

Journal of Zhejiang University SCIENCE A
 ISSN 1009-3095 (Print); ISSN 1862-1775 (Online)
 www.zju.edu.cn/jzus; www.springerlink.com
 E-mail: jzus@zju.edu.cn



Multi-user Motion JPEG2000 over wireless LAN: run-time performance-energy optimization with application-aware cross-layer scheduling*

Ji Xin^{†1,2,3}, POLLIN Sofie^{1,2,3}, LENOIR Gregory^{1,3},
 LAFRUIT Gauthier¹, DEJONGHE Antoine^{1,3}, CATTLOOR Francky^{1,2}

(¹Interuniversity Microelectronics Center, Kapeldreef 75, Leuven 3001, Belgium)

(²Department of Electrical Engineering, K.U.Leuven, ESAT/INSYS, Belgium)

(³Interdisciplinary Institute for Broadband Technology, Gent-Ledeberg B-9050, Belgium)

[†]E-mail: Xin.Ji@imec.be

Received Dec. 4, 2005; revision accepted Feb. 25, 2006

Abstract: This paper introduces a video application-aware cross-layer framework for joint performance-energy optimization, considering the scenario of multiple users upstreaming real-time Motion JPEG2000 video streams to the access point of a WiFi wireless local area network and extends the PHY-MAC run-time cross-layer scheduling strategy that we introduced in (Mangharam *et al.*, 2005; Pollin *et al.*, 2005) to also consider congested network situations where video packets have to be dropped. We show that an optimal solution at PHY-MAC level can be highly suboptimal at application level, and then show that making the cross-layer framework application-aware through a prioritized dropping policy capitalizing on the inherent scalability of Motion JPEG2000 video streams leads to drastic average video quality improvements and inter-user quality variation reductions of as much as 10 dB *PSNR*, without affecting the overall energy consumption requirements.

Key words: Performance-energy optimization, Application-aware scheduling, Motion JPEG2000, WLAN multi-user transmission
doi:10.1631/jzus.2006.AS0151 **Document code:** A **CLC number:** TN919.8

INTRODUCTION

The efficient transmission of video content over wireless communication networks is a challenging goal, especially when considering multiple mobile users equipped with handheld devices and sharing the same channel resources. First, the video application imposes stringent Quality-of-Service (QoS) requirements on the system. Aside from these QoS constraints, wireless transmission can result in highly error-prone and time-varying transmission conditions, which can have a dramatic impact on the video quality. Last but not least, the battery-powered user devices are heavily energy-constrained. Novel optimi-

zation strategies are needed in order to jointly meet these performance and energy challenges. A key solution to achieve this goal is to develop cross-layer systematic techniques that enable us to adapt the system configuration (for both the network and the different terminals) to the varying environment (e.g., instantaneous link reliabilities) and application requirements (e.g., instantaneous video rate demand) so as to provide QoS support to the application, while minimizing energy consumption in the system. In order to enable high quality multimedia applications, in particular for wireless video delivery, scalable video codecs offer a number of very important features, such as easy adaptability to bandwidth variations, robustness to data losses, support for rate scalability, scalable power requirements.

* Project supported by ISBO E2E QoE

The present work is aimed at introducing such a systematic technique enabling reliable and energy-efficient delivery of multiple scalable Motion JPEG2000 video streams over a wireless local area network (WLAN). The proposed solution jointly considers the physical (PHY) layer, the medium access control (MAC) layer and the application layer. It comprises a combination of (1) a cross-layer scheduler that adapts at run-time the parameters of the PHY and MAC layers to meet the video requirements, extending results in (Bougard *et al.*, 2006; Mangharam *et al.*, 2005; Pollin *et al.*, 2005), and (2) an application-aware prioritization strategy enabling us to cope with congestion situations in the network. This strategy exploits the intrinsic scalability of the Motion JPEG2000 bitstream, allocating each video packet to a specific priority queue at MAC-level.

We show that making our cross-layer framework application-aware results in significant end-to-end video quality improvements. It is demonstrated by simulation in the ns-2 network simulation framework augmented with realistic radio transceivers, network and video standards.

Conventional energy management techniques in the present context belong to two major categories: Sleeping (MAC centric), i.e., minimizing the fixed energy consumption of the transceiver circuit by transmitting at the highest rate, allowing the sleep mode (Ye *et al.*, 2002). And Scaling (PHY centric), i.e., scalable transmission control with variable transmission rate, coding and power, spreading the transmission on the complete transmission opportunity time, so as to minimize the transmission energy costs (Uysal-Biyikogly *et al.*, 2002). These two conflicting energy-management approaches present a tradeoff in minimizing the overall system energy and are jointly optimized, in our cross-layer framework.

Scalable video codecs can offer significant advantages in error-prone wireless network applications. The Motion JPEG2000 standard (ISO/IEC 15444-3, 2002), which is an extension of JPEG2000 (Taubman and Marcellin, 2002) for the coding of video sequences, provides a way to perform scalable video coding. A performance comparison between Motion JPEG2000 and the well-known MPEG-4 standard in the framework of video transmission over low bit-rate error-prone wireless channels had already been presented in (Dufaux and Ebrahimi, 2003), where the

authors show that in error-prone conditions intra-frame Motion JPEG2000 can outperform inter-frame MPEG-4. However, the authors only considered bit-level errors, which is not consistent with real packetized network transmission. Regarding the state of the art, an advanced approach of combining MAC level issues with application-awareness through scalable video streams was presented in (Li and van der Schaar, 2004). The authors introduced an adaptive retry-limit algorithm at MAC-level exploiting the features of scalable video streams in a WLAN. In our design, a much more elaborate and realistic system setup is taken into account: a full protocol stack and state-of-the art radio transceivers as opposed to raw transmission were set up; both packet error patterns caused by link failure and congestions are considered. The concern of energy-awareness is also a major differentiator.

The remainder of this paper is organized as follows. Section 2 introduces the considered scenario in more details. The run-time MAC-PHY cross-layer scheduler is introduced in Section 3, capitalizing on our previous work (Mangharam *et al.*, 2005; Pollin *et al.*, 2005). In Section 4, we extend the MAC-PHY framework so as to efficiently cope with congestion situations: a video application-aware dropping strategy relying on the bitstream scalability of Motion JPEG2000 is introduced. The experiments settings and simulation results are described in Section 5, showing a significant gain through the exploitation of the video scalability in our cross-layer framework. Section 6 eventually addresses the conclusion and sketches perspectives for future work.

SYSTEM DESCRIPTION

The considered setup consists of multiple independent users equipped with mobile terminals (MT) who want to upstream real-time video traffic to the access point (AP) of a wireless local area network (WLAN). These users transmit their data over a shared wireless channel, assumed to be slowly fading (typical of indoor propagation conditions). A WiFi-based WLAN is considered: the IEEE 802.11a standard (1999) is taken for the physical layer (OFDM-based transmission in the 5 GHz band), and the QoS-enabled IEEE 802.11e standard is considered for

the MAC functionalities. It is assumed that different video quality levels can be required by the different users. The system setup is depicted in Fig. 1.




Fig.1 The access point (AP) manages several mobile terminals (MT) in a centralized network

The scenario under test presents a number of challenges. First, video traffic is inherently delay-sensitive. As the wireless medium is highly error-prone, reliable transmission over the different links has to be guaranteed through the appropriate choice of the terminal parameters at physical layer-level. As the wireless channel is a broadcast medium, channel resources have also to be properly shared between the different users at MAC-level. On top of that, congestion in the network can result in packet losses, with a significant impact on video quality. Finally, as the mobile terminals are battery-powered, energy is heavily constrained. To meet all of these challenges optimally, our approach is two fold: First, we implement an energy-efficient PHY-MAC cross-layer run-time scheduler located at MAC-level of AP side. Assuming the resource requirements can be satisfied for all flows (i.e., no congestion in the system), the aforementioned scheduler enables guaranteeing the per-flow QoS constraints for multiple users while minimizing energy consumption; second, we utilize the new features of WLAN 802.11e protocol. This enables taking the application priority into consideration so as to efficiently transmit the video data in congested network situations.

RUN-TIME SCHEDULER

The goal of the run-time cross-layer scheduler is to ensure reliable and timely delivery of the different real-time video streams over the wireless links, while minimizing the total energy consumption of the mobile terminals under varying wireless channel conditions.

Conventional energy management techniques in the present context belong to two major categories:

(1) Sleeping (MAC centric), i.e., minimizing the fixed energy consumption of the transceiver circuit by transmitting at the highest rate, allowing the sleep mode of the transceiver. Irrespective of the channel utilization, the highest feasible PHY rate is then always used and the power amplifier operates at the maximum transmit power.

(2) Scaling (PHY centric), i.e., scalable transmission control with variable transmission rate, coding and power, spreading the transmission on the complete transmission opportunity time, so as to minimize the transmission energy costs. These two conflicting energy-management approaches present a trade-off in minimizing the overall system energy and are to be jointly optimized in our cross-layer framework.

In this context, a cross-layer run-time scheduler was developed (Mangharam *et al.*, 2005; Pollin *et al.*, 2005). This scheduler is located at the AP and relies on the HCF functionality of the IEEE 802.11e MAC protocol. Its goal is to optimize the performance (expressed as a job failure rate criterion (Mangharam *et al.*, 2005), where a job is defined as the reliable delivery of the information that has to be transmitted by a terminal during an MAC scheduling period), while minimizing the overall energy consumption in the mobile terminal. The design of the controller is sub-divided into two steps:

Step 1: During the so-called design-time phase, the performance-energy scalability available in the components of the system is first analyzed, based on available knobs (i.e., run-time controllable parameters). The considered knobs (stated here at functional level) are constellation order, code rate, output power of the front-end, linearity of the front-end, communication mode (SISO or MIMO). Performance and energy models of the system's behavior are built to capture the impact of these knobs at system-level. The performance-energy trade-off of each user can then be fully characterized for each possible system state (i.e., the finite set of possible realizations of the external variables).

Step 2: During the run-time phase, knowing the current system state and relying on the characterized trade-off obtained in the design-time phase, the AP can then jointly optimize the aforementioned shut-down and scaling MAC and PHY of the network strategies, define the amount of channel resource

allocated to each user, and the configuration they should use, and inform the terminals of the different schedule they are assigned.

The design step of the run time scheduler is out of the scope of this paper. We refer the interested reader to (Mangharam *et al.*, 2005; Pollin *et al.*, 2005) for more details.

INTRODUCING APPLICATION-AWARENESS BASED ON BITSTREAM SCALABILITY

In the present section, we extend the PHY-MAC run-time scheduler introduced in the previous section to consider also congested network situations.

After a brief description of the Motion JPEG2000 bitstream structure and packetization, we will introduce an application-aware dropping strategy that exploits the inherent scalability of the bitstream so as to cope with these congestion situations. This will enable us to show that making our cross-layer framework application-aware can result in significant improvements in end-to-end video quality compared to the optimal PHY-MAC approach.

Bitstream structure of Motion JPEG2000 and packetization for network transmission

The Motion JPEG2000 standard is an intra frame video coding standard based on wavelet transform. Relying on intra frame encoding, Motion JPEG2000 has less coding efficiency than an inter-frame encoder. On the other hand, because frames are independently coded, the spread of transmission errors is effectively prevented across consecutive frames. Moreover, re-synchronization markers can be inserted within each frame, which limits to a great extent the propagation of errors. Together with the various forms of scalability (e.g., resolution or quality) and precise rate control, Motion JPEG2000 offers attractive features for wireless transmission conditions.

The bitstream of Motion JPEG2000 is composed of video packets. Each video packet corresponds to a specific quality layer, resolution, component and precinct, and is the smallest unit of the bitstream (Taubman and Marcellin, 2002). According to different progress orders, the packets can be concatenated according to Quality, Resolution, Precinct and Components orders. In our case the encoding will be

done with the progress order of Quality, Resolution, Precinct and Components.

Regarding packetization, it is assumed that every packet belonging to one quality layer will be encapsulated into one single UDP packet, which will be transmitted independently. The main headers are encapsulated into the UDP packet associated with quality layer 1.

Application-aware dropping strategy relying on bitstream scalability

The scheduler introduced in the previous section is designed to reach a given transmission error rate under a given delay constraint, assuming no congestion is present in the network.

In case congestion occurs, the PHY+MAC optimal dropping policy is to drop the biggest rate demand in the system first. As it will be shown in the next section, this approach is highly inefficient from the video transmission perspective. Introduction of video-awareness in our cross-layer strategy, based on the features of the Motion JPEG2000 bitstream, enables significant performance improvement. In order to introduce this application-awareness, in our design, we rely on the multiple traffic queues offered at MAC-level by the IEEE 802.11e protocol (QoS extension to WiFi WLANs) (IEEE Std 802.11e/D13.0, 2005).

In Fig.2a, we present the PHY+MAC optimal dropping policy. Two cases are distinguished: one queue per user and multiple queues per user. As mentioned above, when the aggregate rate demand exceeds the available channel resource (i.e., when data cannot be scheduled), the optimal PHY-MAC dropping policy consists (in both cases) of dropping the queue in the system that has the largest rate demand. This process is repeated until the remaining aggregate data rate can be scheduled. The fact that the dropping is performed per queue is motivated by the complexity and tractability considerations. In the 802.11e MAC protocol, the whole queue size is reported to the AP by the terminal at the beginning of each schedule interval. If there is congestion (i.e., if there is not enough network resource), the scheduler needs to recalculate the optimized schedule for each terminal. Doing so at packet-level would significantly increase the complexity. Furthermore, since we design the scheduling period to be approximately similar

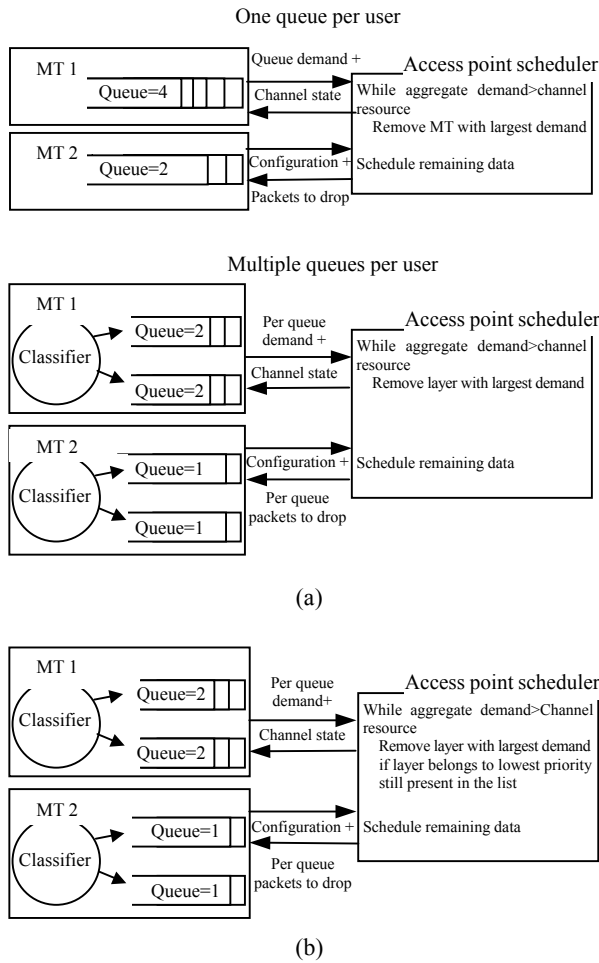


Fig.2 (a) PHY+MAC optimal dropping policy; (b) Application-aware cross-layer dropping policy

to video frame duration, the accumulated packets inside one queue per scheduling period are limited, so that the impact of working at queue-level on the ultimate visual quality should not be dramatic. In the multiple-queue case, the some application awareness is already present in the classifier. This feature is however not taken into account in the dropping policy, which remains PHY+MAC-centric. This setup will be considered for the reference PHY+MAC-centric approach in the sequel (the case of one queue per user would obviously lead to much worse performance, and would not be a fair comparison point).

In Fig.2b, we present the application-aware dropping policy. Prioritized treatment is obtained by associating each of these queues with a given quality layer and by assigning dropping priorities according to the quality layer numbers of the bitstream. When

congestion happens, the dropping policy consists of dropping the queue which has the largest demand among the queues in the system corresponding to the lowest quality layer. If all the queues corresponding to the lowest quality layer have been dropped, the queues corresponding to the second lowest quality layer are then considered. This process is repeated and results in an application-aware cross-layer approach. This simple strategy, as will be shown later, enables us to drastically increase the robustness of the system: it enables a significantly higher number of users for the same quality and energy consumption.

SIMULATION RESULTS

Experiment settings

1. Network part

The transmission of the Motion JPEG2000 video streams over the considered 802.11e wireless error-prone network is simulated by means of an extended ns-2 network simulator, adding performance and energy models of radio transceivers closely matching state-of-the-art implementations. The IEEE 802.11a standard is considered for the physical layer (OFDM-based transmission in the 5 GHz band), and the QoS-enabled IEEE 802.11e standard is considered for the MAC functionalities. The wireless channel is represented by a discrete-state Markov model, closely mapping the associated dynamics, assuming indoor propagation conditions. We refer the interested reader to (Bougard *et al.*, 2006) for a more detailed description of these models. The user number in the following simulation results is increased until 8.

2. Multimedia part

The Motion JPEG2000 Codec used in the simulation relies on Kakadu software (<http://www.kakadusoftware.com>). To have an overview of various kinds of motion degree and bitrate requirements, the tested videos considered in our experiments are: Foreman, Mobile, Bus, Football, with CIF resolution and 30 frames per second. The parameters explored at the encoding side are the number of quality layers and the encoding quality. Each sequence was encoded to force the decoded sequence reach an average *PSNR* without transmission errors around 30 dB and 35 dB, resulting in the bitrate settings reported in Table 1. For a fixed bitrate, the bitstream is encoded respec-

tively in 2, 4 and 8 quality layers, with the bitrate of lowest layer to be around 0.2 Mbps and the intervening layers to be assigned roughly logarithmically spaced bit-rates.

Table 1 Multimedia bitrate settings

Sequence	Bitrate (Mbps)	
	<i>PSNR</i> =30 dB	<i>PSNR</i> =35 dB
Foreman	0.61	2.13
Mobile	5.17	10.64
Bus	2.13	4.56
Football	0.67	2.13

Result analysis

The video sequences have been transmitted 3 times to get a relevant evaluation of the system statistics. The simulation results are shown from two different aspects, which are evaluated as a function of the number of users in the system. The first aspect is the mean *PSNR* value of different users as quality metric (the span of the *PSNR* across different users has also been displayed as an additional metric of quality to show the consistency across different users). The second aspect is the energy consumption in the system (in Joule); both the AP energy and the average user transmission energy are represented. From the results, we find that PHY+MAC optimization with scalable video application-aware dropping policy can improve the end-to-end performance drastically, without adversely affecting the overall energy consumption requirements.

Fig.3 compares the performance of the application-aware cross-layer approach to that of the MAC-PHY-centric cross-layer approach. It shows the average *PSNR* and *PSNR* span across different users within the network, considering the Football sequence encoded at 2.13 Mbps (hence an encoding *PSNR* of around 35 dB). The number of quality layers is set to 4, which according to the global results analysis provides the best quality-energy trade-off (see further). The application-aware cross-layer approach is clearly better, achieving much higher average *PSNR* and less *PSNR* span across different users. When the user number is increased to 8, with the application-aware cross layer approach (MAC+PHY+application), the mean *PSNR* value of the different users can still achieve 32 dB; on the contrary, the PHY+MAC-only approach decreases to about 22

dB for the average user, and the *PSNR* value difference between the best user and the worst user can be as large as 10 dB. Fig.4 shows the impact of the number of quality layers (2, 4 and 8). These comprehensive results strengthen the conclusion that, leveraging the bitstream scalability, the application-aware cross-layer approach guarantees constant high video quality to all users. These observations are confirmed by visual inspection of the decoded sequences, showing much less temporal jitter.

An interesting observation is that the average energy consumptions remain almost the same for both policies. This comparison is provided in Fig.5 showing the AP energy consumption and average transmission energy of each user when considering 4 quality layers. Fig.6 emphasizes the impact of the number of quality layers on these results. Also observe that as the number of users increases, the per-user transmitted energy decreases. This phenomenon is at first sight counter-intuitive, but it can easily be explained on the basis of Fig.3. When more users are added, more congestion occurs in the network: as a result of the dropping policy, the MTs are then required to drop packets up-front, i.e., before sending them to the AP. This on average frees up network resources, which advantageously reduces the energy consumption per user. This phenomenon could be used for the purpose of efficiently optimizing the energy cost by deliberately dropping some low-priority video packets independently of congestion consideration. Even if the energy costs of both approaches are similar, it is of course obvious, in view of the resulting *PSNR* impact depicted in Fig.3. It is obvious that the application-aware cross-layer approach is mandatory for best quality performance, as it offers an important added value. Fig.6 shows that the more the quality layers, the more the energy consumption. Since the difference in *PSNR* between 4 and 8 quality layers is negligible, we recommend for energy reasons to rely on medium scalable bitstreams (i.e., 4 quality layers) in the considered setup.

In Fig.7, the impact of the encoding *PSNR* is considered, considering several video sequences. Triangle markers are used for sequence Football, star markers for sequence Mobile and pentagrams for sequence Bus. Solid and dashed line styles are used to distinguish average encoding *PSNR*s (i.e., without transmission errors) around 35 dB and 30 dB,




Fig.3 Average PSNR and PSNR span vs user numbers

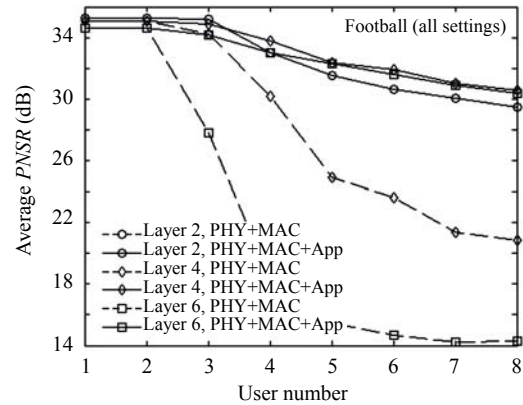


Fig.4 Average PSNR and PSNR span vs user numbers (all settings)




Fig.5 Network transmission energy vs user numbers

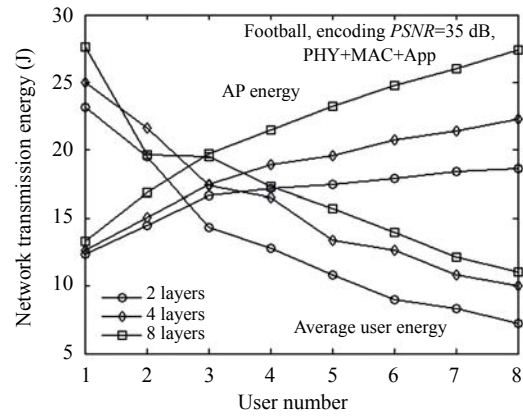


Fig.6 The influence of quality layer number to network transmission energy

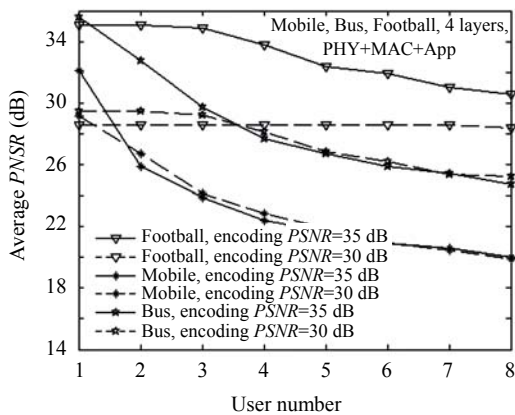


Fig.7 The influence of number of users on the transformation of encoding PSNR into decoded PSNR

respectively. These results show how almost independently of the encoding PSNR (30 dB vs 35 dB), the quality of the decoded video sequence in the pre-

sence of multiple users tend to converge to the same final quality. This corresponds to a situation where only one single quality layer (the highest priority one) is left and needs to be decoded. Depending on the actual video content, the convergence between the 30 dB and 35 dB distortion curves occurs at a different number of users, depending on the considered data: 2 users with sequence Mobile, 3 users with sequence Bus and more than 8 users with sequence Football. Encoding such sequences in scalable format could thus enable automatic adaptation of the quality with the network and users conditions.

CONCLUSION

To meet emerging wireless multimedia application requirements, energy-efficient support for end-

to-end video quality for multi-users in WLAN is a challenging goal. In this context, we have introduced a content-aware run-time cross-layer performance-energy optimization framework. It extends the PHY-MAC run-time cross-layer scheduling strategy that we introduced in (Mangharam *et al.*, 2005; Pollin *et al.*, 2005), to also enable considering congested network situations where video packets have to be dropped. We have shown that an optimal solution at PHY-MAC level can be highly suboptimal at application level. We have then shown that making the cross-layer framework application-aware, capitalizing on the inherent scalability of Motion JPEG2000 video steams, leads to drastic improvement of end-to-end video quality, compared to the optimal PHY-MAC approach: the average video quality improves with as much as 10 dB *PSNR*, while a drastic reduction of 10 dB *PSNR* occurs in inter-user quality variations.

References

- Bougard, B., Pollin, S., Dejonghe, A., Catthoor, F., Dehaene, W., 2006. Cross-layer power management in wireless networks and consequences on system-level architecture. *EURASIP Signal Processing Journal, Special Issue on Advances in Signal Processing*, in Press.
- Dufaux, F., Ebrahimi, T., 2003. Motion JPEG2000 for Wireless Applications. Proc. of First International JPEG2000 Workshop. Lugano, Switzerland.
- IEEE Std 802.11a, 1999. Standard for Telecommunications and Information Exchange Between Systems—LAN/MAN Specific Requirements, Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications.
- IEEE Std 802.11e/D13.0, 2005. Medium Access Control (MAC) Enhancements for Quality of Service (QoS). IEEE 802.11 WG, Draft Supplement.
- ISO/IEC 15444-3, 2002. Information Technology—JPEG2000 Image Coding System, Part 3: Motion JPEG2000.
- Li, Q., van der Schaar, M., 2004. Providing adaptive QoS to layered video over wireless local area networks through real-time retry limit adaptation. *IEEE Trans. on Multimedia*, 6(2):278-290.
- Mangharam, R., Pollin, S., Bougard, B., Rajkumar, R., Catthoor, F., van der Perre, L., Moerman, I., 2005. Optimal Fixed and Scalable Energy Management for Wireless Networks. IEEE Conference on Computer Communications (INFOCOM'05). Miami, FL, USA.
- Pollin, S., Bougard, B., Mangharam, R., Catthoor, F., Rajkumar, R., Moerman, I., van der Perre, L., 2005. Optimizing transmission and shutdown for energy-efficient real-time packet scheduling in clustered ad hoc networks. *EURASIP Journal on Wireless Communications and Networking, Special Issue on Cross-layer Design for Ad Hoc Networks: Cross-Layer Issues*, p.698-711.
- Taubman, D., Marcellin, M., 2002. JPEG2000: Image Compression Fundamentals, Standards and Practice. Kluwer Academic Publishers.
- Uysal-Biyikogly, E., Prabhakar, B., El Gamal, A., 2002. Energy-efficient packet transmission over a wireless link. *ACM/IEEE Transactions on Networking*, 10(4):487-499. [doi:10.1109/TNET.2002.801419]
- Ye, W., Heidemann, J., Estrin, D., 2002. An Energy-Efficient MAC Protocol for Wireless Sensor Networks. IEEE INFOCOM 2002.

Welcome visiting our journal website: <http://www.zju.edu.cn/jzus>
 Welcome contributions & subscription from all over the world
 The editor would welcome your view or comments on any item in the journal, or related matters
 Please write to: Helen Zhang, Managing Editor of JZUS
 E-mail: jzus@zju.edu.cn Tel/Fax: 86-571-87952276/87952331