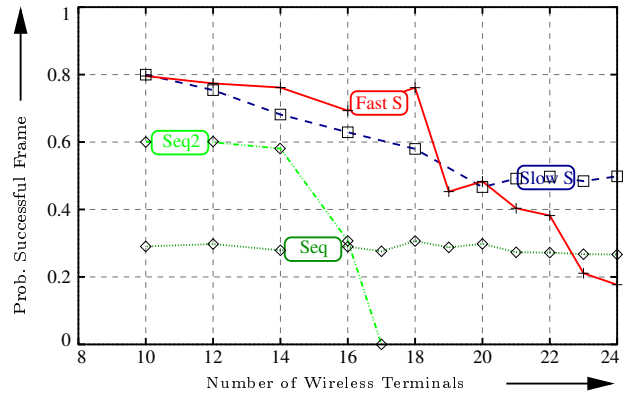


Figure 18: Comparison of different transmission schemes: Probability of successfully delivering a video frame for a delay constraint of $\tau_{\text{delay}} = 150$ msec ($L_{\text{Queue}} = 100$ LPDUs).

observe from Figure 17 that for slow-start SMPT (in contrast to sequential transmission) the probability of successful transmission of a video frame decreases with each additional video stream. However, the success probability with slow-start SMPT is always larger than with sequential transmission.

The probabilities of successful video frame delivery with the different transmission schemes (1) se-

quential transmission with 64 kbps (**Seq**), (2) sequential transmission with 128 kbps (bit rate doubled by using half the spreading gain) (**Seq²**), (3) fast-start SMPT (**FastS**), and (4) slow-start SMPT (**SlowS**) are compared in Figure 18 for the delay constraint $\tau_{delay} = 150$ msec. While fast-start SMPT gives the largest success probability for up to 18 wireless terminals, slow-start SMPT gives only slightly smaller success probabilities in this region. In contrast to the fast-start SMPT mechanism, however, the slow-start SMPT mechanism does not become unstable as the number of wireless terminals increases further. Note that slow-start SMPT performs always significantly better than sequential transmission.

4.3 Impact of Link Layer Buffer Size L_{Queue}

In Figure 19 we plot the average jitter J as a function of the link layer buffer size L_{Queue} for 18 wireless terminals and different transmission schemes. Figure 20 gives the corresponding video frame loss probabilities. (Recall that a video frame is lost only when one of the LPDUs carrying the frame finds the link layer buffer full; there is no delay bound τ_{delay} or jitter bound τ_{jitter} imposed in this experiment.) The buffer length is given in LPDUs, where each LPDU is 128 byte (and carries 80 byte of link layer payload). We observe that the SMPT mechanisms achieve significantly smaller average jitter values and smaller loss probabilities than the sequential transmission mode. Among the SMPT mechanisms, fast-start SMPT performs slightly better than slow-start SMPT.

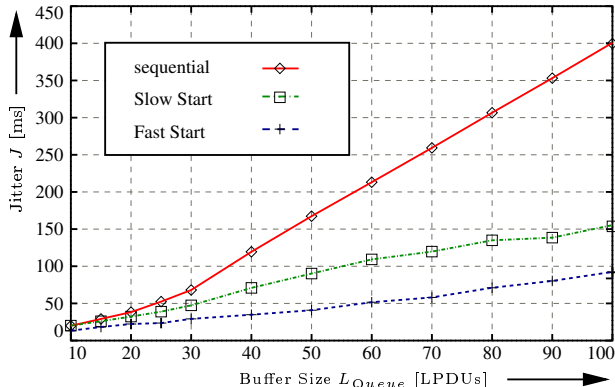


Figure 19: Average jitter J as a function of the link layer buffer size L_{Queue} for 18 wireless terminals and different transmission schemes.

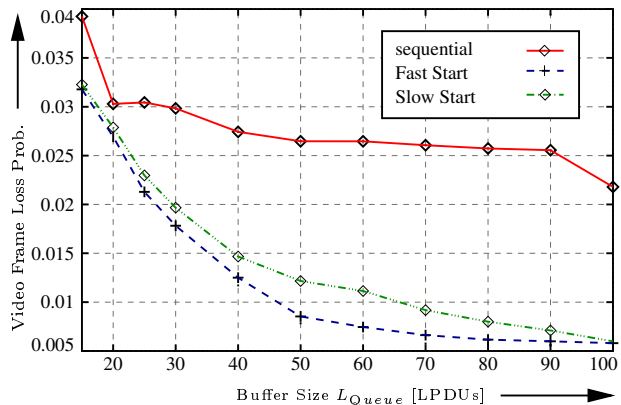


Figure 20: Video frame loss probability as a function of the link layer buffer size L_{Queue} for 18 wireless terminals and different transmission schemes.

While for small buffers the jitter and loss probabilities do not differ significantly among the transmission approaches, for larger buffers the differences are very pronounced. Note that larger buffers result in high cost end-systems. For the sequential transmission scheme a buffer size of $L_{Queue} = 20$ LPDUs (= 2.56 kByte) is a reasonable choice. For this buffer size the average jitter is small and the loss probability is about as small as it can be with sequential transmission. A larger buffer does not provide a significant improvement for sequential transmission. For the SMPT mechanisms, on the other hand, a buffer size between 50 and 70 LPDUs (6.4 – 9 kByte) gives both a small jitter and a small loss probability.

In Figures 21, 22, 23, and 24 we plot the probability masses of the jitter J (in msec) as a function of the video frame size (in byte). We consider the sequential transmission mode and fast-start SMPT in this experiment. The link layer buffer is set to $L_{Queue} = 20$ LPDUs and $L_{Queue} = 40$ LPDUs, respectively. There are 18 wireless terminals sending one video stream each in the cell. Note that fast-start SMPT may transmit a video frame faster than the sequential transmission would over an error-free channel.




Figure 21: Probability masses for video frame jitter J as a function of the video frame size. (Sequential transmission, $L_{Queue} = 20$ LPDUs, 18 video streams.)




Figure 22: Probability masses for video frame jitter J as a function of the video frame size. (fast-start SMPT, $L_{Queue} = 20$ LPDUs, 18 video streams.)

This results in negative jitter values J for the fast-start SMPT mechanism, which we count as a jitter of zero. We observe from the figures that fast-start SMPT achieves generally smaller jitter values than sequential transmission, with the differences being more pronounced for the larger link layer buffer. As we have observed for the rate-controlled video in Section 3.3, it is also the case for video without rate control that smaller video frames experience a larger jitter. This is again due to the fact that smaller video frames fit more easily into the link layer buffer and are more likely to be delayed by video frames ahead of them in the buffer. This effect is more pronounced for sequential transmission than for SMPT. For the large buffer, small frames experience jitter values of up to 500 msec with sequential transmission, whereas the jitter values are less than 200 msec with SMPT.

Overall, we conclude from the simulation results presented in this section that the fast- and slow-start SMPT mechanisms make the efficient transmission of highly variable open-loop encoded video over wireless links possible. The fast-start SMPT mechanism performs very well for a limited number of simultaneous video streams, but becomes unstable when the number of streams grows beyond a certain threshold (of 18 streams with our parameter setting). The slow-start SMPT mechanism performs almost as well as the fast-start SMPT mechanism when the number of video streams is small. In contrast to fast-start SMPT, slow-start SMPT does not become unstable; it achieves significantly higher goodput values as well as smaller jitter values and smaller loss probabilities than sequential transmission even for a large number of video streams. Thus, slow-start SMPT appears to be a good choice as link layer transmission mechanism for video encoded without rate control.

5 Conclusion

We have studied the uncoordinated real-time transmission of video by distributed wireless clients in a wireless cell. We have demonstrated that simple to deploy Simultaneous MAC Packet Transmission (SMPT) mechanisms enable the efficient transmission of video with tight real-time constraints on the




Figure 23: Probability masses for video frame jitter J as a function of the video frame size. (Sequential transmission, $L_{Queue} = 40$ LPDUs, 18 video streams.)

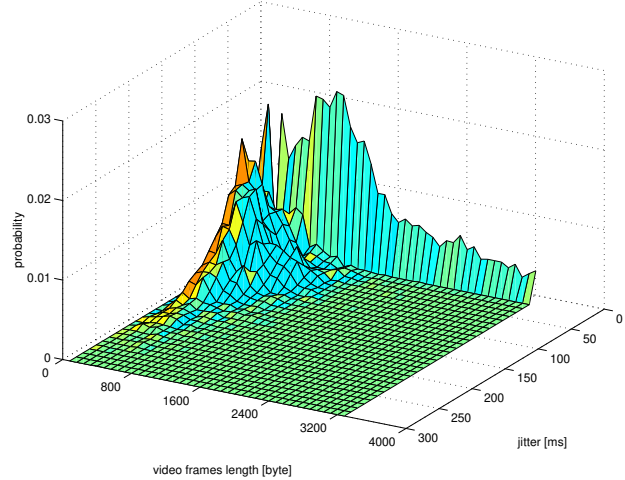


Figure 24: Probability masses for video frame jitter J as a function of the video frame size. (fast-start SMPT, $L_{Queue} = 40$ LPDUs, 18 video streams.)

order of 150 msec, thus enabling real-time applications, such as video conferencing, games, and telemedicine. For video encoded with rate control we found that the slow-healing SMPT mechanism achieves high goodput and small video frame loss probability while supporting a large number of simultaneous video streams in the cell (i.e., giving high cell capacity). For the more bursty video traffic resulting from encoding without rate control we found that the slow- and fast-start SMPT mechanisms provide efficient transmission scheduling to achieve a high cell capacity and good video quality. The slow-start SMPT mechanism is particularly resilient and degrades gracefully as the load on the cell increases. We studied the impact of the link layer buffer, the only hardware component required by the SMPT mechanisms. We gave guidelines for the dimensioning of this buffer.

While we considered the transmissions from distributed wireless terminals to a central base station in a cellular wireless network in our simulations, we emphasize that the presented SMPT mechanisms can be deployed readily in ad-hoc wireless networks. The SMPT mechanisms do not require any coordination of the transmissions among the distributed clients, and can thus be used for point-to-point transmissions in a cluster of a wireless ad-hoc network.

We also note that throughout our focus has been on mechanisms that are low in complexity and cost and easy to deploy yet give tangible performance gains, rather than finding more complex mechanisms that further enhance the performance. The described SMPT mechanisms work exclusively at the link layer of the sending wireless terminal and do not require any information from higher protocol layers. Thus preserving the isolation of the layers of the networking protocol stack and reducing cost and complexity. Given the simplicity of the described mechanisms there are several avenues for future research and refinement. For instance, a refined mechanism could take advantage of the deadlines of the video frames for the scheduling at the link layer. This refined scheduling algorithm would drop a frame that will miss its playback deadline at the receiver and instead start to transmit the next video frame. (Recall that the simple mechanisms studied in this paper do not assume any knowledge of the frames' deadlines and simply transmits all the LPDUs in the link layer buffer.) Note that the refinement would

add complexity since it needs to obtain the frame deadlines (for instance by parsing the RTP header). The refinement would improve the performance and the cluster (or cell) capacity by not transmitting video frames that would miss their playout deadline and thus reducing the interference level.

References

- [1] P.-R. Chang and C.-F. Lin. Wireless ATM-based multicode CDMA transport architecture for MPEG-2 video transmission. *Proceedings of the IEEE*, 87(10):1807–1824, October 1999.
- [2] G. Fankhauser, M. Dasen, N. Weiler, B. Plattner, and B. Stiller. The WaveVideo system and network architecture: Design and implementation. Technical report, Computer Engineering and Networks Laboratory (TIK), ETH Zurich, June 1998.
- [3] D. Ferrari. Client requirements for real-time communication services. RFC 1193, November 1990.
- [4] F. Fitzek, R. Morich, and A. Wolisz. Comparison of multi-code link-layer transmission strategies in 3G wireless CDMA. *IEEE Communication Magazine*, 38(10):58–64, October 2000.
- [5] F. Fitzek, B. Rathke, M. Schläger, and A. Wolisz. Simultaneous MAC-Packet Transmission in Integrated Broadband Mobile System for TCP. In *ACTS SUMMIT 1998*, pages 580–586. ACTS, June 1998.
- [6] F. Fitzek and M. Reisslein. MPEG-4 and H.263 video traces for network performance evaluation. *IEEE Network*, 15(6):40–54, November/December 2001. Video traces available at <http://www.eas.asu.edu/trace>.
- [7] F. Fitzek and M. Reisslein. A prefetching protocol for continuous media streaming in wireless environments. *IEEE Journal on Selected Areas in Communications*, 19(10):2015–2028, October 2001.
- [8] F. H. P. Fitzek. *QoS Support for Multimedia Services in Wireless CDMA Systems*. PhD thesis, Telecommunication Network Group, Dept. of Electrical Eng., Technical University Berlin, June 2002.
- [9] F. H. P. Fitzek, R. Supatrio, A. Wolisz, M. Krishnam, and M. Reisslein. Streaming video applications over TCP in CDMA based networks. In *Proceedings of the 2002 International Conference on Third Generation Wireless and Beyond (3Gwireless 2002)*, pages 755–760, San Francisco, CA, May 2002.
- [10] H. Gharavi and S. M. Alamouti. Multipriority video transmission for third-generation wireless communication systems. *Proceedings of the IEEE*, 87(10):1751–1763, October 1999.
- [11] C. Hsu, A. Ortega, and M. Khansari. Rate control for robust video transmission over burst-error wireless channels. *IEEE Journal on Selected Areas in Communications*, 17(5):756–773, May 1999.
- [12] Z. Jiang and L. Kleinrock. A packet selection algorithm for adaptive transmission of smoothed video over a wireless channel. *IEEE Transactions on Parallel and Distributed Systems*, 60(4):494–509, April 2000.
- [13] W. Kumwilaisak, J.W. Kim, and C.C.J. Kuo. Reliable wireless video transmission via fading channel estimation and adaptation. In *Proceedings of IEEE WCNC*, pages 185–190, Chicago, IL, September 2000.

- [14] H. Liu and M. El Zarki. Performance of H.263 video transmission over wireless networks using hybrid ARQ. *IEEE Journal on Selected Areas in Communications*, 15(9):1775–1786, December 1997.
- [15] M. Naghshineh and M.W. LeMair. End-to-end QoS provisioning in multimedia wireless/mobile networks using an adaptive framework. *IEEE Communications Magazine*, 35(11):72–81, November 1997.
- [16] ns2. The Network Simulator.
- [17] ptolemy. Ptolemy project — heterogeneous modeling and design, UC Berkeley, EECS.
- [18] D. Qiao and K. G. Shin. A two-step adaptive error recovery scheme for video transmission over wireless networks. In *Proceedings of IEEE Infocom*, Tel Aviv, Israel, March 2000.
- [19] R. R. Rao. Higher layer perspectives on modeling the wireless channel. In *Proceedings of IEEE ITW*, pages 137–138, Killarney, Ireland, June 1998.
- [20] A.S. Tosun and W.C. Feng. On improving quality of video for H.263 over wireless CDMA networks. In *Proceedings of IEEE WCNC*, pages 1421–1426, Chicago, IL, September 2000.
- [21] Y. Wang and Q. Zhu. Error control and concealment for video communication: A review. *Proceedings of the IEEE*, 86(5):974–997, May 1998.
- [22] F. Wegner. Personal Communication. Siemens AG, Mobile Radio Access Simulation Group, Berlin, Germany, May 2000.
- [23] D. Wu, Y. T. Hou, and Y.-Q. Zhang. Scalable video coding and transport over broad-band wireless networks. *Proceedings of the IEEE*, 89(1):6–20, January 2001.
- [24] H. Zheng and J. Boyce. An improved UDP protocol for video transmission over internet-to-wireless networks. *IEEE Transactions on Multimedia*, 3(3):356–365, September 2001.
- [25] M. Zorzi and R. R. Rao. Error control and energy consumption in communications for nomadic computing. *IEEE Transactions on Computers*, 46(3):279–289, March 1997.