Networking for Multimedia
Tech Topic #07: Multimedia for Mobile Devices

T. Pering, Y. Agarwal, R. Gupta, and R. Want, "CoolSpots: Reducing the Power Consumption of Wireless Mobile Devices with Multiple Radio Interfaces," ACM Mobisys, Uppsala, Sweden, June 2006.

Coolspots is definitely a catchy name, although what is in a name, after all. Actually, going back through the paper it was difficult to ascertain what exactly the moniker refers to—at first the system as a whole was assumed, but an early figure actually labels the Bluetooth access points individually as the "Coolspots". This was right above another figure that grouped both Bluetooth and Wi-fi together as "Coolspots". It is supposed that such usage reinforces the system's intended ubiquity.

The basic idea behind Coolspots is for mobile devices to dynamically switch between different wireless interfaces (multiple radios) on the basis that shorter range, lower power communication modes can be leveraged, when available, to conserve battery life.

Section 2 does a good job of motivating the work by displaying a pie chart where wi-fi is shown consume greater than 2/3 of the power in mobile devices—at least at first glance. The reviewer felt slightly cheated to find that the chart depicted a connected device in idle mode with all LCDs off, including backlighting. Some reasoning on why this state was chosen would have been nice, given that one not familiar in the field might imagine it would be more common to have a device disconnected/idle or connected/in use, than the previous case. Particularly given that the point (Wi-fi burns a large amount of power) probably could have been made fairly dramatically anyway.

The Coolspots concept is introduced and related work is discussed, both fairly standard. The main contribution of the paper introduces several switching policies and compares them. In addition to baseline policies such as only Wi-fi and only Bluetooth, Coolspots presents switching based on bandwidth and channel capacity thresholds. Benchmarks used to compare these policies are idle, streaming (multimedia), web browsing and all you can eat bandwidth buffet (all that's available).

In the middle of reading the authors' experimental setup it suddenly occurred that Coolspots is oriented more towards laptop-type mobile devices than cell phones. Perhaps because Mobisys always makes the reviewer think smartphone-type devices, but a major critique of the paper was greatly lessened by this realization. The infrastructure required to support adequate Bluetooth coverage in the context of phone-like mobility—think how far people wander when speaking on a cell phone—appeared daunting. However, if for example the specific application is improving laptop battery life when streaming multimedia, it is far easier to imagine a strategically-placed Coolspot or two, say, near one's favourite armchair.

A more general difficulty with Coolspots is the lack of any spectrum between the ranges (and, indirectly, power consumption) of Bluetooth and Wi-fi. An order of magnitude apart with nothing in between makes it difficult to fine-tune switching policies. Although this is almost certainly the result of prevalent wireless technologies, a brief discussion along these lines may have been appropriate.

S. Mohapatra, R. Cornea, N. Dutt, A. Nicolau, and N. Venkatasubramanian, "Integrated Power Management for Video Streaming to Mobile Handheld Devices," ACM Multimedia, Berkeley, California, November 2003.

It was surprising to find that this paper was published in one of the premiere multimedia conferences, ACM Multimedia. The contribution of the content is significant, however the paper is so badly organized as to jeopardize its acceptance, in this reviewers opinion. It is possible, of course, that we have become so accustomed to a particular formula that anything that deviates from it even slightly is classified as "badly organized". It is the first paper seen in quite a while which does not contain the obligatory "the rest/remainder of this paper is organized as follows". It was also somewhat odd to place related work at the end.

Particular formula aside, it was difficult to extract the main objectives amidst what might best be described as poor flow. It took until near the end of the paper to fully understand what the authors were trying to do. There was no question that a significant amount of work had been done; the presentation of this effort was lacking, however. Concepts and goals seemed to be tacked on to random sections, only to appear again later in the paper in a different context. The dynamic voltage scaling notion is a good example of this.

Contribution-wise, the paper attempts to accomplish quite a bit, which makes it somewhat difficult to critique any particular part. The style seems to be in a constant rush to get through all the material, and as a result key details are often glossed over or ignored completely. There is an entire subjective video assessment study thrown into section 2, which will more often than not comprise an entire journal-level paper in its own right. The results of this study are presented almost in passing and detailed results are omitted.

The body of the work involves a cross-layer design incorporating low-level architectural optimizations and middle-ware techniques. The architectural section focuses mostly on cache performance, but also confuses things by throwing in various other tidbits such as compiler optimizations and MPEG decoding tweaks. The middle-ware techniques are essentially integrating the former with connection-admission control (CAC).

The exploration of cache performance is interesting and fairly complete, however, it lacks generalization. All is well and good to perform a detailed analysis of cache behavior on an iPaq, but nothing in the paper implies that any of this effort is applicable to other devices. Rather, it seems to be the case that in order to use the feedback mechanism employed in the system, an exhaustive search must first be conducted on each particular hardware platform in the network, as the middle-ware proxy must be aware of this information. Cross-layer design can be very useful; it can also be dangerous. Layered designs are usually in place for a reason, and the paper does not effectively discuss this trade off.

Another point assumed throughout is full network connectivity of all nodes at all times, without justification.

On a lighter note, the authors found a way to usefully present 3D graphs in black and white.

H. Shojania and B. Li, "Random Network Coding on the iPhone: Fact or Fiction?," ACM NOSSDAV, Williamsburg, Virginia, June 2009.

Not as much to say about this paper but for a desire to express how well-written it was found to be, in direct contrast with the previous review. Excellent flow with motivation, background, objectives, and results all nicely presented.

The hand-tuned optimization section was particularly elegant (perhaps the reviewer is too much of a software geek).

It may have been a bit more thorough to comment on what was lost, either in code density or otherwise, by compiling with the superset ARM ISA rather than only mentioning what was gained.

Towards the end the content got a bit less clear as the discussion turned more towards work-in-progress types of things, but this is perfectly fine for a workshop paper. The metrics used for results were exceptionally straightforward and easy to understand.

This paper reflects well on what is probably one of if not the premiere workshop in multimedia networking as it celebrates its 20th anniversary this year.