

Co-optimization of Streaming Multimedia QOS and Backlight Power Consumption for Mobile Handheld Devices

Sudeep Pasricha, Shivajit Mohapatra, Manev Luthra, Nikil Dutt & Nalini Venkatasubramanian
School of Information & Computer Science
University of California, Irvine, CA 92697-3425
{sudeep,mopy,mluthra,dutt,nalini}@ics.uci.edu

Abstract

Mobile handheld devices have stringent constraints on power consumption because they run on batteries that have a limited lifetime. Conserving power to prolong battery life is of primary importance for these devices. Several factors such as backlight intensity, the hard disk, the CPU, the network interface and the nature of the application contribute significantly towards power consumption for a mobile device. While significant research effort has been made to optimize power consumption at the application, network and processor levels, comparatively little work has been done to reduce or adapt to the power consumed by the backlight. We propose an adaptive middleware based approach to optimize backlight power consumption for mobile handheld devices when playing streaming MPEG-1 video, without significantly compromising on video quality. Our performance results indicate that up to 60% of the power consumed by the backlight can be saved by using the proposed approach.

1. Introduction

Modern handheld mobile devices are increasingly being used to provide users with streaming multimedia content. These devices due to their modest sizes and weights have limited computing, storage and battery resources. Moreover multimedia applications tend to be extremely resource hungry, due to which streaming multimedia content onto low-power mobile handheld devices introduces a significant research challenge. Typically, the CPU, the network interface and the display constitute the three primary sources of power consumption in low power devices. In this work, we focus on achieving energy savings from the backlight display of the device, without significantly compromising on the quality of the streamed video.

Researchers have proposed several schemes in recent years, to optimize power consumption of low-power devices in mobile environments. Techniques such as dynamic voltage scaling (DVS) [2,3], dynamic power management of disks [10], network interfaces [8] and compiler/OS/middleware based adaptation [4,5,11] attempt to reduce power consumption at various computational levels. However, efforts to reduce the power consumption of the backlight

have received relatively little research interest. Choi et. al. [1] have done some interesting work in this area and attempted to address this problem. They propose compensating the brightness of a still image and simultaneously reducing the backlight level, such that there is no perceptible difference in luminous intensity, compared to viewing the unmodified image at the original backlight level. They also suggest that more aggressive compensation is possible if the contrast of the image is enhanced. However their proposed contrast compensation distorts the original image (since the original color of the image is not preserved), which limits the practical application of their scheme. Moreover, they implement their approach in the limited context of still images.

In this work, we explore a more aggressive approach to brightness compensation and device backlight control for streaming video. Furthermore, the adaptation is shifted away from the low-power device and performed at a network proxy server, obviating the need for the decoder on the device to be modified. We have found that aggressive brightness compensation is possible for streaming video as compared to still images, without considerably impacting the video quality. This is because small defects (introduced due to aggressive compensation) that might be noticeable in a still image are less discernable in streaming video where several frames (images) are displayed on the screen every second. We also propose an effective brightness compensation algorithm for optimized power savings. Finally, we introduce middleware based adaptation schemes which integrate our compensation algorithm to achieve low power backlight operation for streaming video content to mobile handheld devices. Our proposed approach gives significant power reductions, up to 60% of the power consumption attributed to the backlight, depending on the chosen adaptation scheme and the characteristics of the streamed video.

The rest of the paper is organized as follows. Section 2 describes the system architecture used. In Section 3 we discuss some issues relevant to backlight compensation. Section 4 presents the dual compensation algorithm in detail. Section 5 describes the middleware adaptation policies for low power backlight operation. In Section 6, we present our performance evaluation. Section 7 presents concluding remarks and future research directions.

2. System Architecture

We assume a system model depicted in Figure 1. The system entities include a multimedia server, a proxy server that has access to a database of profiled luminosity values for various video streams and device specific parameters (e.g. number of backlight levels, average luminosity at each level etc.), a rule base to determine compensation values and a video transcoder (Figure 1); and low-power wireless devices capable of displaying streaming MPEG video content. Moreover, all communication between the handhelds and the multimedia server are routed through the proxy server that can change the video stream in real-time.

Each client has an application layer where the video stream is decoded and a middleware layer which routes the information flowing from and to the video decoder application. The client middleware layer has access to system parameters such as the backlight levels, the current battery level and information identifying the type and make of the handheld (e.g. iPAQ, Jornada etc). In addition to accessing these system parameters, the middleware layer on the client can change these parameters (e.g. operating backlight level) through API calls to the underlying OS.

The middleware on the proxy performs the dynamic adaptation of the streaming video content (brightness compensation) and communicates control information to the client middleware (operating backlight levels) through the low bandwidth control stream. The proxy maintains a database of information about the videos available at the server and information specific to different handheld types such as the number, luminous intensity and average power consumption of the backlight levels. Additionally, the proxy also employs a static rule base which specifies conditions which determine values for backlight and video compensation. The database and certain parameters of the rule base are populated by extensive profiling and subjective assessment of videos on different handhelds.

3. Backlight Compensation

Table 1 gives the power readings for power consumed at different backlight levels on a Compaq iPAQ handheld device running Windows CE. The power readings have been averaged for several different MPEG-1 video streams. From the table, we see that significant energy savings are possible by operating the device at a lower backlight intensity level.

When operating in the Super Bright mode for instance, the backlight contributes almost 40% to the total system power consumption! The inherent problem with proposing schemes to reduce power at this level is that the backlight directly affects the display quality and the user experience. For example, even a slight reduction in backlight intensity during multimedia playback on the handheld (with the intention of saving power) can degrade the human

perception of multimedia quality. Indeed, simply reducing the backlight is not a viable solution. We explore the use of a video compensation algorithm that induces power savings without noticeably affecting video quality.




Figure 1: System Architecture

Our prior work [11] has focused on several aspects of video quality on handheld devices. We have determined that user perception of video playback quality is significantly influenced by the environment and the type of handheld used to view the video. Consequently, objective assessment [7] of video quality is extremely difficult and subjective assessment [6] is still the primary method for assessing video quality. In our profiling and surveys, we attempted to follow the recommendations given in [6] and chose a diverse collection of video streams (movie clips, animations, sports, documentaries etc) to use in assessing the suitability and effectiveness of our proposed schemes.

Backlight Modes	Power Consumed (in Watts)
Super Bright	2.80
High Bright	2.51
Medium Bright	2.32
Low Bright	2.16
Power Save	1.72

Table 1: Power consumed at various backlight levels during streaming multimedia playback

To validate our assessments, we conducted an extensive survey to subjectively assess the human perception of video quality when viewing streaming video on a handheld device. The subjects were first shown a full screen version of an original unaltered video stream. Next, they were shown the compensated stream and asked to record their observations pertaining to differences perceived in the video quality. We then attempted several different combinations of parameters used in our compensation algorithm to determine the values which would give the most power saving without noticeably degrading video quality. This phase was repeated for several different video streams and feedback from the subjects was recorded.

Based on our extensive analysis of the feedback received, we arrived at a set of values for the parameters in our compensation algorithm which gave us substantial power savings with very little degradation in video quality. The following sections describe our compensation algorithm and the middleware adaptations used in our system.

4. Dual Compensation Algorithm

In this section, we characterize the problem and propose our compensation algorithm. Let n be the number of backlight levels supported and $P(n)$ be the power at each level. Then $P_{\text{save}} = P(n) - P(n-k)$ denotes the power savings when dimming the backlight from level n to $n-k$. The perceived intensity of an image is denoted by

$$I = \rho LY_{\text{fr}}$$

where ρ is the transmittance of the LCD panel, L is the backlight luminance and Y_{fr} is the average luminance value of the frame [1]. The luminance value for a pixel (Y_{pix}) can be obtained from its RGB values, after applying standard conversion functions to convert it from RGB to the Y_C, C_b coordinate space [9]. Let the luminance of the backlight at level n be given by L and the luminance at level $n-k$ be given by L' . If we decrease the backlight level from n to $n-k$, then in order to preserve the perceived intensity, the new luminance value Y_{pix}' for each pixel in the frame is given by

$$\xi(Y_{\text{pix}}') = \min(1, \xi(Y_{\text{pix}}) + \Delta L)$$

where $\xi(Y_{\text{pix}}')$ gives the normalized value of Y_{pix}' , and

$$\Delta L = (1 - L'/L).$$

However, pixels already having a high luminosity value cannot be compensated adequately resulting in a loss of contrast due to saturation and we observe degradation in video quality. This limits the amount of compensation that can be applied to an image without degrading its quality intolerably. However, the loss in contrast can also be addressed and compensated for, as will be explained later, which then allows even further savings in power resulting from a larger reduction in backlight intensity. Next we introduce the concept of a group of scenes (GOS) which defines the granularity at which backlight compensation is performed. We define a group of scenes as a group of contiguous frames in a video stream such that the variance of the average luminosity values of each frame belonging to the group is less than a threshold value α . The average luminosity value Y_{fr} of a frame can be calculated as

$$Y_{\text{fr}} = \sum_{i=0}^{(w-1)(h-1)} \frac{Y_{w,h}}{wxh}$$

where w and h are the width and height of the image in terms of number of pixels and $Y_{w,h}$ is the luminosity of the pixel at (w,h) . The concept of GOS is used to split a video stream into several groups of frames. These groups form the basic entities on which compensation is performed. Video streams in general have the property that a lot of frames having similar average luminosity values are clustered together and this provides ample scope for optimization for low power by uniformly compensating entire GOS entities and reducing the handheld backlight level. Optimizing video streams where every frame is vastly different in its brightness from the previous frame is extremely challenging, but fortunately such video streams are not too common.

There must be a minimum number of frames β in a GOS for it to be eligible for compensation. The reason for introducing this parameter is that there needs to be a minimum duration between changes in backlight levels; otherwise the overhead of frequent switching of the backlight can result in a sharp increase in power consumed, due to increased CPU activity, another major source of power consumption in handhelds.

We introduce a function ' Ω ' that provides the backlight level compensation factor,

$$\Omega(k_i, \Gamma, Y_{\text{gos}}) = k'$$

The input parameters to Ω are the current backlight level (k_i), the type of the handheld (Γ) and the average luminosity of the GOS being considered (Y_{gos}) for which the number of frames $> \beta$ and the function returns the optimal backlight level to be set for the GOS (k'). The average GOS luminosity value is given by

$$Y_{\text{GOS}} = \sum_{i=0}^{n-1} \frac{Y_{\text{avg}}}{n}$$

where n is the number of frames in the GOS. This function uses the extensively profiled video information for a particular type of handheld stored in the proxy database to select and return a suitable value for backlight level based on the value of Y_{gos} and k_i .

Next we introduce another function ' σ ', that provides the video luminosity compensation factor

$$\sigma(k_i, \Gamma, \Omega(k_i, \Gamma, Y_{\text{gos}})) = c' = \Delta L$$

The input parameters for ' σ ' are the current backlight level (k_i), the type of the handheld (Γ) and the value for the next backlight level (returned by the function Ω) and it returns the brightness compensation value for the GOS being considered. The return value is the difference in luminosity for the two backlight levels, for the particular handheld device being used as a client. The average luminosity values for different backlight levels supported by the handheld are obtained from the database at the proxy, which stores these default values for several handheld

devices. If k_g is the current backlight level on a handheld of type Γ and Y_{GOS} is the average luminosity of an eligible GOS streaming to the proxy from the video server, then a control message is sent to the client asking it to set its backlight level to $\Omega(k_g, \Gamma, Y_{gos}) = n_g$ while at the same time the group of pictures is compensated with a brightness of $\sigma(k_i, T, \Omega(k_i, \Gamma, Y_{gos})) = c'$ before it is sent to the client. This results in power savings of $P_{save} = P(k_g) - P(n_g)$ over the time interval when the GOS is played back on the client.

Balancing contrast and luminosity: Increasing the luminosity of a frame can cause a loss in contrast between different regions in the frame, which makes it difficult to identify edges of objects in the frame and degrades picture quality. To overcome this, we propose performing an additional compensation step. In this step, the luminosity compensated frame is passed through a high pass filter which performs a spatial convolution on the luminosity values of the frame. This convolution step sharpens the edges and makes objects in the frame more recognizable. A convolution kernel is of the form

$$\mathbf{E}_k = \begin{bmatrix} c_1 & c_2 & c_3 \\ c_4 & c_{pix} & c_5 \\ c_6 & c_7 & c_8 \end{bmatrix}$$

where c_1, c_2, \dots, c_8 are values carefully selected to increase the amplitude of the high-frequency content in the frame. The convolution kernel is not limited to a 3x3 matrix and can be larger (e.g. 5x5). Now let the luminosity value of the pixel under consideration be L_{pix} and that of its 8 neighboring pixels be as shown:

$$L_{3 \times 3} = \begin{bmatrix} L_1 & L_2 & L_3 \\ L_4 & L_{pix} & L_5 \\ L_6 & L_7 & L_8 \end{bmatrix}$$

Then the modified pixel value after convolution is given by:

$$L_{pix}' = \frac{1}{\sum_{k=1}^8 c_k} \left(L_{pix} * c_5 + \sum_{k=1}^8 L_k * c_k \right)$$

$$\text{if } \sum_{k=1}^8 c_k > 0$$

and

$$L_{pix}' = \left(L_{pix} * c_5 + \sum_{k=1}^8 L_k * c_k \right) + 128$$

$$\text{if } \sum_{k=1}^8 c_k = 0$$

An alternative to the high pass filter is a median filter which is a non-linear filter that also sharpens a frame and also possesses the additional property of noise tolerance. This is generally harder and more time consuming to implement than the high pass filter described above. Our experiments have shown that this particular filter produces images that are of the same quality as those produced by the high pass filter for the video streams considered.

5. Proxy based middleware adaptation

A high level overview of the proxy based adaptation algorithm is presented in Figure 2. Note that the proxy maintains a database of the number of backlight intensity levels for different types of low-power devices and the average luminosity due to backlight at each of these levels for reference images (e.g. all black and all white). This database also maintains a record of the group of scenes (GOS) and the average luminosity values (Y_{GOS}) for every GOS in the requested video for frequently used values of α and β .

Using the control information, the client middleware sets the backlight levels and displays the corresponding GOS of the video stream. We now present three middleware adaptation policies that utilize the compensation algorithm. The first two policies can be implemented with limited operating system interface control and the third is our proposed dual compensation algorithm that requires both proxy and OS interfaces for optimal operation.

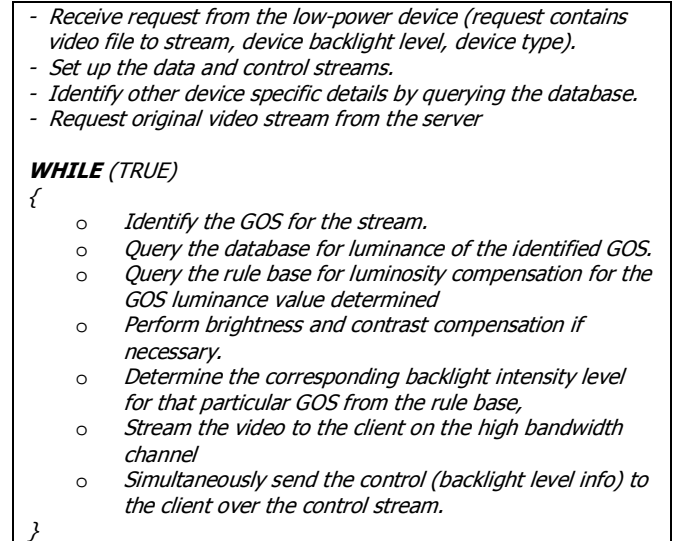


Figure 2: Overview of the Proxy based Adaptation algorithm

Simple Backlight Compensation (SBC): Using this policy, power can be saved by identifying GOS entities for which Y_{GOS} is high (above a threshold level τ) and reducing the backlight level at the client for these GOS entities. Note that the proxy simply calculates new backlight levels

MPEG Video	Resolution	FPS	Duration (sec)	Luminosity Variation	Video Type
bipolar.mpg	320 x 240	30	41	Little	Dark, 3D animation
iceegg.mpg	240 x 136	30	59	Moderate	Bright, 3D animation
intro.mpg	160 x 120	30	59	Very High	Flashy, TV show clip
simpsons.mpg	192 x 144	30	27	High	Colorful, 2D animation

Table 3: Characteristics of video streams used in experiment




Figure 3: Snapshots of MPEG-1 video streams

without compensating the brightness of the video stream. The proxy has a record of the GOS entities in the video stream and whenever it expects to be sending a GOS with $Y_{GOS} > \tau$, it sends control information to the client to reduce its backlight level, such that the perceived difference in quality is minimal. This scheme has the disadvantage that it slightly degrades video quality in every case and is attractive only because of its simplicity (no video compensation performed).

Parameter	Description	Value
β	Minimum number of frames in GOS	60
α	Variance threshold for Y_{GOS}	40
τ	Threshold level for SBC scheme	220
\mathcal{E}_k	Convolution kernel used in high pass filter	$\begin{bmatrix} 0 & -1 & 0 \\ -1 & 5 & -1 \\ 0 & -1 & 0 \end{bmatrix}$

Table 2: Parameter values for compensation algorithm used in experiments

Constant Backlight with Video Luminosity Compensation (CBVLC): A more interesting and practical approach is to set a constant backlight level at the start of the video stream and then dynamically compensate different GOS entities based on their Y_{GOS} values. However, since the backlight value is fixed for the entire duration of the video, the level has to be chosen conservatively so that video quality is not affected adversely if there are dramatic variations in luminous intensity of consecutive GOS entities in the video stream. This is a static approach and would provide a constant power saving (depending on the initial backlight level chosen), with no dynamic adaptation of backlight intensity levels. The limitations of the previous approaches can be overcome with a hybrid dual

compensation approach described below.

Dual Compensation Approach (DCA): In this technique, we simultaneously compensate the video stream and the backlight levels for different GOS entities. The proxy dynamically compensates the GOS entities in the video stream and begins streaming the video to the client, simultaneously directing the client to change its backlight level through the control stream. The client middleware sets the appropriate backlight intensity levels for the video playback. This approach provides more flexibility for aggressive optimizations with much greater power savings.

6. Performance Evaluation

To evaluate these approaches, we conducted experiments with several different MPEG-1 video streams and a Compaq iPAQ 3600 series mobile handheld device, which comes with a color reflective thin film transistor LCD screen with a pixel pitch of .24mm and a display resolution of 240 x 320 pixels.

The parameters we used for our compensation algorithm are given in Table 2. In addition, the rule base for determining $\Omega(k_x)$ and $\sigma(k_x)$ for the five backlight levels (k_0, k_1, k_2, k_3 and k_4) for the Compaq iPAQ is given below

$$\begin{aligned} \Omega(k_0) &= k_0; \sigma(k_0) = 0; && \text{for all } Y_{GOS} \\ \Omega(k_1) &= k_1; \sigma(k_1) = 0; && \text{for all } Y_{GOS} \\ \Omega(k_2) &= k_1; \sigma(k_2) = 30; && \text{for } Y_{GOS} < 140 \\ \Omega(k_3) &= k_2; \sigma(k_3) = 30; && \text{for } 80 < Y_{GOS} < 190 \\ &= k_1; \sigma(k_3) = 55; && \text{for } 190 < Y_{GOS} < 220 \\ \Omega(k_4) &= k_3; \sigma(k_4) = 30; && \text{for } 190 < Y_{GOS} < 220 \\ &= k_2; \sigma(k_4) = 55; && \text{for } 60 < Y_{GOS} < 190 \\ &= k_1; \sigma(k_4) = 65; && \text{for } Y_{GOS} < 60 \end{aligned}$$

The luminance threshold values and the parameters in Table 2 were determined after extensive profiling and subjective assessment of several video streams to achieve substantial

power savings without significant degradation in video quality and stored in the database at the proxy. Our compensation algorithm then queries these parameters for different GOS entities to determine the amount of compensation to perform.

We analyzed the effectiveness of the three schemes described in the previous section for several different video streams of different durations (number of frames), resolutions (frame width, height) and types (containing dynamically/statically changing scenes). Due to lack of space we present the results for a limited set of streams. Figure 3 shows the video streams used and Table 3 describes them in detail. For all our experiments, we assume that the initial level of the handheld is set to a brightness level that takes into account the brightness of the client's environment, such that the quality of the video playback is deemed acceptable by the user at the set backlight level. Our schemes are not applicable if the initial backlight setting is low or on power save mode.




Figure 4: Power saving for initial reference point = super bright

The power save mode is the theoretical limit beyond which power savings cannot be expected since the backlight is switched off in this mode. For the low power mode, we found it to be very difficult to maintain video quality after reducing the backlight by switching to the power save (no backlight) mode, hence we chose not to apply our schemes in this case. This information is captured in the rule base described earlier.

The first experiment assumes that the initial backlight level of the handheld is set to super bright. Figure 4 shows the results for this case. It is easy to see that the DCA scheme outperforms the other two schemes in all the cases. The CBVLC scheme outperforms the simplistic SBC scheme in all cases except for *intro.mpg*. This video has a very high variation of average luminous intensity from one frame to the next and very high intensity values for some frames. Consequently there are very few GOS entities eligible for compensation in this case.

Since the CBVLC scheme sets the client backlight level just once (in the beginning), lowering the backlight level can result in significant degradation of quality for many of the

frames in this case, which the compensation cannot rectify. As a result, the CBVLC scheme takes the conservative approach and does not request the client to lower its backlight level from its preset value. It is interesting to note that for the case of *bipolar.mpg*, the SBC scheme does not provide any power saving. This is because the video is very dark (low luminance) on an average and there is no GOS entity with a value of $Y_{GOS} > \tau$. Consequently, for this video, the backlight level is never reduced in this scheme.




Figure 5: Power saving for initial reference point = high bright

Figures 5, 6 show the results when we assume that the initial backlight level of the handheld is set to high bright and medium bright respectively. As we lower the initial backlight levels, we expect that the scope for reducing power consumption in the handheld decreases, which can be seen from the results. For the experiment with the initial level set to medium bright, it is interesting to note from the figure that both the CBVLC and scheme DCA schemes perform more or less the same. An exception to this observation is the case of *intro.mpg* (for reasons mentioned earlier).




Figure 6: Power saving for initial reference point = medium bright

The reason for the similar performance of the two schemes is that the DCA scheme can only reduce the backlight by 1 level at most – further reduction makes it difficult to maintain quality even with video compensation. This fact is

reflected in our rule base. The CBVLC scheme also manages to lower the backlight level by 1 level for all the cases except *intro.mpg* and hence performs similar to the DCA scheme. Overall, it can be seen from these experiments that the DCA scheme performs just as well as the other two schemes in a few cases, but performs much better compared to them in a majority of the cases. The power savings for the DCA scheme ranges from 100 mW to 625 mW depending on the type of video being played and the initial backlight setting, which is a range that corresponds to roughly 9% to 60% reduction in power consumed by the backlight on the mobile handheld client.

7. Conclusion

We have proposed middleware adaptation policies that are used in conjunction with a brightness compensation algorithm to reduce backlight power consumption for video playback on low-power mobile devices. For the Compaq iPAQ the backlight can consume as much as 40% of the total power when playing streaming MPEG video. To reduce the contribution of the backlight to overall power consumption, we proposed reducing the handheld backlight level while simultaneously compensating the video stream by increasing average frame luminosity, and convolving the frame with a high pass filter to improve picture detail after aggressive luminosity compensation. We presented three middleware adaptation techniques that utilize a proxy server to implement our approach. Our experiments with several different MPEG-1 streaming videos on the Compaq iPAQ handheld show that the dual compensation approach (DCA) – which is a hybrid of the other two proposed middleware adaptation approaches – is effective in reducing as much as 60% of the power consumed by the backlight on the handheld for certain video streams.

8. References

- [1] Inseok Choi, Hojun Shim and Naehyuck Chang, Low-Power Color TFT LCD Display for Hand-Held Embedded Systems, In ISLPED, 2002.
- [2] C. J. Hughes, J. Srinivasan, and S. V. Adve. Saving Energy with Architectural and Frequency Adaptations for Multimedia Applications. In MICRO-34, Dec. 2001.
- [3] Pillai, P., and Shin, K. G Real Time Dynamic Voltage Scaling for Low Power Embedded Operating Systems. In SOSP (2001)
- [4] Mohapatra, S., and Venkatasubramanian, N. PARM: Power-Aware Reconfigurable Middleware. In ICDCS-23 (2003).
- [5] Noble, B. D., Satyanarayanan, M., D.Narayanan, J.E.Tilton, and Flinn, J. Agile Application-Aware Adaptation for Mobility. In SOSP (October 1997).
- [6] ITU-R Recommendation BT-500.7, Methodology for the

subjective assessment of the quality of television pictures. In *ITU Geneva Switzerland* (1995)

[7] Sarnoff® Corporation white paper, JND: a human vision system model for objective picture quality measurements. In <http://www.sarnoff.com> (2001).

[8] S. Chandra. Wireless network interface energy consumption implications of popular streaming formats. In *Multimedia Computing and Networking MMCN'02*, 2002

[9] Ralf Steinmetz, Klara Nahrstedt. *MULTIMEDIA: Computing, Communications and Applications*. Prentice Hall, July, 1995

[10] Fred Douglass and P. Krishnan and B. Marsh, "Thwarting the power hungry disk", WINTER USENIX conference, Jan 1994

[11] Shivajit Mohapatra, Radu Cornea, Nikil Dutt, Alex Nicolau, Nalini Venkatasubramanian, "Integrated Power Management for Video Streaming to Mobile Handheld Devices", ACM Multimedia 2003, Berkeley, CA, 02-08 November 2003



Sudeep Pasricha received the B. E degree in Electronics and Communications from Delhi Institute of Technology, Delhi, India in 2000. From 2000 to 2002 he worked as a design engineer in the Center for Research and Development at STMicroelectronics, Noida, India and Crolles, France where he worked on an Architecture Description Language (ADL) driven processor exploration framework. He was also involved in developing a High Level System-on-Chip modeling methodology. He is currently pursuing his Ph.D. degree in Information and Computer Science from the University of California, Irvine. His research interests include high level modeling methodologies for SoC architecture exploration, branch prediction schemes and adaptive middleware based methods for low power distributed operation. He is a member of the Architecture and Compilers for Embedded Systems (ACES) group and is actively involved in the development of the EXPRESSION ADL driven Design Space Exploration framework.

Sudeep was awarded a scholarship for outstanding academic achievement in Electronics and Communications at Delhi Institute of Technology in 1997. He is currently on a fellowship from the School of Information and Computer Science at UCI for the academic year 2002-03.



Shivajit Mohapatra received his Masters Degree in Computer Systems and Networking at University of California, Irvine in 2001 and his B.Tech degree in Computer Science from Birla Institute of Technology and Science, Pilani, India in 1997. He worked at Wipro Technologies from 1997 to 1999, where he was involved with the design and development of U/WIN

(Unix for Windows) and the Wipro H.323 protocol stack. He is currently working towards his Ph.D. in Information and Computer Science at University of California, Irvine. His research interests

include Power-Aware Middleware frameworks, Low-power systems and Distributed Systems.

Shivajit received a fellowship from the School of Information and Computer Science for the academic year 1999-2000.



Manev Luthra received the B.Tech. degree in Electrical Engineering from the Indian Institute of Technology, Delhi, India, in August 2001. From January 2001 to September 2002 he worked on the design and development of wireless systems-on-chip and home gateways, initially at Angeles Design Systems and then at Atheros Communications Inc. Since then he

has been pursuing the Ph.D. degree in Information and Computer Science at the University of California at Irvine. His research interests include embedded systems design, low power system design and distributed systems.

Manev received a fellowship from the School of Information and Computer Science at UCI for the year 2002-03.



Nikil D. Dutt received a Ph.D. in Computer Science from the University of Illinois at Urbana-Champaign in 1989, and is currently Professor in the Center for Embedded Computer Systems at the University of California, Irvine, with academic appointments in the CS and EECS departments. His research interests are in embedded systems design automation,

computer architecture, optimizing compilers, low power design and system specification techniques. He is a coauthor of three books: "High-Level Synthesis: Introduction to Chip and System Design", Kluwer 1992, "Memory Issues in Embedded Systems-on-Chip: Optimizations and Exploration", Kluwer Academic Publishers, 1999, and "Memory Architecture Exploration for Programmable Embedded Systems", Kluwer Academic Publishers (March 2003). He received best paper awards at CHDL89, CHDL91 and VLSIDesign2003. Professor Dutt currently serves as Associate Editor of ACM Transactions on Design Automation of Electronic Systems (TODAES). He was an ACM SIGDA Distinguished Lecturer during 2001-2002, and is currently an IEEE Computer Society Distinguished Visitor for 2003-2005.

He has served on the steering, organizing, and program committees of several premier CAD and Embedded System Design conferences, including ASPDAC, DATE, ICCAD, CODES/ISSS, CASES, ISLPED and LCTES. He is a senior member of the IEEE, serves on the advisory boards of ACM SIGBED and ACM SIGDA, and is Vice-Chair of IFIP WG 10.5.



Nalini Venkatasubramanian is an assistant professor at the Department of Information and Computer Science, University of California, Irvine. Her research interests include distributed and parallel systems, middleware, real-time multimedia systems, mobile environments and formal reasoning of distributed systems.

She is specifically interested in developing safe and flexible middleware technology for highly dynamic environments. Nalini was a member of technical staff at Hewlett-

Packard Laboratories in Palo Alto, California for several years where she worked on large scale distributed systems and interactive multimedia applications. Nalini has also worked on various database management systems and on programming languages/compilers for high performance machines. She has an M.S. and Ph.D. in Computer Science from the University of Illinois, Urbana-Champaign and is a member of the IEEE and ACM. She is a recipient of the NSF Career Award and a Teaching Excellence Award from the University of California, Irvine.