

Energy consumption savings with 3G offload

M. Isabel Sánchez^{*†}, Carlos J. Bernardos[†] and Antonio de la Oliva[†], Pablo Serrano[†]

^{*}Institute IMDEA Networks, Spain
Email: mariaisabel.sanchez@imdea.org

[†]Departamento de Ingenieria Telematica, Universidad Carlos III de Madrid, Spain
Email: {cjbc, aoliva, pablo}@it.uc3m.es

Abstract—Current trends on mobile traffic show an exponential grow of the traffic consumed by users from smartphones and other portable devices. The explosion of traffic in cellular networks has forced operators to start deploying solutions to alleviate the congestion on their capacity-limited and expensive radio access networks. One of the solutions being discussed is the so called 3G offload that enables the terminals to use other technologies such as WiFi to offload some of the traffic. IP flow mobility is one mechanism providing 3G offload, by enabling selected flows to be moved among network interfaces. Although this is a very promising technology, it is not clear yet how it will affect the protocols currently in use to provide IP mobility in cellular networks, e.g., Proxy Mobile IPv6. The use of 3G offloading does not only benefits the operators, but also the final user, as it might extend the battery lifetime of its terminal. In this paper we first describe some network-based IP flow mobility extensions, highlighting important design choices. Secondly, we focus on providing experimental measurements showing how the use of this technology can result in an extended battery life for the case of 3G and WiFi enabled terminals.

I. INTRODUCTION AND MOTIVATION

In the recent years we have witnessed an exponential growth in mobile data applications and the subsequent increase in the data traffic volume to be handled by cellular network operators every day. The considerable bandwidth demand from mobile users makes the operators consider additional access networks to alleviate the overloaded cellular access. In addition, the high penetration rate of smartphones is being fostered by high-profile terminals that upgrade their software and hardware capabilities very fast, and therefore, trigger the creation of applications and services customizable to every mobile user. Interworking between 3G and WLAN access networks is a widely addressed research topic, but the availability of mobile terminals, their high adoption and market-penetration rates, the continuous upgrade of their capabilities, and the plethora of applications are renewing the interest of researchers and operators on this topic. The design of an efficient mechanism to have the cellular and the WLAN connections sharing the traffic load benefits both the network operators and the final users. By offloading the traffic from the cellular network to a WLAN, a network operator can reduce the load on its network and reuse the freed resources for users that cannot handoff their traffic. On the other side, offloading and flow mobility in general can also improve the user experience, although usually this part is overlooked since the main driver of it is to alleviate the operator's problems. For example, the mobile users can experience a better quality due to the higher bandwidth that a WLAN can offer compared to 3G, or even use

both interfaces at the same time in order to achieve a higher available bandwidth.

Although 3G offload can be used to just refer to the simple handover of all IP flows from one interface, e.g., 3G to a secondary one, e.g., WiFi, the opportunities that this technology enables are maximized when a fine-grained flow selection is allowed. For example, an operator might prefer not to offload VoIP flows, due to the inherited difficulties in providing QoS guarantees on an unmanaged WiFi access, while video traffic might be always offloaded to a technology providing higher bandwidth.

This paper aims at providing a high level view of the flow mobility mechanisms, specifically for PMIPv6, highlighting the key challenges and design choices encountered while providing this new functionality (Section II). In addition, it also provides insights of an extra-benefit to the end user that has not been previously studied, the increased energy efficiency at the terminal that can be achieved by access technologies such as WiFi that consumes less energy than 3G (Section III). Through several measurements comparing the energy consumed by each interface on its different transmission states, in Section III-B we prove that flow mobility allows a longer battery lifetime. Finally we conclude this work with Section IV.

II. 3G OFFLOAD

Flow mobility targets the offloading of certain flows, selected according to some optimization criteria, from one network connection to another. This change needs to be performed so it minimizes the impact on the quality perceived by the user. In a real scenario, this change would happen between different access technologies, mainly because current smartphones only have one interface per access technology.

Due to the market penetration and the availability of smartphones featuring both WLAN and 3G interfaces, we are focusing on this study case. A very basic approach, consists on performing an inter-technology handover whenever a WLAN connection becomes available, moving all the traffic from the default cellular access. However, the ability in selecting the traffic that is going to be moved by supporting simultaneous 3G and WLAN connections is a much more appealing solution. In this scenario, the operator can choose the kind of traffic to route over the WLAN, offloading the congested 3G network by developing policies for IP flow mobility.

A. Flow mobility for PMIPv6

Proxy Mobile IPv6 [1] is the network-based solution for mobility management proposed by the IETF. It extends Mobile IPv6 defining additional functional entities: *i*) the Mobile Access Gateway (MAG) and *ii*) the Local Mobility Anchor (LMA). The MAG is usually the access router of the mobile node, being responsible for tracking the mobile node in the access link. Therefore, the MAG takes care of the signaling on behalf of the mobile node attached to it. The LMA is the entity at the backbone network that gathers the routing information for the mobile nodes, storing the MAG to which the mobile nodes are connected at every moment. Note that the LMA is also the entity that assigns the IPv6 prefixes to the mobile nodes. After the mobile node is attached to a MAG, the MAG and LMA exchange the signaling messages to update the mobile node's location and then, they set up a bidirectional tunnel between them. From that moment on, all the traffic exchanged by the mobile node will be encapsulated into this tunnel.

The basic requirement for a mobile device to support flow mobility is to have several physical network interfaces. If the mobile node connects to the PMIPv6 domain through multiple interfaces, each of the interfaces will be assigned a unique set of home network prefixes, and each set will be managed under a different mobility session¹. There are three possible scenarios to connect through different interfaces into the same PMIPv6 domain:

- Unique set of prefixes per interface. This is the default mode of operation for PMIPv6. Each attached interface is assigned a different set of prefixes, and the LMA maintains a mobility session (i.e., a binding cache entry) per MN's interface.
- Same prefix but different global addresses per interface. In this case the same prefix is assigned to multiple interfaces, though a different address is configured on each interface. This mode is not completely supported by PMIPv6.
- Shared address across multiple interfaces. In this scenario, the MN is assigned the same IP address across multiple interfaces. This enables applications on the terminal to see and use only one address, and therefore the MN could be able to benefit from transparent mobility of flows between interfaces. This scenario is not supported by current PMIPv6,

As it is currently defined, PMIPv6 does not provide flow mobility support in any of the previous flavors, so we described a set of extensions in [2], mainly affecting the mobility signaling exchanged between LMA and MAG and the data structures maintained by each element. We summarize these extensions in the following section. It is important to note that these extensions enable a mobile device to support flow mobility without introducing any kind of mobility support in the terminal.

1) *Single IP interface case, the logical interface model:* The *logical interface* [3], [4] at the IP layer is the most

¹The term mobility session refers to the creation of a state associated to the mobility binding of a mobile node on the LMA and the MAG.

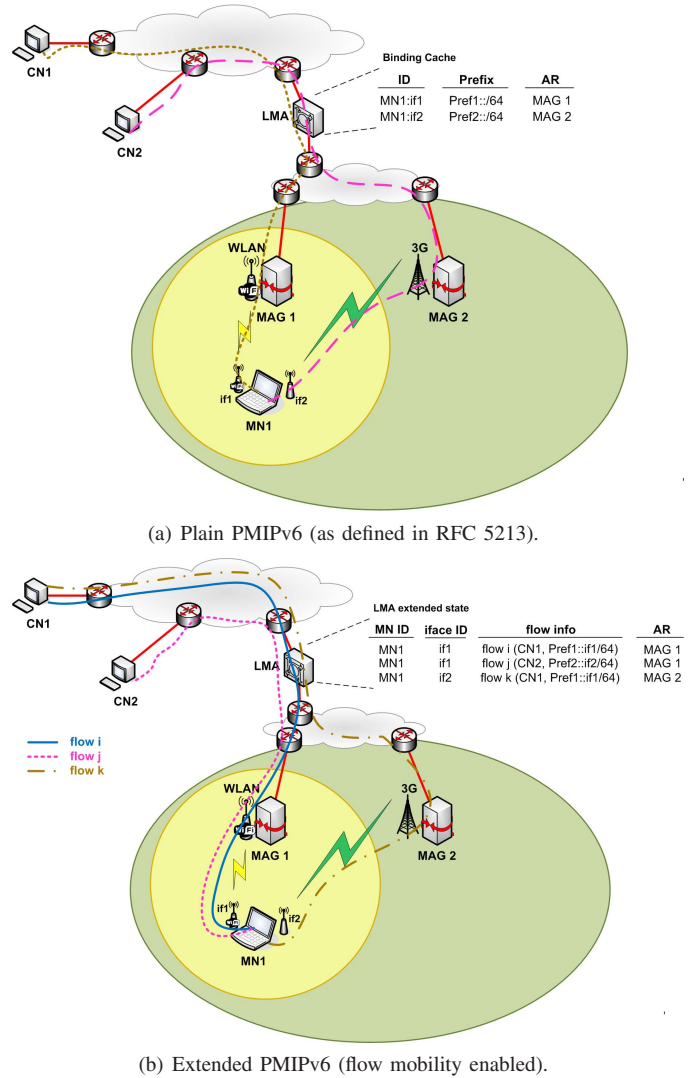


Fig. 1. PMIPv6 extensions for flow mobility support.

complete approach from those link-layer implementations that hide the use of multiple physical interfaces to the IP layer [5]. When a mobile node connects to the same PMIPv6 domain via multiple physical interfaces using the logical interface, it appears to the rest of the network as a set of different endpoints with the same layer 2 and layer 3 addresses. For the LMA, once the mobile node has registered one of its interfaces, the subsequent attachments using different interfaces will be seen as handover requests. Our extension, *i*) extends the original PMIPv6 for the MAG to specify that the mobile node is attaching through a physical interface belonging to a logical interface, and *ii*) modifies the data structure at the LMA so it stores information about every MAG that connects to the same mobile node.

2) *Multiple IP interfaces case, the weak host model:* In case the terminal does not follow the logical interface approach, it is possible to enable flow mobility support if the terminal implements the *weak host* model [6], [7]. This model allows the mobile node to receive and process packets whose IP address corresponds to that one of any of its local interfaces, instead of only accepting the packets addressed to the IP of the

interface receiving the traffic. In our experience, the weak host model is implemented in Linux (tested with Linux-2.6.26) and Mac OS X (Leopard), both for IPv4 and IPv6 traffic. When an MN connects through different interfaces, each of them is treated separately, as if each of them were a different MN. This solution enables the LMA to group together in a single mobility session all the information referring to the same MN.

III. ENERGY CONSUMPTION ASSESSMENT

Enabling flow mobility implies benefits both for the operator, that can save resources in the radio access, and also for the user, that is able to take advantage of an increased available bandwidth. However, in order to fully assess the suitability of this mechanism, it is essential to evaluate it in terms of complexity and of another component usually forgotten, its energy consumption. We focus in this paper on the energy efficiency. Energy consumption is specially critical for mobile devices and smartphones, which already suffer from battery-drain issues due to continuous and exhaustive use along the day. Despite the fact that 3G connection is heavily consuming the battery of the device, it is generally configured to be the default access connection and is almost always on. Therefore, in order to implement a flow mobility solution, it is reasonable to assume that in addition to this intensive usage of the 3G interface, we will need to add the energy consumption corresponding to additional network interfaces. Nevertheless, our experimental results show that the energy consumed by the 3G interface is higher than the one by the WLAN interface, so offloading the 3G connection helps reducing this consumption. Through this section we provide experimental results supporting the claim that flow mobility can also be beneficial for the user in terms of achievable energy consumption savings.

Modern terminals such as Android or iPhone smartphones do not allow by default the simultaneous use of 3G and WiFi interfaces. To overcome this issue and perform an experimental assessment of the energy cost derived from enabling IP flow mobility (i.e., use of multiple network interfaces at the same time) we perform real power consumption measurements on a multi-mode device, equipped with a WLAN IEEE 802.11a/b/g and a 3G UMTS (HSDPA capable) interface. In order to be able to control as much as possible the used devices, capture traffic sent and received at the network interfaces, as well as closely monitor the device, we decided to use a small residential router, Asus WL-500GP v1.0, based on a Linux firmware. The measurements provided through this procedure are later validated by the analysis of the battery lifetime of an android smartphone while using 3G and WLAN interfaces separately. Finally we derive our main conclusions through a synthetic use case that allows us to provide quantitative gains on the percentage of battery spent through the use of the proposed flow mobility mechanism.

A. Assessing the power consumption of the joint operation of 3G and WiFi

The following section is devoted to perform the experimental assessment of the energy consumption associated to our flow mobility solution. To measure the power consumption of each technology we have chosen a small residential router: the Asus WL-500g Premium. This router is equipped with a 266

TABLE I. POWER CONSUMPTION RESULTS

3G ON		WLAN ON	
WLAN OFF	1.80 ± 0.10 W	3G OFF	1.03 ± 0.08 W
WLAN IDLE	1.86 ± 0.08 W	3G IDLE	1.21 ± 0.16 W
WLAN ON	2.16 ± 0.13 W	3G ON	2.16 ± 0.13 W

MHz processor, an IEEE 802.11b/g WLAN interface and an IEEE 802.3 Ethernet interface connected to a VLAN capable 5-port switch. This version of the router has a mini-PCI slot that allows changing the original wireless card. We replaced the original Broadcom card an Atheros based 802.11a/b/g (Alfa Networks AWPCI085S) one, which is supported by the Madwifi² driver. In order to mitigate as much as possible the impact of collisions and interference in the power consumption measurements, we avoided the 2.4GHz band (IEEE 802.11b/g) – which is very crowded in our lab, as reported in [8] – and configured the WLAN interface in 802.11a mode.

We replaced the original firmware of the router by installing a lightweight Linux-based version, which gives us more flexibility in the configuration. We choose the distribution Kamikaze 8.09.2 of OpenWRT³ with a Linux-2.6 kernel and this allows the support of a 3G USB modem. In our tests, we used a Huawei E160 HSDPA USB stick⁴.

Power consumption was measured using a PCE-PA 6000 power analyzer⁵. Power measurements were carried out using a PCE-PA-ADP current adapter where the power supply of the router was plugged in. Measurement data was transferred from the power analyzer to a computer via an RS-232 interface for its processing.

Using this setup we performed the measurements described next. We first calibrated the power analyzer by measuring the consumption when both the WLAN and 3G interfaces are switched off. All reported results are relative to this level. For the actual measurements, we are interested in the power consumption when the network interfaces are in the following states:

- OFF: the interface is switched off.
- IDLE: the interface is on but it does not send or receive any data traffic. For the case of WLAN, this means that the card is associated to an access point (so the card is receiving beacon frames) without sending or receiving any user data traffic. For the case of 3G, this means that the interface is up, a PDP context has been activated and a PPP interface has been set up, but no data is exchanged.
- ON: the interface is on and engaged in a data traffic exchange. In our tests, this means that a file is downloaded from a server using HTTP. By using TCP, the card is receiving at the maximum available rate, and traffic is sent in both directions (downlink: mostly data segments, uplink: mostly TCP acknowledgments).

We measured the power consumption for the different states of the WLAN and 3G interfaces. Table I shows the

²<http://www.madwifi.org/>

³<http://www.openwrt.org/>

⁴<http://www.huawei.com/mobileweb/en/products/view.do?id=1960>

⁵<http://www.industrial-needs.com/technical-data/power-analyser-PCE-PA-6000.htm>

mean and 95% confidence interval of the results extracted from five 300-second experiments. We focus on the scenarios where at least one of the interfaces is actively sending or receiving traffic, as those are the cases in which it is important to evaluate the energy cost associated with having a second active interface. This second interface may be either receiving or sending traffic or just idle, ready to operate. Results show that the 3G interface consumes more energy than the WLAN interface, but the difference between using only the 3G interface and using simultaneously the 3G and the WLAN interfaces is only of 20%. Note that this additional cost is only incurred when both interfaces are actively engaged in a data transfer, and that by using them simultaneously, the time required to send a given amount of data via WLAN would be shorter – since the throughput obtained via a WLAN network is typically higher than the one that can be obtained via a 3G network – and this would also contribute to a lower power consumption. The extra power consumption caused by activating the WLAN interface (IDLE state) is just around 3%, which besides would only be needed when the mobile is sending or receiving traffic, as it is then when the network operator and the user may benefit from offloading traffic from the 3G infrastructure to a WLAN hotspot, if available. It is important to highlight that the actual values of energy consumption of the device are not directly comparable with the results we would obtain with a smartphone, since the level of integration provided in such a platform is much higher, allowing further improvements in the energy consumed by the device. Due to this, and to be able to compare, we only focus on the relative difference between the 3G and WiFi consumption profiles which, as we will see in the next section, follow the same trend in both device families.

B. Energy consumption profiles of 3G and WiFi in an Android device

In order to confirm the results obtained in section III-A, we measured the battery duration of an HTC Legend device. The operating system of the device under test is Android 2.1 (Eclair). Apart from the device, we used a desktop to monitor and configure the parameters of interest in the mobile terminal. We also configured the WLAN interface of this desktop as Access Point to which the mobile terminal would associate. To measure the energy consumption we developed an application running in the background to monitor the battery continuously, keeping track of the voltage level in order to compute the power consumed. As this application is a service running in the background, it consumes negligible CPU resources, minimizing the impact on the energy consumption measurements. In addition, no interaction with the user (or the tester in this case) is required, as all the information is saved to a text file. All the measurements have been performed keeping one element active and the rest inactive, in order to isolate the contribution of each individual element to the total power consumption of the device.

Figs. 2 and 3 show the battery drainage curves for WiFi and 3G respectively. Fig. 2 presents results for each of the possible states of the WiFi interface, as explained in section III-A, while Fig. 3 considers only the states we can control on the 3G interface, namely: transmission, reception and disconnection states. The results from these measurements show that the battery life of the device is much shorter when the 3G interface is on than when the WLAN interface is at its maximum battery

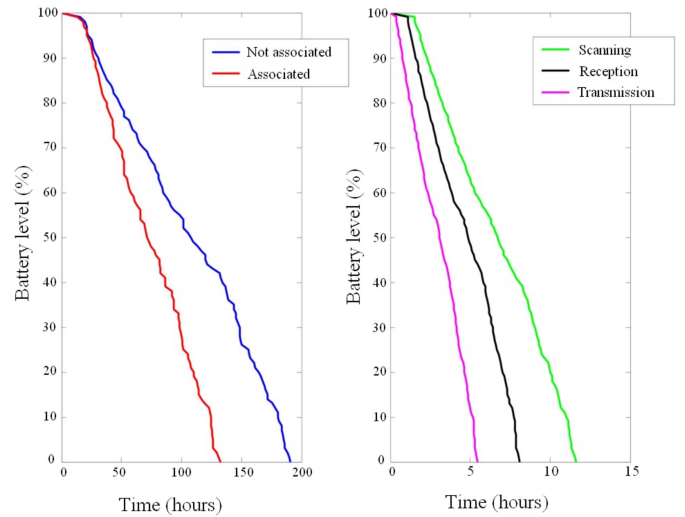


Fig. 2. Battery drainage for the different WiFi states.

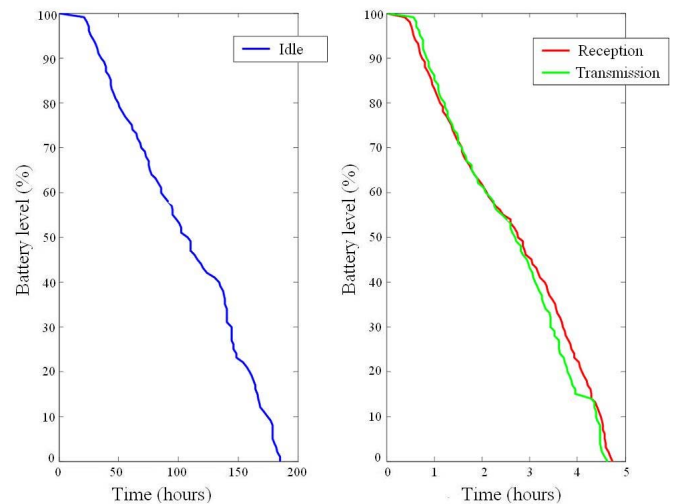


Fig. 3. Battery drainage for the different 3G states.

consumption task, which is transmitting packets. The battery of the mobile terminal can last 200 hours when the 3G interface is inactive, oppositely to the less than 5-hour duration when there is incoming or outgoing traffic. However, in the case of the WLAN interface, the difference between the transmission and reception states are much more evident than those of the 3G interface, shrinking from 8 hours to 5 hours, respectively.

In addition, through the analysis of the slopes of each curve, we can obtain the relative difference between the cases of a single active interface (3G) and the case when both interfaces (3G+WiFi) are simultaneously used. Supposing a transmission and reception cycle of a 50%, the overall difference between both cases is approximately 15%. As expected, the relative difference between both cases is lower for the smartphone device, we argue that this difference is due to a higher integration of the components in the smartphone compared to the router.

Finally, and in order to conclude this analysis, let us make a synthetic example of the energy saving that such an approach

would provide to the end user. First let us consider some assumptions, for the sake of the simplicity of this analysis, which aims at assessing if a typical mobile user could afford the additional power consumption introduced by the use of flow mobility extensions. Several studies, such as [9], point out that users of smart hand-held devices download an average of 20 MBytes per day via 3G. Considering Fig. 3, and an average 3G speed of 1 Mbps, the download would take 160 seconds and consume around an 0.8% of the battery.⁶ In case a flow mobility solution was deployed and the terminal was able to use WiFi to download the same amount of information, it would use the WiFi interface for approximately 6.4 seconds (assuming IEEE 802.11a extended rates and a real throughput of approximately 25 Mbps). During this time, the terminal will use a 15-20% more energy compared to the case of using only 3G, but the overall time would be highly reduced. This implies that the terminal would have spent less than a 0.1% of the battery downloading the file. This simple analysis does not aim at providing rigorous and precise figures, but just at roughly assessing if a flow mobility solution is affordable from the perspective of power consumption. Based on the obtained results, we can conclude that selectively using more than one network interface results in an affordable additional cost.

From these experimental results we can derive that the use of the WLAN interface is considerably more efficient in terms of energy consumption than the use of the cellular 3G connection. In addition, the throughput and the achievable bandwidth by using a WLAN access are also higher than the ones that the 3G connection can offer. Therefore, we can take advantage from the higher bandwidth offered by the IEEE 802.11 access network, offloading the cellular connection and freeing resources for other users while reducing the energy consumption of our devices.

IV. CONCLUSION AND FUTURE WORK

In this work we have focused on the IP flow mobility and we analyzed it from the perspective of energy efficiency. First of all, we motivate the need for flow mobility and we provide a flow mobility solution built upon PMIPv6, defining the extensions needed by the protocol to support flow mobility, as the current definition does not have this capability. We have identified the advantages that flow mobility can bring both for the network operator and for the end user, and in order to argue the energy consumption we present some experimental results on commercial devices. These results enable us to claim that the flow mobility solution provided is also affordable in terms of battery consumption, which is a key element under study in the research community.

Networking research is evolving towards a greener framework, analyzing the causes of battery draining and searching for optimizations or new solutions that allow to diminish the power consumption of networking protocols and communication devices. Our next steps intend to evaluate several of these commercial devices, such as an Android smartphone, an iPhone or a Windows Phone, so as to evaluate the performance in terms of energy efficiency of real devices available in the market and for different operating systems and architectures.

ACKNOWLEDGMENTS

The authors would like to thank Fabio Giust, Andres Garcia Saavedra and Elena Lopez Orgaz for their valuable contributions to the final results presented in this work. The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7-ICT-2011-8) under grant agreement n. FP7-ICT-318115 (CROWD project). This work has also been supported by the Spanish Government, MICINN, under research grant TIN2010-20136-C03.

REFERENCES

- [1] S. Gundavelli, K. Leung, V. Devarapalli, K. Chowdhury, and B. Patil, "Proxy Mobile IPv6," RFC 5213, August 2008.
- [2] T. Melia, C. Bernardos, A. Oliva, F. Giust, and M. Calderon, "IP Flow Mobility in PMIPv6 Based Networks: Solution Design and Experimental Evaluation," *Wireless Personal Communications*, vol. 61, no. 4, pp. 603–627, 2011. [Online]. Available: <http://dx.doi.org/10.1007/s11277-011-0423-3>
- [3] R. Wakikawa, S. Kiriya, and S. Gundavelli, "The applicability of virtual interface for inter-technology handoffs in Proxy Mobile IPv6," *Wireless Communications and Mobile Computing*, 2009.
- [4] H. Yokota, S. Gundavelli, T. Trung, Y. Hong, and K. Leung, "Virtual Interface Support for IP Hosts," Internet Engineering Task Force, draft-yokota-netlmm-pmipv6-mn-itho-support-03.txt (work-in-progress), March 2010.
- [5] S. Gundavelli and T. Melia, "Logical Interface Support for multi-mode IP Hosts," Internet Engineering Task Force, draft-ietf-netext-logical-interface-support-01.txt (work-in-progress), October 2010.
- [6] R. Braden, "Requirements for Internet Hosts - Communication Layers," Internet Engineering Task Force, RFC 1122, October 1989.
- [7] D. Thaler, "Evolution of the IP Model," Internet Engineering Task Force, draft-thaler-ip-model-evolution-01.txt (work-in-progress), July 2008.
- [8] P. Serrano, C. J. Bernardos, A. de la Oliva, A. Banchs, I. Soto, and M. Zink, "FloorNet: Deployment and Evaluation of a Multihop Wireless 802.11 Testbed," *EURASIP Journal on Wireless Communications and Networking*, vol. 2010, 2010.
- [9] P. Rysavy, "Mobile Broadband Capacity Constraints and the Need for Optimization," February 2010.

⁶This value matches perfectly with a real measurement of the power consumed by an iPhone 3GS downloading a 20 MBytes file.