Towards an Energy-Aware Network Activation Strategy for Multi-Homed Mobile Devices

M.S. Bargh and A.J.H. Peddemors

Telematica Instituut,
P.O. Box 589,
7500AN Enschede, the Netherlands;
Tel.: +31 53 4850{443, 421}; Fax: +31 53 4850400;
E-mails: {Mortaza.Bargh, Arjan.Peddemors}@telin.nl

Abstract— Simultaneous operation of multiple network interfaces can benefit mobile devices with diverse network interfaces. This paper outlines an IP-level network activation strategy that aims at optimum use of the complementary features of the underlying networks. A key design issue in such strategies is the energy cost of different network interfaces. As a step towards validating and fine-tuning our strategy, we devised an experimental setting to measure the relative energy consumption of WiFi and GPRS interfaces in two modern mobile devices. The paper reports the experimental method and results. Our study indicates that the proposed strategy –in this case, reaching devices from the GPRS interface and, whenever appropriate, handing over some sessions to the WiFi interface– is plausible from energy consumption viewpoint and in some cases it outperforms the existing methods. The energy consumption model obtained will be used to make the network activation strategy energy aware.

Key words—Mobile devices, multiple network interfaces, connectivity management, energy cost, WiFi-GPRS.

1 Introduction

BEYOND 3G mobile systems aim at seamless integration of heterogeneous wireless networks with wire-line backbone networks. The range of wireless networking technologies includes Wireless Wide Area Networks (WWANs) such as GPRS and UMTS cellular networks, and Wireless Local Area Networks (WLANs) such as Wi-Fi, HiperLAN/2 and HomeRF networks. As the main element of this seamless integration, mobile devices with multiple network interfaces emerge in the market at a fairly fast pace. Such devices are able to connect to the Internet via multiple wireless network interfaces sequentially or simultaneously. Internet access, i.e., having IP level connectivity, via multiple interfaces simultaneously is called “end-host multi-homing” [1], or “multi-homing” in short. Motivations for multi-homing are: having smoother hand-offs, offering better QoS (due to accessing multiple networks), exploiting access networks with link asymmetry (e.g., satellite systems), reducing cost, and improving privacy and security [2]. A main issue of multi homing is to optimize the way network interface resources are used in order for such devices to be reachable by others, and to occasionally switch the path of ongoing sessions from one network interface to another, i.e., to do a vertical handover.

This paper investigates the energy consumption characteristics of different networks interfaces in multi-homed mobile devices as a first step towards optimizing the network activation process. Motivated by ongoing trends, we consider mobile handheld devices with GPRS and Wi-Fi interfaces, and evaluate the energy efficiency of these network interfaces from the perspective of mobility management functionality. Our measurements for two commercial mobile devices show that the GPRS network interface is energy efficient for reaching these devices and the WiFi network interface is energy efficient for sending/receiving large volumes of IP data. The study, moreover, provides insight into the energy trade-offs for these interfaces in different operational modes. The obtained results also validate the energy efficiently of a network activation strategy proposed in our previous work [3] for pervasive communication. This strategy exploits the WiFi and GPRS network characteristics –for example, better link-layer mobility management, lower bandwidth per unit area and larger coverage area of GPRS networks compared to WiFi networks– in order to facilitate, for instance, IP-level mobility management and vertical handovers.

The rest of this paper first lays down the problem context and provides some background information in Section 2. Section 3 describes our network activation strategy as the main motivation behind the study reported in this paper. Section 4 explains our approach for energy consumption measurements and presents the results for a commercial device. Section 5 summarizes the measurement results and analyzes them for two devices. Finally, Section 6 presents some conclusions and future work.

2 Selecting Network Interfaces

This section reviews a number of solutions proposed to deal with activating multiple network interfaces in an Always Best

This work is part of the Freeband AWARENESS project (http://awareness.freeband.nl). Freeband is sponsored by the Dutch government under contract BSIK 03025.

Connected (ABC) way, i.e., being connected in best possible way [4].

2.1 Activating Interfaces Sequentially

The ABC vision strives for connecting the user with the best possible QoS at any time and place. Network selection algorithms proposed for vertical handover take various (QoS) parameters as input to decide over the best network. Such parameters for decision making include radio signal strength, mobile device speed, network traffic load and application QoS requirements for throughput and handover delay, for the references see [5]. The authors of [5] consider network condition, application QoS requirements, and user preferences as the decision factors. The network condition is assessed based on network QoS parameters of throughput, timeliness (i.e., delay, response time, jitter), reliability (i.e., bit error rate, burst error, average retransmission), security, and cost. The strategy of [5] assumes that the UMTS network is always on and available. Upon discovering a WLAN, the decision process selects one of the WLAN or UMTS networks for connectivity based the criteria mentioned.

All approaches mentioned actually choose for one network interface at a time and opt for a sequential operation mode for interfaces. Such sequential activation of network interfaces saves battery energy by avoiding two or more network interfaces being on simultaneously. On the other hand, however, these algorithms do not take into account the relative energy cost of these network interfaces in the decision process. Moreover, this sequential mode of operation does not allow finer distribution of different (application) flows between network interfaces as suggested in [6].

2.2 Activating Interfaces Simultaneously

Per flow selection of network interfaces requires simultaneous activation of two or more interfaces. Such simultaneous operation is the core property required for multi-homing in order to fully exploit the capabilities of heterogeneous networks. Such network selection algorithms consider application QoS requirements, the mobility management features [7][3], and/or the reliable performance and reliability guarantees of the underlying networks [6][3], to map different data/signaling flows onto network interfaces. In the following we review one of these approaches since it is akin to the topic of this paper.

Generally the IP-level mobility solutions proposed, e.g., Mobile IP, do not deal with issues such as access network detection, access network selection, and network interface configuration that are prerequisite to mobility management. However, the approach of the MIRAI project [7] convolves access network selection with mobility management signaling. Aiming at seamless integration of heterogeneous wireless systems, the MIRAI project assumes the simultaneous operation of network interfaces in multi-homed mobile hosts. Specifically, the MIRAI architecture introduces a Basic Access Network (BAN) to carry out signaling for network discovery, network selection, inter-network handover, location updates, paging, authentication, authorization, and accounting. The BAN should be reliable and have a broad coverage area, preferably larger than that of the other networks. The MIRAI researchers designed a special network as the BAN in their earlier work [8] and recently in [7] they also enabled selecting a GSM network as the BAN. Such a static selection of access networks for mobility management signaling overlaps partially with our approach in Section 3. Note that they did not carry out energy consumption measurements of the GSM network interface when used as the BAN.

3 A Network Activation Strategy

This section provides an overview of our IP-level network activation strategy together with a high-level system view. The previous work in [3] describes a similar IP-level mobility management strategy in detail and analyzes its grounds from the perspective of the mobility management requirements of different applications. In this section we position it as a strategy governing the IP level network activation in mobile devices with multiple interfaces. This underscores the tight binding that basically exists between the mobility management and network selection functionalities in such mobile devices.

Assume a mobile device with two generic types of air interfaces: a WWAN interface, which can connect to, e.g., GPRS or UMTS cellular networks, and a WLAN interface, which can connect to, e.g., Wi-Fi, HiperLAN/2 or HomeRF networks. Figure 1 illustrates the state diagram governing the operation of the IP-level network activation strategy with transition edges marked by the Event (E) causing the state change.

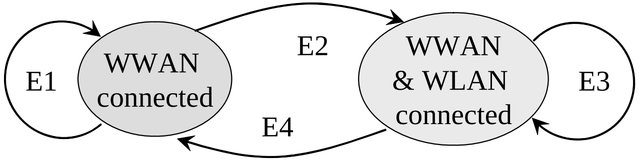

Figure 1: the state diagram of the strategy.

The strategy assumes that the mobile device operates in the default state of “WWAN connected”. In this state, for example, all IP-level location management signaling is carried out via the WWAN interface. This location management signaling looks like the way proposed in the latest variant of the MIRAI in [7] with BAN being cellular GSM (i.e., the WWAN). After reaching a mobile node via its WWAN interface, a session can be initiated via the WWAN interface (corresponding to edge E1) or the session can be handed over to the WLAN interface (corresponding to edge E2). The event E2 leads to a vertical handover of the session and depends on the context parameters like availability of a hotspot, requirements of the (ongoing) session, cost, remaining energy level, etc. The handover management signaling, e.g., informing the correspondent hosts over the mobile node’s new IP address in order to change the route of ongoing sessions, is carried out via the WLAN interface. In this respect our approach differs from the MIRAI one that requires also handover signaling to go through the BAN interface). Event E3 marks having at least one session that communicates via the WLAN interface and event E4 denotes deactivating the WLAN interfaces based on the context, e.g., unavailability of the hotspot, termination of all sessions that use the WLAN interface, cost, remaining energy level, etc. Basically all events E1-E4 may depend on the context in which the mobile device operates. This context dependency makes the network activation strategy to be application-aware (e.g., the QoS requirements of running sessions), infrastructure-aware (e.g., which network has best mobility management features), movement-aware (e.g., at which speed to handover sessions), etc.

The network activation strategy described attempts to make optimal use of the features embedded in the existing technologies. Particularly, it realizes IP-level mobility management without requiring any IP-level paging capability because it relies on the sophisticated and energy efficient link layer mobility features of the cellular WWANs. Thus, the IP level Location Management, namely IP-address update and paging signaling, becomes lightweight due to reliance on L2 WWAN location management mechanisms.

A partial implementation of the strategy described is reported in [9], which switches a multimedia streaming session between the WiFi and GPRS interfaces using a proprietary signaling protocol. In near future we intend to extend the current implementation and fully realize the strategy mentioned. To this end, a special service for mobile devices is developed (see [10]) that reflects the availability and capacity of access networks and offers various protocols such as the Mobile IP, SIP or mSCTP to realize the strategy. Unlike the approach of the MIRAI project, our design does not require any change to the standardized Mobile IP protocol.

4 Measurements

The strategy presented in Section 3 assumes the simultaneous operation mode for network interfaces on some occasions. In this mode the mobile device consumes more battery energy than when the network interfaces are activated sequentially. The study in this paper mainly aims at validating the feasibility of the proposed strategy and gaining insight on how the strategy can be fine-tuned from the energy efficiency viewpoint. The results of these experiments can moreover be used to evaluate any network activation strategy that adopts simultaneous operation mode, i.e., all those used in multi-homed devices. This section describes the method user for measuring the energy consumption, explains the tools developed, and presents the graphs obtained for a commercially available mobile device.

4.1 Experimental Tools

The experiments were carried out on two commercially available mobile handheld devices: Qtek 9090 and iPAQ 6515. Both devices run Microsoft Windows Mobile™ 2003-second edition Operating System (OS) and have WiFi (IEEE802.11b) and GPRS network interfaces. The iPAQ device uses an SDIO (input/output) card for the IEEE802.11b interface.

To measure the energy consumption we used the GetSystemPowerStatusEx function of Microsoft Windows Mobile to retrieve the power status of the system [11]. The power status is the form of the SYSTEM_POWER_STATUS_EX structure, from which we polled and recorded the BatteryLifePercent member in every 60 seconds during every experiment. The BatteryLifePercent member can be a value in the range 0 to 100 indicating the percentage of full battery charge remaining [11].

Our measurements aimed at capturing the energy-draining rate during various device operation modes that concern the evaluation of our network activation strategy. During each operation mode, we measured the normalized energy remaining, i.e., the remaining energy in percentage, as explained above. We determined these operation modes based on whether the WiFi or GPRS network interface had: no connectivity (i.e., being off), idle-IP connectivity or active-IP connectivity. The no connectivity denotes the case where the mobile device is not connected to a given network interface. The idle-IP connectivity denotes the case where the interface has IP connectivity, i.e., configured with an IP address, but does not send/receive application IP packets. The active-IP state denotes the case where the interface sends or receives TCP/IP packets continuously. We used the Netperf tool [12] with a client running on the mobile device to send dummy TCP packets over a given network interface. All these three connectivity cases are defined from the viewpoint of applications running on top of the OS. The exact physical layer state in which an (WiFi or GPRS) interface operated in each of these high-level states was not relevant in our study.

In case of the Qtek device, the Windows Mobile OS does not allow to turn on and off the WiFi and GPRS network interfaces at the IP level independently. Moreover, for both Qtek and IPAQ devices the OS does not allow to willingly direct IP packets via either of these interfaces if ever both are activated at the IP level simultaneously. We developed special tools to realize independent operation of WiFi and GPRS interfaces in the Qtek device, for detail see [13].

4.2 Experimental Setting

The experiments for various network interface connectivity cases were carried out in two generic modes: the screen backlight being on or off. The motivation was to observe the backlight energy cost relative to transmission energy cost and, moreover, to validate our experimentation method by comparing the corresponding values (see Subsection 4.3.1).

We ran the experiments until either the remaining energy fell below 25% or a period of 6 hours elapsed. During each experiment we kept the operation of the mobile device steady, for example, by not displacing the device or initiating any other application. The objective behind running the experiments for such a long duration and in a steady mode was to minimize and to cancel out the fluctuations of energy consumption due to short-lived and impulsive terminal activities. In other words, we performed an integration operation on the sporadic energy variations stemmed from periodic terminal activities like location updates. We observed the overall energy consumption behavior of the device for a given operation mode, i.e., when the device network interfaces have no, idle-IP, and active-IP connectivity. Note that our measurements do not include averaging over different devices or over operation instances. For the Qtek Figure 2 shows the graphs of three typical measurements versus time in seconds.

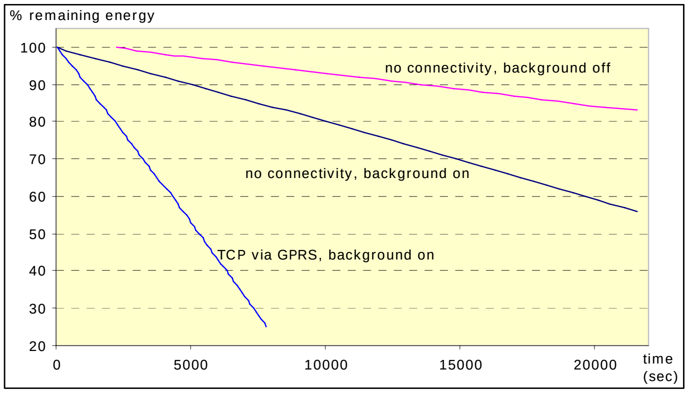

Figure 2: graphs of three typical experiments.

As seen from the graphs the energy consumption graphs behave linearly with respect to time. Therefore, we model the energy consumption in a particular operation mode as $E_{rem}^{n}(t)=100-P_{av}^{n}\cdot t$, where $E_{rem}^{n}(t)$ is the normalized energy remaining at time t and P_{av}^{n} is the normalized average power of the battery, i.e., the slope of the line. The normalization facture is $0.01E_{full}$ for E_{rem}^{n} and P_{av}^{n}, where E_{full} is the energy level when the battery is fully charged. Measuring the normalized energy is sufficient for the purpose of examining our strategy because we are interested in investigating the relative energy cost of WiFi versus GPRS interfaces in a particular device.

4.3 Qtek Results

This section reviews the energy usage graphs in different operation modes experimented for the Qtek 9090 device.

4.3.1 Backlight Effect

To examine the effect of the backlight on energy consumption we carried out all our experiments in two generic modes: with backlight being on and off. While most applications require the backlight being on for the purpose of user-device interaction, there exist applications that run in the background and do not require having the backlight on (e.g., healthcare monitoring applications). This motivated us to investigate the energy cost of the screen backlight and, moreover, to validate our measurement method.

Figure 3 presents the screen backlight effect on energy usage when the device isn't in IP-active state. Figure 4 shows the same when the device is in IP-active state. The no connectivity graphs in Figure 3 indicate the case of both WiFi and GPRS interfaces having no connectivity. In experiments the Netperf program delivered upstream TCP transmission rates of almost 2 Mbs and 25 Kbs for WiFi and GPRS interfaces, respectively. The ratios of the transmission rate to the reception rate were about 50 and 25 for WiFi and GPRS interfaces, respectively.

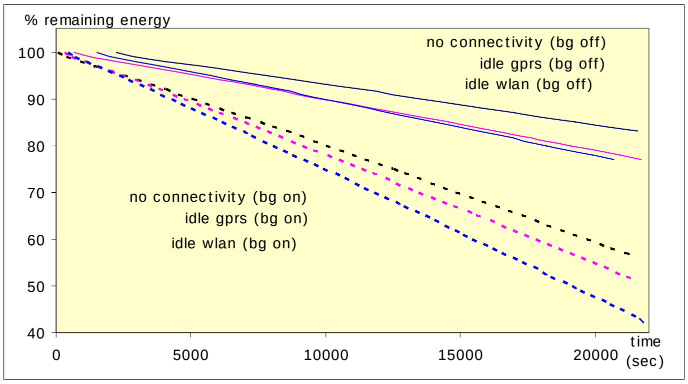

Figure 3: energy consumption of background (bg) light.

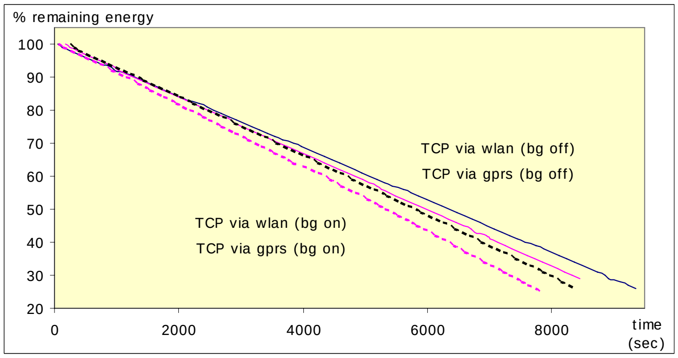

Figure 4: energy consumption of background (bg) light.

Table 1 presents the normalized average power consumed by the backlight in different operation modes. These values are obtained by subtracting the slope of the backlight-off-graph from the slope of the corresponding backlight-on-graph for the graphs shown in Figure 3 and Figure 4. The results in Table 1 show the consistency of our method in differentiating the relative energy cost inflicted by the screen backlight. From now on, we only consider the screen backlight being on case to capture the aspects relevant to our network activation strategy.

Table 1: backlight power consumption.

Operation modes	Normalized average power (i.e., P_{av}^{n}) used due to backlight
No connectivity	11.76 × 10-4
Idle GPRS	12.19 × 10-4
Idle WiFi	12.28 × 10-4
Idle WiFi & GPRS	12.05 × 10-4
TCP via WiFi	11.34 × 10-4
TCP via GPRS	11.31 × 10-4
TCP via WiFi with idle GPRS	10.89 ×10-4

4.3.2 *Idle-IP Connectivity*

In this paper we define the idle-IP connectivity to a network as the case in which a network interface is (a) turned on and (b) configured with an IP address. Basically a device is potentially able to send and receive IP packets through an interface having idle-IP connectivity. The notion of idle-IP connectivity enables us to abstract from the power saving mechanisms deployed in, for instance, Data-Link Layer and Physical Layer (e.g., the power saving mode of the IEEE802.11 standard and other solutions [14][15][16] that exploit the difference between the amount of energy consumed in the different operation modes of WiFi network interfaces). In our view, all such solutions may be deployed in a device when having idle-IP connectivity. In other words, we aim at investigating the energy consumption behavior of existing mobile devices from the perspective of running applications, and any interlayer power optimization is out of our scope.

Figure 5 shows the result of our experiments for the case of idle IP connectivity when backlight being on. The figure shows the energy consumption graphs: the GPRS interface has idle-IP connectivity, the WiFi interface has idle-IP connectivity, and both WiFi and GPRS network interfaces have idle IP connectivity. As a base of comparison, the figure also depicts the no connectivity graph, i.e., when both interfaces are turned off. Note that the Windows Mobile OS did not allow us to have IP level connectivity for both network interfaces simultaneously. For this end we modified the setting of the OS and implemented special tools to activate GPRS and WiFi interfaces at the IP level independently, see Section 4.1.

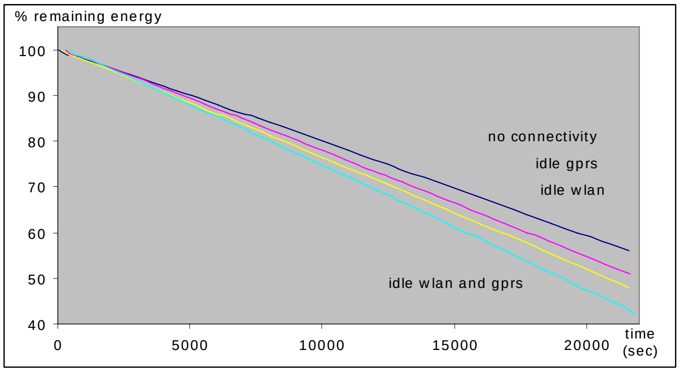

Figure 5: idle connectivity in three modes and no connectivity (the top one), all with background light on.

Table 2 summarizes the values of the normalized power for the graphs shown in Figure 5. These values are relative to the normalized power of the no connectivity case, i.e., the difference between the slopes of the corresponding graphs. The normalized power values are roughly the average values over the period of each experiment. The main conclusions to be drawn from these experiments are:

- During idle-IP connectivity the GPRS network interface consumes less or comparable amount of energy relative to the WiFi one.
- In order to be reachable at the IP level, a mobile device just needs to have idle IP connectivity. Thus, the GPRS network interface is relatively more energy efficient for reaching the mobile device at the IP level.

Table 2: Normalized power for idle IP connectivity.

Operation modes	relative to no connectivity, i.e., $P^{n}_{av,\mathrm{mod}e} - P^{n}_{av,no-conn}$
Idle GPRS	2.56 × 10-4
Idle WiFi	3.82 × 10-4
Idle WiFi & GPRS	6.52 × 10-4

4.3.3 *Active-IP Connectivity*

A network interface has active-IP connectivity to a network when it is sending or receiving IP packets continuously. For example, a multimedia application running on the mobile device receives stream of IP packets from another node in the Internet. The definition presented above does not distinguish between sending and receiving IP packets intentionally. This is because there is a negligible difference in the consumed energy between sending and receiving data in typical network interface cards compared to that between sending/receiving and being idle, see e.g. [17].

Figure 6 shows the result of our experiments for the case of active-IP connectivity by using Netperf program and when backlight being on. The figure shows the graphs of energy consumption when the WiFi interface has active IP connectivity, the GPRS interface has active IP connectivity, and the WiFi interface has active IP connectivity while the GPRS interface having idle-IP connectivity. Note that the latter graph represent the case the mobile is reachable from GPRS interface while a session is going on through the WiFi interface. As a base of comparison, the figure also shows the no connectivity case, i.e., when both interfaces are turned off.

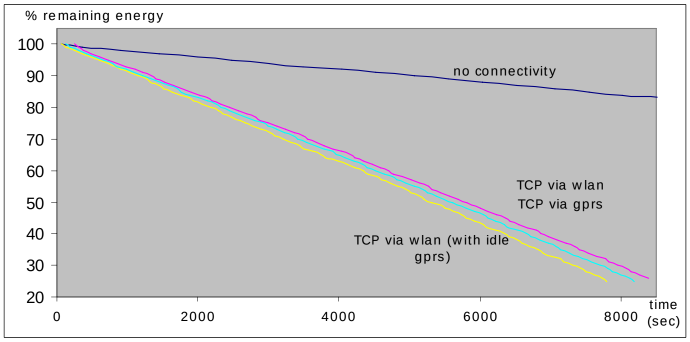

Figure 6: active connectivity in three modes and no connectivity (the top one), all with background light on.

Table 3 summarizes the values of the normalized power for the graphs shown in Figure 6. These values are relative to the normalized power of the no connectivity case, i.e., the difference between the slopes of the corresponding graphs. The conclusions we draw are:

- In a given period of time, the WiFi interface consumes less or comparable amount of energy for sending TCP traffic relative to the GPRS one.
- Considering the fact the WiFi and GPRS interfaces send 2 Mbps and 25 Kbps TCP data, respectively, the WiFi interface consumes much less energy per bit than the GPRS one (almost two order of magnitude).
- In order to send/receive large volumes of IP packets, it is energy efficient to route such session via (or to hand over them to) the WiFi interface.

Table 3: Normalized power for active IP connectivity.

Operation modes	relative to no connectivity case, i.e., $P^{n}_{av,\mathrm{mod}e} - P^{n}_{av,no-conn}$
TCP via WiFi	69.77 × 10-4
TCP via GPRS	76.53 × 10-4
TCP via WiFi with idle GPRS	72.19 × 10-4

4.3.4 *Energy Cost of WiFi Switching*

The strategy proposed in this paper requires that the WiFi interface to be turned on, i.e., to have idle-IP connectivity, and be turned off, i.e., to have no connectivity, whenever appropriate. In order to examine the amount of energy consumed by such operation we devised a software tool to turn on the WiFi interface and immediately turn it off. We repeated this operation 3000 times and monitored the rate of energy decrease in regular intervals. The normalized average power for this operation, i.e., $P^{n}_{av,on-off}$, became 53.74 × 10-4, which is comparable to that of the TCP via the WiFi interface case, shown in Table 3. The average time needed to turn on and off the WiFi interface became $t_{av,on-off} =$ 3.83 seconds. Note that in this experiment the WiFi network was configured for the Open System Authentication mode, i.e., no authentication was needed at all. The expected normalized energy needed for every activation-deactivation of the WiFi interface can be obtained from $t_{av,on-off} \cdot P^{n}_{av,on-off}$. This value should be added to the total normalized energy cost of the WiFi interface in any network activation strategy.

5 Summary of Results and Discussions

This section summarizes the values of the normalized power for different operations modes of two Qtek and iPAQ devices in Table 4 and Table 5, respectively (when the backlight is on). These values are relative to the normalized power of the no connectivity case for each device, i.e., the difference between the slopes of the corresponding graphs. Table 4 integrates the values shown in Table 2 and Table 3. The first impression from the results of our measurements for the iPAQ device is that the battery capacity is less than that of the Qtek device (so for the same operation iPAQ battery discharges faster). Another observation is that the ratio of WiFi average power to GPRS average power in idle-IP connectivity is an order of magnitude larger for the iPAQ device compared to that ratio for the Qtek device. This significant difference can be associated with the fact that Qtek has an integrated WiFi interface card while iPAQ uses a removable SDIO one.

Table 4: Normalized power for idle IP & active IP (Qtek).

Operation modes	relative to no connectivity case, i.e. $P^{n}_{av,\mathrm{mod}e} - P^{n}_{av,no-conn}$	
	Idle-IP	Active-IP
WiFi is …	3.82 × 10-4	69.77 × 10-4
GPRS is …	2.56 × 10-4	76.53 × 10-4
WiFi is … while GPRS idle-IP	6.52 × 10-4	72.19 × 10-4

Table 5: Normalized power for idle IP & active IP (iPAQ).

Operation modes	relative to no connectivity case, i.e. $P^{n}_{av,\mathrm{mod}e} - P^{n}_{av,no-conn}$	
	Idle-IP	Active-IP
WiFi is …	39.70 × 10-4	62.94 × 10-4
GPRS is …	1.34 × 10-4	52.50 × 10-4
WiFi is … while GPRS[1] idle-IP	40.83 × 10-4	65.40 × 10-4

The results shown in Table 4 and Table 5 indicate that, during idle-IP connectivity, the GPRS interface is more energy efficient than the WiFi one (for the Qtek the difference is small). Thus in long periods of being idle and being IP-level reachable it is better to switch on the GPRS interface and allow the WiFi interface to have no connectivity. However during active-IP connectivity, the WiFi interface is more energy efficient relative to the GPRS one. Thus when having a large volume of IP traffic to transfer, it is efficient to activate the WiFi interface and handover the session to it. Maintaining the GPRS interface in idle-IP connectivity mode while WiFi interface has active-IP connectivity consumes relatively small amount of energy[2] while it simplifies the IP-level location management in our strategy[3]. This location management simplification may also result in reduced energy consumption for IP-level location updates. The latter benefit is not observed by our experiments and requires a full implementation of the strategy (our future work).

An open issue is when a volume of IP packets to be transferred is considered large enough to activate the WiFi interface, i.e., to make the E2 transition in the state diagram of Figure 1. In addition to considering all other contextual parameters like application QoS, cost, mobility, etc, one should take into account the energy cost of turning the WiFi interface on and off for sending these packets. All these

---

[1] For the iPAQ device no PDP context is established, i.e., there is only GSM connectivity and no IP connectivity to the GPRS network.

[2] From Tables 4 and 5 the amount of power increase for Qtek is (72.19 - 69.77) / 69.77 = 3.5% and for iPAQ is (65.40 - 62.94) / 62.94 = 3.9%.

[3] This is not the case in, e.g., the connectivity management strategy used in the Microsoft Windows Mobile OS used in Qtek 9090 devices.

experiments confirm that our IP-level network activation strategy is also energy efficient as explained in the previous paragraph. For intermediary cases that require turning the WiFi interface on and off the strategy needs to consider the extra energy needed for switching the WiFi interface.

6 Conclusion

Mobile devices with multiple network interfaces can benefit from interface diversity when these interfaces are operating simultaneously. In this contribution we outlined an IP-level network activation strategy that aims at optimum use of the complementary features of the underlying heterogeneous networks. A key issue in deploying any strategy requiring the simultaneous operation of multiple interfaces is the increased battery energy consumption due to having more interfaces turned on simultaneously. Therefore, as a step towards validating our strategy we designed an experimental setting and carried out a number of experiments. Our objective was to see the relative energy cost for different network interfaces in two modern handheld devices. The paper reports the gained insight on how much WiFi and GPRS interfaces consume energy relative to each other in a particular device.

Our measurements indicate that our strategy is plausible from energy consumption viewpoint. Interestingly, early observations indicate that in some operation modes, e.g., when the mobile device is idle and reachable at the IP level or when the device intends to send/receive high volumes of data, the strategy outperforms those strategies that activate interfaces sequentially. This hints that, as our future work, we have to intelligently adapt the strategy to the context in which the terminal operates. The adaptation goal could be optimizing energy and monetary cost for having connectivity to cellular and local networks. Our work gives a clear view on the energy consumption behavior of different interfaces as compared with the total energy consumed in a mobile device.

7 Acknowledgment

The authors would like to thank Hans Zandbelts and Martin Wibbels for their helpful comments and implementation assistance.

8 References

[1] P.Nikander, J.Ylitalo, and J.Wall. Integrating Security, Mobility, and Multi-homing in a HIP Way. Proceedings of Network and Distributed Systems Security Symposium (NDSS'03), San Diego, USA, Feb 2003.

[2] X.Zhao, C.Castelluccia, and M.Baker. Flexible Network Support for Mobility. Proceedings of the 4th Annual International Conference on Mobile Computing and Networking (MobiCom), Dallas, Oct 1998.

[3] M.S.Bargh, H.Zandbelt, and A.Peddemors. Managing Mobility in Beyond-3G Environments. Proceedings of 7th IEEE International Conference on High Speed Networks and Multimedia Communications (HSNMC), Toulouse, France, Jun 30-Jul 2 2004.

[4] E.Gustafsson and A.Jonsson. Always best connected. IEEE Wireless Communications, vol.10, no.1 (Feb), pp.49- 55, 2003.

[5] Q.Song and A.Jamalipour. Network Selection in an Integrated Wireless LAN and UMTS Environment Using Mathematical Modeling and Computing Techniques. IEEE Wireless Communications, vol. 12, issue 3 (Jun), pp. 42-48, 2005.

[6] V.Gazis, N.Alonistioti, and L.Merakos, Toward a Generic "Always Best Connected" Capability in Integrated WLAN/UMTS Cellular Mobile Networks (and Beyond), IEEE Wireless Communications, vol. 12, no 3 (Jun), pp. 20-28, 2005.

[7] M.Inoue, K.Mahmud, H.Murakami, M.Hasegawa, and H.Morikawa, "Novel Out-of-Band Signaling for Seamless Interworking Between Heterogeneous Networks," IEEE Wireless Communications Magazine (Special Issue on 4G Mobile Communications), Apr 2004.

[8] G.Wu, M.Mizuno, and P.J.M.Havinga, "MIRAI Architecture for Heterogeneous Network," IEEE Communications Magazine, pp. 126-134, Feb 2002.

[9] A.J.H.Peddemors, H.Zandbelt, & M.S.Bargh (2004). A Mechanism for Host Mobility Management supporting Application Awareness. In Proceedings of the Second International Conference on Mobile Systems, Applications, and Services (MobiSys), Boston, Mass., USA, 2004.

[10] A.J.H. Peddemors, H. Eertink, & I.G. Niemegeers. Communication Context for Adaptive Mobile Applications. 3rd IEEE Conference on Pervasive Computing and Communications Workshops (PerCom 2005 Workshops), pp173-177, Kauai Island, HI, USA, 8-12 Mar 2005.

[11] Microsoft Developer Network (MSDN). http://msdn.microsoft.com/library/.

[12] Netperf homepage: http://www.netperf.org/netperf/NetperfPage.html, (retrieved on 25/11/2005)

[13] A.J.H. Peddemors, H. Eertink, & I.G. Niemegeers. CoSphere: an extensible network resource abstraction for applications on mobile devices. Under submission to MobiSys’06.

[14] S.Chandra. "Wireless network interface energy consumption Implications for popular streaming formats," Multimedia Systems, Springer-Verlag, vol. 9, no. 2, pp. 185-201, 2003.

[15] E.Shih, S.-H.Cho, N.Ickes, R.Min, A.Sinha, A.Wang, and A.Chandrakasan. Physical layer driven protocol and algorithm design for energy-efficient wireless sensor networks. Proceedings of the Seventh Annual ACM Conference on Mobile Computing and Networking, pp. 272-287, Rome, Italy, July 2001.

[16] R.Kravets and P.Krishnan. Power Management Techniques for Mobile Communication. Proceedings of the Fourth ACM International Conference on Mobile Computing and Networking (MOBICOM), 1998.

[17] E.Shih, P.Bahl, and M.J.Sinclair. Wake on Wireless: An Event Driven Energy Saving Strategy for Battery Operated Devices. Proceedings of the Eighth Annual ACM Conference on Mobile Computing and Networking, Altanta, Georgia, USA, Sep 2002.cc