Energy Efficiency of Dynamic Interface Selection Mechanisms in Wireless Ad-Hoc Networks

Luis Sanchez, Jorge Lanza, and Luis Muñoz

Laboratorios I+D Telecomunicación, Plaza de la Ciencia s/n, 39005 Santander, Spain
`{lsanchez,jlanza,luis}@tlmat.unican.es`

Abstract. Energy efficiency is critical to ensuring scalability, embedding, and portability of emerging computing and communication systems. It is of particular interest in the design of mobile computing systems because of the limitations in energy and power availability. This paper presents and compares in terms of energy efficiency two strategies for the dynamic selection of the outbound interface on multi-radio devices in wireless ad-hoc networks. Findings from the studies show that intelligent selection of communication interface in heterogeneous ad-hoc networks leads to more efficient use of the energy consumed while assuring the quality of service parameters necessary for the correct provision of applications running on top of wireless ad-hoc mobile networks.

1 Introduction

Although wireless networks have existed for many years already, explicit concern about their energy efficient operation has emerged only recently. It is quite evident that when the power source is either costly or in short supply, energy efficiency is of paramount importance. In some wireless network applications, energy is actually entirely non-renewable and is thus an overriding constraint for the design and operation of the network.

However, things are not that simple. First of all, if energy efficiency is the only concern in a communication system, one might as well transmit nothing. Energy reserves would thus remain intact perpetually. Clearly communication performance is also of paramount interest. Thus, the choice of how to incorporate energy efficiency in the overall design is far from clear. One approach is to try to minimize energy consumption subject to throughput (or delay) staying above (or below) a certain threshold. Alternatively, one can try to maximize throughput (or minimize delay) per joule of expended energy. Neither of these approaches led to simple precise formulations or easy solutions.

Different approaches have been proposed to provide more efficient energy consumption at different layers. At the link layer, transmissions may be avoided when channel conditions are poor, as studied in [1]. Also, error control schemes that combine Automatic Repeat Request (ARQ) and Forward Error Correction (FEC) mechanisms may be used to conserve power (i.e. trade off retransmissions with ARQ versus longer packets with FEC) as in [2]. Energy efficient routing protocols may be achieved by establishing routes that ensure that all nodes equally deplete their battery

K. Pentikousis et al. (Eds.): MONAMI 2010, LNICST 68, pp. 188–197, 2011.

power, as studied in [3]-[4] or that avoid routing through nodes with lower battery power. More complex solutions such as [5] exploit cross-layer operation and control the network topology by varying the transmitted power of the nodes so that certain network properties are satisfied.

Cross-layer design is particularly interesting under energy constraints, since not only energy across the entire protocol stack must be minimized, but also system performance must be optimized. While layer-specific solutions might disregard valuable information residing in other layers, cross-layer solutions can exploit global knowledge of the system to provide sub-optimal solutions at each of the different layers that result in optimal solutions at system level.

In this paper, the framework over which dynamic interface selection mechanisms have been implemented is briefly presented and two different outbound interface dynamic selection strategies are evaluated from an energy efficiency point of view.

2 Universal Convergence Layer

The concept of isolating the upper-layers from the underlying wireless technologies and thus providing real multi-mode can be achieved by introducing a Universal Convergence Layer (UCL) within the protocol stack. The UCL can be seen as a twofold approach. It will mainly act as an enabler for backward and forward compatibility by defining a common interface towards the network layer while managing several different wireless access technologies independently of their PHY and MAC layers. On the other hand, UCL also enables the cross-layer optimisation paradigm. Its privileged location within the protocol stack gives the UCL the possibility to support the information flow both bottom-up (e.g. use of SNR information for enriching the decision-making process in an ad hoc routing algorithm) and top-down (e.g. tuning of MAC parameters depending on the battery status or QoS requirements).

Figure 1 presents the different pieces in which the UCL can be divided. Each of these modules is specialized in providing the different features the UCL offers. This approach allows the easy addition and removal of functionalities depending on the requirements and characteristics of the system on which it will run. For the sake of simplicity and due to the scope of this paper, the building blocks presented in Figure 1 are the ones that are involved in the dynamic interface selection strategies that are evaluated in this paper.

- **Multi-radio Management module.** The UCL hides the complexity of the available air interfaces and offers a unique interface to the upper layers. UCL aims at masquerading multihoming by aggregating the different network interfaces (one per access technology the node is equipped with) on a single interface. Multi-radio devices in wireless ad-hoc networks will therefore keep only one network identifier (i.e. IP address) while being able to exploit the opportunities offered by its multi-radio nature. The possibility of using different links to one destination allows the UCL to intelligently modify the output interface according to the requirements and needs of the system.

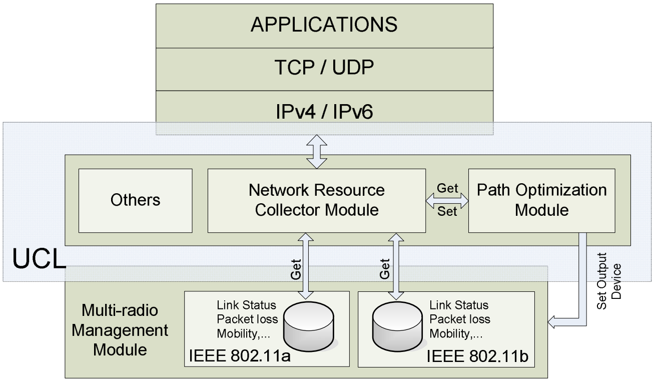

Fig. 1. UCL simplified high-level architecture

- **Network Resource Collector.** Data exported by the underlying technologies may offer important information to other layers in the stack. As this information has different meaning depending on where and by whom it is used, this module is specialized on collecting and translating data to make it available to the rest of the modules thus enabling the cross-layer optimization. As the different optimization rules are specific depending on the final purpose, this task is decentralized on different modules.
- **Path Optimization module.** Selecting the optimal output interface may depend on user preferences as well as the current status of the network. The outcome from the Network Resource Collector Module is analyzed and a decision is taken based on that.

Coordinated operation of these modules allows implementing communications management strategies as described in the following section.

3 Power-Aware Optimization Based on Dynamic Outbound Interface Selection

The UCL leverages Multiple Attribute Decision Making (MADM) algorithm in order to decide which of the available outgoing interfaces to use in order to guarantee system best possible performance. A utility function is evaluated for all the possible link layer interfaces and the one with the best result is selected.

Several policies and parameters might be used for implementing the utility function calculation. In this paper two of them are evaluated and compared. In both cases, the sender is the one that selects the interface on which it is transmitting the packets.

In the first of them, the UCL decides which interface to send the packet through taking into account the Signal to Noise Ratio (SNR) observed in each of the wireless channels available. Thresholds are set in advance for each wireless technology the device is equipped with. This way, the selection of the outbound interface maximizes the performance of the system given the SNR observed on each of the interfaces the node can use for transmitting each IP datagram. It is important to note that SNR is a mean for assessing channel status. When the 802.11a channel becomes is bad (i.e. SNR goes below a pre-defined threshold), the status of the 802.11b channel is typically good.

The second adaptation technique is based on assessing the channel conditions using the number of lost packets. More specifically the UCL evaluates the bursts of correct and lost packets in order to know the actual status of the channel, and decides which radio interface to use correspondingly.

Wireless channels typically behave in a burst manner due to the fading processes they suffer. Hence, even when the SNR is relatively low we can correctly receive packets. This strategy attempts to take advantage of this behaviour and react quickly to these slots. Since it is not possible to get information about the packet los rate in the RAT that is not active, UCL takes decisions based on the bursts of erroneous or correct frames transmitted through the active RAT. In this sense, when a pre-determined number of frames are lost in a row, the UCL switches the transmission to a technology that is more robust. Of course, if channel conditions are so bad that even the most robust RAT suffers high FER, UCL cannot do anything. On the opposite direction, when a burst of frames are correctly sent, the UCL changes to the RAT that provides faster and/or energy efficient behaviour. To avoid as much as possible ping-point effects, the size of the bursts can be dynamically adapted so that it rises upon the detection of such abnormal behaviour.

The number of packets lost on which the threshold is set for swapping to use a more robust technology (downgrading threshold) can be adapted accordingly to have a trade-off between the throughput and the application loss. Similarly, the number of correct packets, received in a row through the more robust technology, with which we decide to use a faster technology (upgrading threshold) can also be tuned.

It is important to mention that the processing needed for this decision making is negligible since it is based on data exported by the underlying technologies and utility function calculation at the Path Optimization module can be greatly reduced. Additionally, since decisions are locally taken at the transmitter, no further signalling is necessary between the nodes involved in the communication.

4 Dynamic Interface Selection Mechanisms Power Efficiency Assessment

This section presents the results obtained from the evaluation analyses carried out in order to prove and validate the benefits introduced by the UCL solution for the selection of the most appropriate network interface for outgoing traffic.

The scenario selected for the analyses offered a range of situations from very good channel conditions (Location 1 in Figure 2) to poor ones that will produce a deep degradation of the communications (Location 3 in Figure 2). Basically, the charter of selecting this environment was to test the UCL on a real-world scenario which would allow us to extract conclusions that can be directly mapped on real user experiences. Two laptops were used, each of them equipped with one IEEE 802.11a and one IEEE 802.11b interface. Main rationale for this selection is that these technologies are

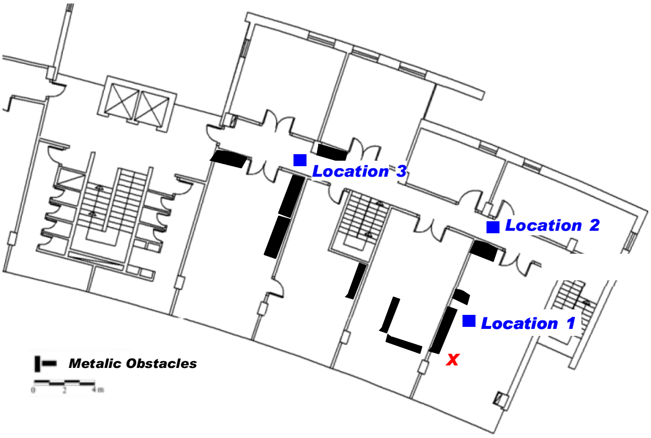

Fig. 2. Measurement campaign environment

nowadays the most widely used for Wireless Personal Area Networks (WPAN) and Wireless Local Area Network (WLAN) scenarios. Moreover, they operate in different frequency spectrum ranges, use different PHY and have slightly different MAC mechanisms, so they represent a good example of heterogeneity. The tests consisted on moving one of the laptops from Location 1 to Location 3 and back again while the other stays fixed in Location 1 (actually on the red cross). Transmission was done from the fixed laptop towards the mobile one. For both technologies we forced the transmission rate to be always the maximum possible, this is, 11 Mbps for IEEE 802.11b and 54 Mbps for IEEE 802.11a.

We simulated in Matlab® this moving scenario by generating a sequence of states which modelled the reception status of the transmitted frames. In our case we have used a Gilbert-Elliot model in which the parameters have been derived from [6] and tuned with our own experimental characterization of the office scenario shown in Figure 2. A total of 60000 link-layer frames were transmitted per simulation. In the scenario simulated, two thirds of the frames where transmitted over the best channel (Location 1) while the other third was equally transmitted over the other two channels (Locations 2 and 3). It is important to note that IEEE 802.11 MAC implements an ARQ scheme by which each frame is retransmitted a given number of times if an error occurs. Typically this number is fixed to 4. Thus, Frame Error Rate (FER) is not equivalent to Packet Error Rate (PER). This fact was also taken into account in the simulations. A Monte Carlo approach was taken so that simulations were run 1000 times.

Based on power consumption per bit values taken from [7] we can compare the efficiency of the different strategies in terms of power consumption and measure the optimization when weighting them against blind selection of the wireless interfaces.

As is shown in Figure 3, a blind selection of the 802.11a interface leads to slightly better use of power. This is due to its better natural energy efficiency (see Figure 4).

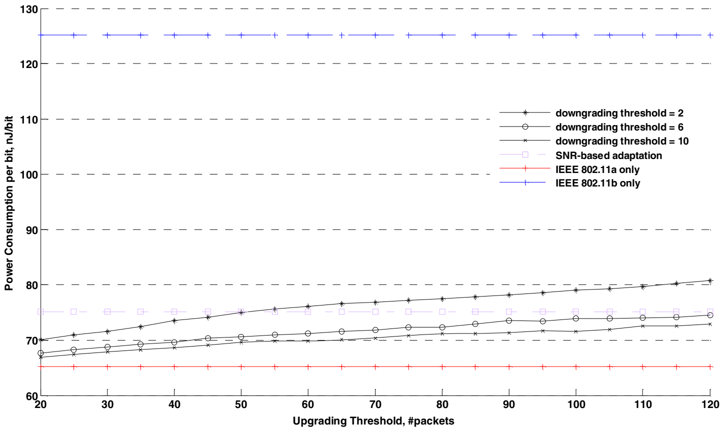

Fig. 3. Power consumption efficiency of UCL vs non-UCL approaches

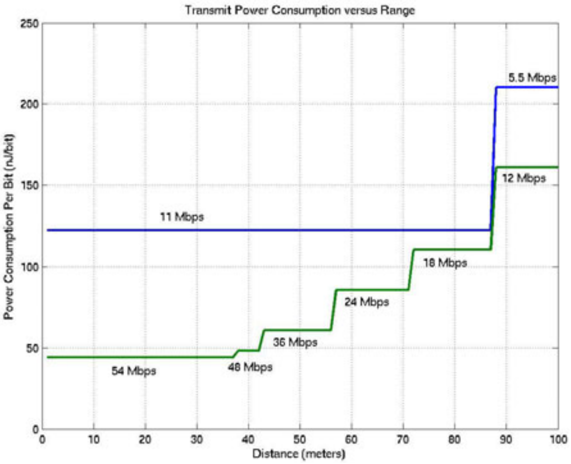

Fig. 4. Energy consumed per transmitted bit (802.11b vs. 802.11a) (source [7])

The consumed battery per unit of information is lower than the one of the other strategies. However, IEEE 802.11a is a less robust technology and suffers from higher FER under the same conditions. Thus, as it is shown in Figure 5, during the simulation a significant packet loss is experienced when using only the 802.11a. Packet loss shown in Figure 5 corresponds with loss at application level, this is, although the frame containing the packet was retransmitted 4 times at link level, none of these retransmissions arrived at the receiver error-free. Despite this packet loss, that would ruin the communication, and the associated number of retransmissions, inherent energy efficiency of 802.11a interface for transmission compensates it and results on slightly better results when looking only at the energy efficiency metrics.

This is also true because of the ratio chosen in the scenario between good and bad channel conditions (i.e. 2/3 vs 1/3). A different ratio would change the results in terms of energy efficiency of using IEEE 802.11a only. In contrast, when only using 802.11b the efficiency is much lower even though the low frame error rate experienced is much better on any of the channels that compose the moving scenario.

Having said this, it is not enough for making a correct comparison to look at only one metric but is necessary to assess the gain of the proposed solutions both from an energy efficient point of view and from a quality of service standpoint.

This is better seen in Figure 5 where the left Y-axis represents the power consumption per bit and the right Y-axis the packet loss at application level obtained in each of the cases. When selecting the appropriate parameters for the packet loss based adaptation strategy the UCL achieved relatively small power consumption per bit figures, but highly reduced the packet loss at application level. For example, using 6 packets as the downgrading threshold and 65 packets as the upgrading one we can obtain three times less packets lost at the application level, while still keeping the overall power consumption per bit only 10% higher.

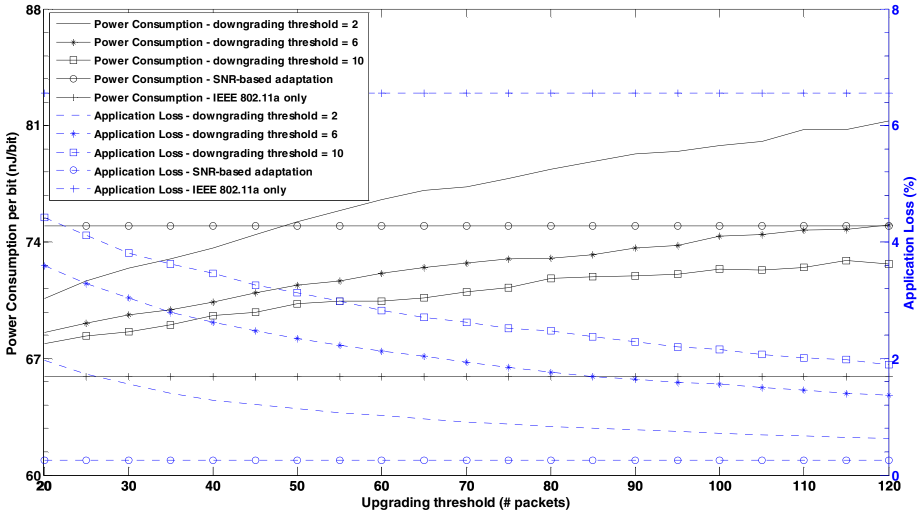

Fig. 5. Power consumption efficiency and packet loss of UCL and 802.11a only approaches

When the SNR-based approach is used the packet loss is reduced drastically (around 20 times less), while keeping the power consumption only 15% higher than when 802.11a interface is used all the time.

Making similar studies with the achieved throughput, we were able to obtain almost 2.5 times more throughput than using 802.11b only while still reducing the overall power consumption per bit by 40% as it can be seen in Figure 6 for the case when 6 packets are used as the downgrading threshold and 65 packets as the upgrading one. This means almost halving the total power consumed during the simulation. When the SNR-based approach is used we also double the throughput while the power consumption is still 40% less than when using the 802.11b interface only.

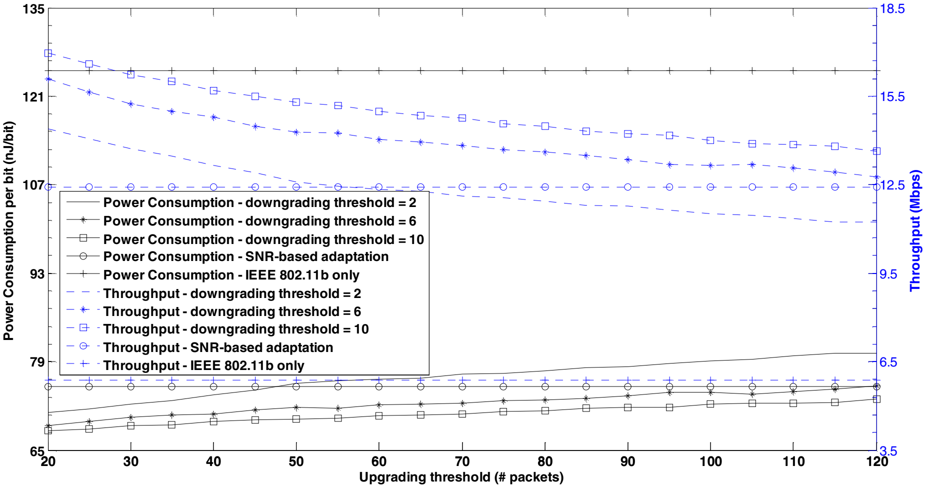

Fig. 6. Throughput versus power consumption efficiency of UCL and 802.11b only approaches

We can conclude that an intelligent use of the available interface leads to a sub-optimal power consumption, but does not jeopardize the system performance as would be the case with an 802.11b only approach, where throughput is reduced or with an 802.11a only approach, where application-level loss is much higher.

Only UDP transport protocol has been used during the assessment since the TCP congestion avoidance mechanisms would have stopped the transmitter when the channel conditions were poor, thus making it impossible for us to compare the energy efficiency of the different alternatives shown. UCL would keep performance on a reasonable level by using the most robust wireless technology while the device is facing bad channel conditions [8]. On contrary, when only IEEE 802.11a is used, as soon as the communication experience, TCP emitter is stopped [9]. Therefore, in terms of energy efficiency, this latter approach would be better, but at the cost of ruining the throughput and hence the services being provided.

5 Conclusions

It has been demonstrated that wireless ad-hoc networks energy efficiency can be enhanced by taking advantage of multi-radio capabilities of the devices. Opposite to other approaches, like rate adaptation mechanisms [10], that focus only on one technology, UCL allows seamless service provision in multi-radio mobile nodes. This enables larger versatility since the specific advantages of all the available WPAN and WLAN technologies can be exploited depending on the application and user requirements. Even though a vertical handover might be forced in order to use the most appropriate interface during the actual service provision, the session is not affected.

The solution proposed has proven to be able to offer a good range of possible outputs depending on the parameters chosen for each of the adaptation strategies. In this sense, the focus has not been so much on finding the optimal configuration but on assessing the optimization capacity of the solutions proposed. This is important since it would be able to fit the demands of a wide range of services, users and channel conditions.

Most interesting finding is that taking advantage of the global view that the convergence layer provides (access to information from multiple layers and access technologies) a complete map of the system can be inferred. After matching it with the user requirements, the UCL can take the appropriate decisions in order to best serve the final users' wishes, optimizing their quality of experience not only by providing enhanced bandwidth but also improving the power consumption efficiency or enhancing the service provision by lowering the packet loss due to wireless channel impairments. In this sense, the UCL selection strategies does not look for optimization of throughput only but the decisions taken depend on a multi-parametric matrix that include battery status, type of application or user preferences for example.

Implementation over real testbed has been done for the SNR-based adaptation approach. The results from experimental validation support the conclusions from this paper.

Acknowledgements

The authors would like to thank the Spanish Government (Ministerio de Ciencia e Innovación) for its funding in the project "Cognitive, Cooperative Communications and autonomous SErvice Management", C3SEM (TEC2009-14598-C02-01).

References

1. Zorzi, M., Rao, R.R.: Energy constrained error control for wireless channels. IEEE Personal Communications Magazine 4(6), 27–33 (1997)
2. Lettieri, P., Fragouli, C., Srivastava, M.B.: Low power error control for wireless links. In: Proceedings from the 3rd Annual ACM/IEEE International Conference on Mobile Computing and Networking, pp. 139–150 (September 1997)
3. Woo, M., Singh, S., Raghavendra, C.S.: Power Aware routing in mobile ad hoc networks. In: Proceedings from the 4th Annual ACM/IEEE International Conference on Mobile Computing and Networking, pp. 181–190 (October 1998)
4. Chang, J.-H., Tassiulas, L.: Energy conserving routing in wireless ad-hoc networks. In: Proceedings from the 19th Annual Joint Conference of the IEEE Computer and Communications Societies, pp. 22–31 (March 2000)
5. Ramanathan, R., Rosales-Hain, R.: Topology Control of Multihop Wireless Networks using Transmit Power Adjustment. In: Proceedings from the 19th Annual Joint Conference of the IEEE Computer and Communications Societies, pp. 404–413 (March 2000)
6. Arauz, J., Krishnamurthy, P.: Markov modeling of 802.11 channels. In: Proceedings from the 58th IEEE Vehicular Technology Conference (October 2003)
7. Texas Instruments White Paper, Low Power Advantage of 802.11a/g vs. 802.11b (December 2003)

8. Sánchez, L., Lanza, J., Muñoz, L.: Experimental assessment of a cross-layer solution for TCP/IP traffic optimization on heterogeneous personal networking environments. In: Cuenca, P., Orozco-Barbosa, L. (eds.) PWC 2006. LNCS, vol. 4217, pp. 284–296. Springer, Heidelberg (2006)
9. García, M., Agüero, R., Muñoz, L.: On the unsuitability of TCP RTO estimation over bursty error channels. In: Niemegeers, I.G.M.M., de Groot, S.H. (eds.) PWC 2004. LNCS, vol. 3260, pp. 343–348. Springer, Heidelberg (2004)
10. Lacage, M., Manshej, M.H., Turletti, T.: IEEE 802.11 Rate Adaptation: A Practical Approach. In: Proc. MSWiM 2004, Venice, pp. 126–134 (June 2004)