Topic Review: Multimedia for Mobile Devices

Review: Integrated Power Management for Video Streaming to Mobile Handheld Devices

Published at the ACM Multimedia conference in 2003, "Integrated Power Management for Video Streaming to Mobile Handheld Devices" presents a relatively thorough treatment of the problems affecting user experience of multimedia on handheld devices. While it has its shortcomings, I believe it is a good paper that makes a solid contribution.

The paper does a good job at the beginning of laying out the motivation for the research by describing an inherent technological conflict, which is that the resource constraints of handheld devices and the resource requirements of multimedia are fundamentally at odds. This problem will continue to increase as both streaming media and the use of mobile devices continue to proliferate. The paper mentions that this problem has been addressed in the past by attempts to optimize at individual computation layers (for example, dynamic voltage scaling, or cache optimizations). In contrast, this paper takes the approach of optimizing multiple layers with respect to each other, thereby achieving greater overall gains.

An obvious question arises whenever a paper discusses optimization techniques, which is what metrics are being optimized for, and are they useful ones to target? The authors of this paper address this up front, stating that they are optimizing for "user experience" and power consumption. The user experience of streaming media is indeed a useful metric to optimize for since this most affected by the particular constraints of the handheld platform. However, it is also a relatively "soft" metric that is difficult to quantify. The authors appear to be aware of this and take two approaches to measuring user experience:

1. They define a "utility factor" function that is intended to provide a quantitative measure of user experience by combining several other performance metrics
2. They conduct user satisfaction surveys to qualitatively assess the user experience improvements of their optimizations.

While the follow-through on these measurements in the evaluation section falls short, to some extent, of what is promised in the earlier sections (which I'll discuss later), the paper does a good job at the beginning of establishing a foundation for evaluating the effectiveness of optimizations for streaming media on handheld devices.

The authors implement their optimization approach in two complementary strategies. First, they identify and then tuning hardware-level "knobs" on the handheld device. Specifically, they focus on two aspects of the underlying hardware: the cache configuration (size and associativity) and the CPU voltage,

both which they adjust dynamically to target specific quality levels for the streamed video. Second, they develop middleware software that runs both on handheld devices and on a proxy server through which media streams to devices are routed. The middleware running on each device senses the level of residual energy and communicates this information to the proxy server. The proxy server then transcodes the video stream into a format that optimizes user experience given the available energy on the device.

The authors do a good job of demonstrating in the paper's evaluation section that the optimal cache configuration, DVS settings, and video quality to maintain a certain energy threshold are all highly interdependent, which establishes the case that a strategy that can dynamically optimize for all these parameters at once would be beneficial. Further they demonstrate significant reductions in consumed energy in their simulations, which makes their solution compelling. One disadvantage to the system they have proposed is its complexity, particularly the middleware infrastructure that establishes a feedback channel between each device and a proxy server capable of transcoding media streams. The authors do not discuss the algorithms used by the proxy server to interpret device feedback and choose an appropriate transcoding scheme. Neither do they discuss other challenges with architecture they have proposed, such as scalability and the proxy server as a single point of failure.

They also do not sufficiently justify their choice of "knobs". For example, they state that in the cache reconfiguration experiments they only consider the data cache, and do not consider reconfiguration of the instruction cache because it is out of the scope of the paper. Yet figure 4 clearly indicates the instruction cache is by far the most significant power consumer among the CPU's functional units. One wonders, therefore, why it is out of the scope of this paper.

In addition, the paper is unconvincing in its argument that the proposed solution increases user experience or satisfaction with viewing video streams on their mobile devices. The utility factor function (U_F) to objectively measure user satisfaction is an interesting and compelling idea. However, the definition given for it in the paper confusing and little support is given that it is in fact an accurate measure. The U_F is defined in the paper as the difference between the maximum possible quality level (Q_{MAX}), and the quality level at which a video is currently playing (Q_{PLAY}) while staying beneath an acceptable power consumption threshold (E_{RES}) and above an acceptable quality threshold (Q_A). If either threshold is crossed, U_F is -1. This definition is confusing because it means that U_F decreases as the quality of play increases, so that it the same video could be played at a higher quality while maintaining the same power consumption rate, it would have a lower U_F. It could be the case that U_F is inversely proportional to user satisfaction; however, figures 16 and 17 clearly indicate that a higher U_F is desirable.

Though the paper does cite user satisfaction surveys that were conducted, it does not mention the types of questions that were asked, and the conclusions listed in section 2 appear only tangentially related to main focus of the paper. In particular, I was looking for results comparing the level user satisfaction with the proposed design against satisfaction levels with a typical system. All that is given are figures 16 and 17 which show that the proposed design attains a higher U_F. However, there is no demonstrated correlation between U_F and actual user satisfaction. The user surveys would seem to have been the

perfect opportunity to show that U_F is indeed an accurate measure of user experience, but this was not done.

In summary, the paper does a good job of establishing a necessary and viable area of research, and proposes an interesting solution that achieves good results under simulation. In addition, the techniques proposed to dynamically adjust at multiple computation layers are compelling. However, the paper is not as thorough or holistic in its evaluation as it purports to be, and it would have been stronger if it more sufficiently backed up its claims.