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.

This paper proposes a new method, called CoolSpots, for dealing with network traffic on mobile phone devices. The purpose of CoolSpots is to reduce battery consumption since "the utility of mobile devices is directly impacted by their operating lifetime…" The idea behind CoolSpots is to implement a switching policy that will determine whether to use Bluetooth or WiFi for network traffic since the paper claims that "wireless communication subsystems account for a major component of the total power consumption… of these devices." The paper claims that such a switching policy would be effective in reducing battery consumption because of how much better Bluetooth performs in idle mode compared to WiFi. The paper also makes a pretty bold claim saying that "more than a 50% reduction in energy consumption by the wireless subsystem is possible, which can effectively double the system battery lifetime."

The paper then goes on to discuss some details in regards to the mobile computing platforms. The authors mention that latency and bandwidth requirements change depending on the context or application – a relatively trivial statement. They then discuss Bluetooth and Wifi in slightly more detail in terms of power consumption and bandwidth capabilities. Basically, WiFi offers coverage for up to 100 meters and is most well utilized in high bandwidth situations (up to 11 Mb/s). Bluetooth on the other hand, only works up to 10 meters but consumes much less power (bandwidth restricted to 1 Mb/s). A measurement is then provided for the two radio interfaces while they are in an active transfer state which is 890 mW and 120 mW for WiFi and Bluetooth respectively.

After reading the introduction, I was very impressed with the proposed idea. After examining figure 1, it seems like the advantages of deploying such a system are significant. The authors brought up a lot of interesting points such as how Bluetooth consumes much less battery but transfers at a much slower rate compared to WiFi which is significantly faster but consumes much more battery, especially in idle mode. Later in the paper it is mentioned that WiFi was not designed with battery constraint in mind which I also believe is very valid in a study such as this one. Figure 2 was a nice addition as well, however I did not quite understand what the authors meant by "connected" in the figure's description. I assume it means connected to a power supply but it could also mean connected to an access point, Bluetooth station, etc. The authors should have been clearer when it came to this detail. Another interesting point that was discussed was the fact that since WiFi covers a larger range. This is an important factor to the power consumption since there is the potential for more noise on the wireless channel which would set off CSMA counter-measures, hence using up more battery for re-transmissions.

The paper then goes on to talk about CoolSpots in more detail. The authors mention that based on table 1, "this concept has the potential to realize 10x reduction in power consumption for an idle system…" After examining table 1, it was not clear to me how they drew this conclusion and I wish they would have discussed it in more detail. The authors then go on to claim that deploying CoolSpots would be "relatively inexpensive." I think this is a bit of a bold statement. Most homes already have a wireless access point set up. Deploying a CoolSpot in these homes would require the purchase of a completely new device (that does both WiFi and Bluetooth) or the purchase of just the Bluetooth device. Once the new device is obtained, it would still require new installation and configuration (since the general population probably would not be able to do this on their own) which is an additional cost.

One metric which I found particularly interesting was that even though Bluetooth is much more efficient when it comes to power consumption, it is actually less efficient when you look at it from an energy perspective – the combination of power and time. The paper states that WiFi consumes 81 nJ/Bit while Bluetooth consumes 120 nJ/bit. This claim almost makes it seem like using Bluetooth for network traffic is irrelevant due to the speed gains by using WiFi. I think the paper also does a good job of acknowledging that since much of mobile phone usage is subjective, it is hard to truly gauge what each user expects out of their experience with applications requiring network traffic.

The paper then goes on to discuss the three different switching policies that were implemented – bandwidth, cap-static, and cap-dynamic. The main challenge is to determine when to switch from Bluetooth to WiFi and vice versa. The authors also mention that you do not want to implement a policy which will switch too often because that can lead to increased battery consumption which is clearly undesired. I was glad when the authors mention that "the real world is less ideal as the underlying channel capacity can change due to distance from the base station, interference, obstacles or other circumstances…" It seems like many of the previous papers we have read assume perfect environments as their test bed where this clearly is not the case in the real world. One problem that I had with their evaluation of the different policies was that they did not perform an exhaustive search of the parameterization space. The authors say they chose the parameters they did because the "behave relatively well for the given workload." I believe it would have been more beneficial to find the parameters that behave the best.

In section 6.1, the authors discuss the two basic benchmarks, idle and transfer, and then state that "these two benchmarks are not indicative or a real system…" I found this a little odd. If these benchmarks do not behave like real system, why are they being used? After explaining the experimental setup which I thought was well-done, the paper goes ahead and explains the results. I thought this section was somewhat lacking. For instance, I think figure 7 really could have used a time axis. If you look at the graph as it is, it seems like bluetooth-fixed performs the best for almost all the benchmarks with the exception of the two file transfers where it does not seem to perform much worse than wifi-fixed. Based on this it seems like there is no need for a switching policy. Also, based on the authors' previous descriptions of the different policies, it seemed like the expected cap-dynamic to perform the best which is not evident in the graph depicted in figure 7.

Overall, the paper was very well written and easy to follow which I appreciated. I think the potential for a significant contribution exists in this paper but requires a little more work. There are a few more topics I would have liked to see discussed in this paper such as the following: As previously mentioned, a description of what it would cost to deploy such an infrastructure would have been beneficial. Also, I do not think the authors spent enough time talking about how the base station would work. They keep mentioning that implementing the switch between Bluetooth and WiFi at these stations is easy and trivial but never go into detail as to why that is. Additionally, the authors do not mention how much overhead is introduced due to the fact that the mobile device is continuously measuring whether or not to switch between the wireless interfaces. I believe that if this overhead is significant than the whole infrastructure that is CoolSpots becomes irrelevant.

Another discussion that would have been interesting would be to compare how the system reacts if the phone is in Bluetooth mode but also has to deal with another Bluetooth application being used concurrently (i.e. how does the device handle receiving traffic over Bluetooth while also being paired with a bluetooth headset). Lastly, I was curious to know what happens should the mobile device go out of Bluetooth range. Does the base station automatically know to switch to WiFi or will it wait to receive notification from the device that a switch is necessary as is the case when the device is in Bluetooth range? I think that the paper's quality would have been much greater had these issues also been addressed.