Paper Analysis

CoolSpots: Reducing the power consumption of wireless mobile devices with multiple radio interfaces

This paper addresses the biggest problem (in my opinion at least!) plaguing the smart phones and other handheld/mobile devices: battery (or power) consumption. Communication is the biggest energy guzzling component of a mobile device with Wifi, 3G, GPRS, Bluetooth eating most of the power from the battery than any other process, unless someone develops an atrocious application for the device. Amongst them, Wifi is the most power consuming whereas Bluetooth is the least. CoolSpots is a system which enables a mobile system to dynamically switch between the modes of communication depending upon basically the bandwidth requirements enabling reduction in power consumption.

The paper is pretty well organized and is quite an easy read, except in places where technical specifications are thrown at the user. The authors start off by introducing the reader to the problem and giving a brief overview of previous research in solving it. They also provide a brief overview of the contributions of the paper. Next, the motivation behind the idea is explained. The authors argue that in the future, this strategy will be more important because many devices will have both bluetooth and wifi capabilities. However, the authors do not touch upon 3G networks which I feel are way more important than just the bluetooth and wifi. After all, a mobile user tends to use networking when he is away from his stationary computer or laptop, whence 3G comes into action. It is at the very end that they mention that this could be used for 3G. Nevertheless, I like the way in which the authors explore the two modes (bluetooth and wireless) and lead the reader into their motivation behind CoolSpots taking into account the three factors that matter: bandwidth requirement, power and distance.

The core idea in CoolSpots is to 'switch down' to a lower power consuming mode of communication (bluetooth) when the bandwidth requirement is low and 'switch up' to a higher bandwidth mode of communication (wifi) when requirement is high. For example, an application like email client which downloads new emails does not require as much bandwidth as a youtube application (or any other video streaming app). Even though this idea sounds effective, I am concerned about the overhead caused by this switching. Wouldn't it cause additional power consumption? Especially with irregular data rate data, the switching would happen regularly. The authors do mention something about it in the results section, but not very clearly. This is something similar to the application I am working on where we have to be careful about the fact that our application which aims to improve battery life by doing 'some operations', doesn't itself eat too much power.

The references mentioned by the authors in the 'related work' section are a bit confusing. They do not directly relate to the topic. They site that some power optimization work has been done in the application layer and the MAC layer and their work provides a solution in the network layer. However, they do explain very well how their solution is different from other related work (which it would be since other works are not very related!).

The authors then explain the switching policies used to make the decision of when to 'switch up' and when to 'switch down'. These policies are executed by the mobile device which makes me come back to the same question of power overhead when computing these policies. The mobile device also needs to let the base station know that it has switched. The CPU time used by this switching process could use up a considerable amount of energy, especially when the data rate is not regular and somewhat bursty. Also, in such circumstances, this network route change traffic could create additional network overhead (which I suspect would not be a very big overhead). Switching would also mean that the devices, both mobile and the base station should know about both the connections: BT and WiFi.

Next the authors establish some Baselines, listing some baseline policies starting with wifi-CAM whose full form is not listed anywhere in the paper (or I missed it) or maybe it is something which the reader should know? So the policies used/proposed by the authors are: bandwidth-X, cap-static, cap-dynamic in addition to the baseline policies. Cap-static and Cap-dynamic policies use a frequent ping to determine the situation of the network and decide whether to switch up/down or stick at the same level. I would like to mention a point that the authors mention here and which I found very cool if one thinks about the motive and theme of the whole paper: “the true strength of the switching policies is their ability to work well across all benchmarks and work with dynamic workloads”. And they have gone ahead and provided results to corroborate it which goes to show that they had a clear goal in mind. I also like the way they have honestly put forward a shortcoming of the cap-dynamic policy, that it assumes that the channel characteristics do not change significantly during the switch-up state, which may cause sub-optimal configurations.

The Benchmarks section is an interesting section in the way that it encompasses a good variety of scenarios possible with the system and the authors test the system using these benchmarks. The metric used is not power consumed but power consumed over average time, which they call it by a new term: “Communications Energy”. They calculate only the energy spent due to communications. This metric however, does not involve the power consumed due to switching between communication modes and the computation of the data received. Data received over a slow connection (bluetooth) would take more time to be processed and thus would eat more power. However, this is not included as a metric for some unknown reason. The scenarios used by the authors are: streaming videos, web traffic, file transfers and and idle state. I particularly like the way authors have explained the reasoning behind their choices. For example, they chose the bit rates of streaming to be 128, 250 and 384 kbps because higher bit rates are seldom used for mobile devices (even though bluetooth can support data rates upto 780kbps). Maybe they should have included a benchmark of multiplayer gaming? Gaming requirements are somewhat different from all these benchmarks in the way that the data transfer rate is quite irregular, it is real time and is highly user interactive. Plus, these games are pretty popular in the mobile devices.

As I started reading the experimental setup section, I expected the authors to have forgotten about the scenario where obstacles are present. However, it was great to find the situation accounted for in the chosen locations. But I would refrain from saying that the setup was through for I believe that they should also have had some testing scenario with interference from some other devices. With more active wifi devices around, wouldn't there be some effect on the transfer rate? Especially the 'ping' mechanism to find the RTT of the network. What if the ping packet was lost or corrupted due to interference? They later mention channel interference in the passing while talking about location configurations: “although not directly measured, increasing the distance simulates the effect of other kinds of channel interference”. But I still feel, such interference mechanism should have been used for explicit results.

Another point which I found a bit misleading was that the measurements are done at various locations and not while moving from one location to another, i.e. dynamic mobility. Their argument is somewhat weird: generally people stay in one spot while carrying out such operations. A user might move around while doing a file transfer from his mobile device, which is why it is a mobile device in the first place! A test case considering this scenario should have been there. Another test case the authors could have added is testing on different mobile devices platforms. The use of at least two different platforms could have provided useful insights into the overall power consumption and not just the communication energy spent.

The results from the experiment are somewhat favorable for the cap-dynamic policy. As expected (and mentioned), there is not much improvement in the file transfers. In fact the wifi fixed policy performs much better than the dynamic policies. The cap-dynamic policy performs better only in the higher data rates video streamings, other than that it doesn't seem to offer much benefits. It does offer some help when the distance between the mobile device and base station is greater, where some of the static policies fail to perform.

The paper is quite well written and the evaluation has been done well enough. But some of the reasonings are weird, such as not playing around with 3G, which in fact would be way more important than just wireless and bluetooth. A user is more likely to use networking on his mobile device when he is not around his primary machine to access networks. Or, he could use the device to play multiplayer games, but that case also has not been looked into. Nevertheless I totally agree with the idea that not a great deal of infrastructure change is required (only some switching mechanism needs to be introduced in both base station and mobile devices) as the devices nowadays are equipped with both BT and wifi. Therefore, I think that the idea would be economically viable as well!