

Energy-efficient TDMA Medium Access Control protocol scheduling

Paul J.M. Havinga & Gerard J.M. Smit

*University of Twente
Department of Computer Science
Enschede, the Netherlands*

{havinga, smit}@cs.utwente.nl

Abstract - *In this paper we study the energy efficiency and channel efficiency of TDMA MAC protocol scheduling mechanisms. Most MAC protocols are based on phase grouping that basically has three phases in a frame: uplink, downlink and reservation. We propose a new mechanism in which we have multiple uplink and downlink phases. These phases are grouped per mobile in a frame. Although this has a negative effect on the capacity of the channel, it allows the mobile to turn the power off from the wireless interface for a longer period. We made this choice since in a mobile multimedia environment it is more important that connections have a certain QoS, than highest possible bandwidth. We present an analysis in which these two basic mechanisms are compared in respect to bandwidth efficiency and energy efficiency. We have developed and implemented a novel MAC protocol based on mobile grouping that provides Quality of Service (QoS) support for diverse traffic types.*

I. INTRODUCTION

The energy consumption of portable computers like PDAs and laptops is a limiting factor in the amount of functionality that can be placed in these devices. The wireless network interface of a mobile computer consumes a significant fraction of the total energy consumed by a mobile computer. More extensive and continuous use of network services will aggravate this problem. Energy efficiency can be improved at various layers of the communication protocol stack. However, even today, research is still focused on performance and (low power) circuit design. There has been substantial research in the hardware aspects of mobile communications energy-efficiency, such as low-power electronics, power-down modes, and energy efficient modulation. Due to fundamental physical limitations though, progress towards

further energy-efficiency will become mostly an architectural and software-level issue.

The context of this paper is data link-level communication protocols for wireless networks which provide multimedia services to mobile users. Portable devices have severe constraints on the size, the energy consumption, and the communication bandwidth available, and are required to handle many classes of data transfer over a limited bandwidth wireless connection, including delay sensitive, real-time traffic such as speech and video. Our approach is driven by two major factors. The first factor is that the design should be *energy-efficient* since the mobiles typically have limited energy capacity. The second factor is that it should provide support for multiple traffic types, with appropriate Quality of Service levels for each type. The aim is to meet the required QoS, while minimising the required amount of energy.

The notion of QoS over a wireless link has been the focus of much recent research, and several scheduling algorithms have been proposed (e.g. [5][6][16][17][21]). Access protocols for these systems typically only address network performance metrics such as throughput, efficiency, and packet delay (e.g. [1][2][7][9][10][13]). However, thus far, little attention is given to energy conserving protocols, and researchers mainly focuses their effort on energy reduction by circuit design. Very recently there is a growing interest in energy-efficient design, although mainly concentrating on Medium Access Control (MAC) and link-layer energy reduction techniques ([3][4][11][12][18][19][20][23][24]).

In this paper we present a MAC scheduling principle for a TDMA system that does not provide the most efficient bandwidth utilisation, but reduces significantly the energy consumption that is needed for the mobile to communicate. We believe that this is a valid choice since in a mobile multimedia environment it is more important that

connections have a certain QoS, than pure raw bandwidth. Performance sufficiency and energy efficiency have become the predominant platform requirements for battery-powered mobile multimedia computing devices [15].

II. ENERGY DISSIPATION IN WIRELESS COMMUNICATION

A significant part of the power consumption needed for wireless communication is due to the wireless interface, the transceiver. Typically, the transceiver can be in five modes (see Figure 1); in order of increasing energy consumption these are: off, sleep, idle, receive, and transmit. In transmit mode, the device is transmitting data; in receive mode, the receiver is receiving data; in idle mode, it is doing neither, but the transceiver circuit is still powered and ready to receive or transmit; in sleep mode, the transceiver circuitry is powered down, although in some implementations a small amount of circuitry is still listening for incoming transmissions.




Figure 1: Typical operating modes of a wireless modem.

A wide variety of reasons for unessential energy consumption exist. These reasons include the overhead of the protocol, the high error rate on wireless channels, the inactivity threshold time after which a radio will enter a low energy consuming operation mode, the need to receive messages, the occurrence of collisions, and the turnaround time between various operating modes. The data link layer, and in particular the Medium Access Control protocol, can alleviate these problems significantly. In [8] the causes of unessential energy consumption and main principles of energy-efficient MAC protocol design have been explored in more detail. In this paper we concentrate on ways to reduce the effect of turnaround time, and on minimising the time a radio needs to be in high power mode (idle, receive, transmit).

There are basically three effects that contribute to the required energy for a transition from sleep to transmission or reception:

1. the required time and energy to change the operating mode from sleep to idle.
2. the required time and energy the interface has to be either in idle, receive and transmit mode, but is not

transmitting or receiving actual data. This is the overhead required to initiate and terminate the actual transmission. This time includes the required gap (guard time), interfacing delay, preamble, and the postamble.

3. the required time and energy to switch to the sleep mode after transmission or reception.

We assume the wireless physical header and trailer to be a fact that cannot be changed or improved with a MAC protocol, although the protocol can try to minimise the number of times that these are required.

The overhead introduced in the physical layer can be significant, e.g. for WaveLAN [22] it can be up to *virtual* 58.25 bytes (i.e. the time normally 58.25 bytes could have been sent) (for guard space (in which the silence level is measured), interfacing delay (required to synchronise to the internal slotsync moments), preamble and postamble, see Figure 2). Moreover, with this interface that has a throughput of 2 Mbit/s, a transition time from sleep to idle of 250 μ s already takes *virtually* 62.5 bytes (500 bits). The power consumption of the WaveLAN modem when transmitting is typical 1675 mW, 1425 mW when receiving, and 80 mW when in sleep mode (according to the specs [22]). Minimising the on time of the radio will thus significantly improve the energy efficiency of wireless communication.




Figure 2: WaveLAN physical layer block format.

This also shows that efficient data transmission (in terms of channel utilisation and energy consumption) can only be achieved if the amount of data that is transmitted in one burst is not too small. Scheduling traffic into bursts in which a mobile can continuously transmit or receive data – possibly bundled for different connections –, can reduce the number of transitions. Notice, however, that there is a trade-off with QoS parameters like delay and jitter.

To be fully effective, both the base station and the mobile must be synchronised [14]. When the base-station and mobile are synchronised in time, the mobile can go in standby or off mode, and wake up just in time to communicate with the base-station. The application running on a mobile with the least tolerable delay determines the frequency by which a mobile needs to turn its receiver on.

III. ENERGY-EFFICIENT TRAFFIC SCHEDULING

In this study we concentrate on *Time Division Multiple Access* (TDMA) schemes where a base-station co-ordinates access to one or more channels for mobiles in its cell. The channels can be individual frequencies in FDMA, time slots in TDMA, or orthogonal codes or hopping patterns in case of spread-spectrum. Hybrid TDMA/CDMA schemes benefit from both the capacity of TDMA schemes to handle high bit-rate packet-switched services, and the flexibility of CDMA techniques that allow smooth coexistence of different types of traffic [1].

Definitions

In TDMA a channel is divided into *slots*, and these slots are grouped into *frames*. The payload transmitted in one physical block is referred to as *packet*, and can span multiple slots (see Figure 2). We concentrate on a system with QoS provisions. In these systems mobiles can reserve resources (slots) for connections. A QoS manager (typically located on the base-station) receives transmission requests from the mobiles. The key to providing QoS for these connections will be the scheduling algorithm that assigns the bandwidth.

The *QoS manager* establishes, maintains and releases wireless connections between the base-station and the mobile and also provides support for handover and mobility services. Multimedia networking requires at least a certain minimum QoS and bandwidth allocation for satisfactory application performance. This minimum QoS requirement has a wide dynamic range depending on the user's quality expectations, application usage modes, and application's tolerance to degradation. In addition, some applications can gracefully adapt to sporadic network congestion while still providing acceptable performance.

The *slot scheduler* assigns bandwidth for connections. A schedule is broadcast to all mobiles so that they know when they should transmit or receive data. This schedule is called *Traffic Control Slot* (TCS). The slot scheduler is designed to preserve the admitted connections as much as possible within the negotiated connection QoS parameters.

Frame scheduling

In general we can distinguish three different phases: *uplink* phases, *downlink* phases, and *reservation* phases. In downlink phases the base station transmits data to the mobiles, and in uplink phases the mobiles transmit data to the base station. In the reservation phase mobiles can request new connections. In almost all MAC protocols a frame basically is divided into these three phases

concatenated. We will refer to this mechanism as *phase grouping*.

An alternative approach is to have in principle similar phases, but not grouped together in a frame according to the phase, but grouped together according to the mobile involved. The uplink and downlink phase of *one mobile* are grouped. So, as a result we will have in general multiple uplink and downlink phases in one frame. We will refer to this mechanism as *mobile grouping*.

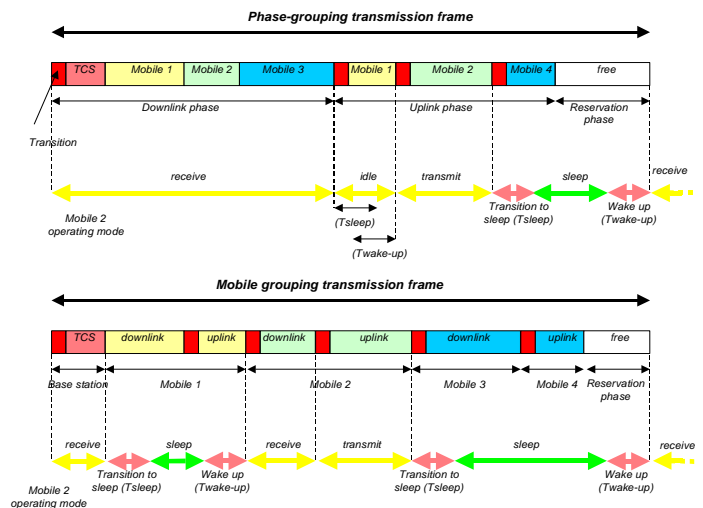


Figure 3: Grouping strategies in a transmission frame.

Figure 3 depicts the two grouping strategies. The upper figure shows a phase-grouping transmission frame, and the lower figure shows a mobile-grouping transmission frame.

Both strategies have positive and negative implications. Mobile grouping strives to minimise the power consumption, while phase grouping strives for an optimal channel utilisation.

In the mobile grouping strategy the slot-scheduler tries to group the transmissions and receptions of a mobile as much as possible according to the service classes, QoS and current load. Grouping of uplink and downlink traffic for one mobile implicates that there is some space between sending and receiving to allow the transceiver to switch its operating mode from sending to receiving (i.e. guard space, preamble, postamble). This has a negative effect on the capacity of the wireless channel. The advantage is that it allows the mobile (i.e. the radio device) to turn its power off for a longer period, and that it makes less operating-mode transitions. If we would, in contrast to the mobile grouping of uplink and downlink traffic, group the downlink traffic from the base-station to all mobiles (*phase grouping*), then the space between sending and receiving that is required for mobile grouping, is not present for the

downlink traffic. This is the reason why the available bandwidth of mobile grouping is less than the available bandwidth in phase grouping. Most MAC protocols group the traffic from the base-station, mainly because of its efficient use of the available bandwidth. However, there are consequences of phase grouping related to the energy consumption:

- 1) With phase grouping, the mobile is in general forced to receive the complete downlink packet, and will ignore the data not destined for the mobile. The receiver of the mobile must be on for a longer period (i.e. during the whole downlink period because it needs to synchronise using the preamble of the radio packet)¹. If the radio would be capable of synchronising during the transmission of a downlink packet, then the mobile might be able to power down during non-relevant parts of the downlink phase.
- 2) The period between two operations is too small to enter sleep mode. This period is determined by the time needed to enter sleep mode (T_{sleep}) and the time needed to wake-up ($T_{wake-up}$). This is shown in the upper part of Figure 3 where Mobile 2 cannot enter sleep mode after reception of the downlink packet, but is forced to idling. Because the operating modes of phase grouping for a mobile are spread in the frame, the power-mode transition times T_{sleep} to enter sleep mode, and $T_{wake-up}$ to wake from sleep mode limits the time a mobile can stay in sleep mode.

Both effects lead to higher energy consumption for the mobile. This shows that there is a trade-off between performance (channel efficiency) and energy consumption. The energy gained with mobile grouping depends on 1) the amount of data in the downlink phase that is not destined to the mobile, but must be received by the radio device because it is present in the downlink transmission frame. And 2) the amount of time between receiving the downlink packet and the uplink packet. Only when this time exceeds the time required to enter sleep mode (T_{sleep}) plus the time needed to wake-up ($T_{wake-up}$), energy can be saved. Otherwise, it is more profitable for the mobile to remain idle, waiting for its time to transmit its uplink packet.

In the next section we will provide a quantitative analysis of the channel efficiency and energy efficiency of the two strategies.

¹ A power-optimised network interface could stop receiving the downlink packet after it has received data for mobile 2, and thus also enter sleep mode. In this case the downlink schedule order determines the amount of data to be received.

IV. ANALYSIS

We will now evaluate the effects of mobile grouping on the available bandwidth and on the energy consumption. We will compare this with the phase grouping mechanism. The properties of interest are:

TCS	The size of the Traffic Control Slot.
O	The overhead needed to transmit a packet. The overhead O consists of the overhead when the interface must be idle O_{idle} (required for guard space and interfacing delay), plus the overhead O_p that is required for preamble and postamble. The interfacing delay is caused by two factors. First, the delay caused by the wireless interface to synchronise to its internal syncslots. Second, we have an additional delay because we must also synchronise to the time slots that the MAC protocol uses. We assume that the MAC protocol uses fixed time slots, but since each packet is not a multiple of this slot size (because of the overhead in the wireless interface) we need to incorporate a delay with an average length of the size of a time slot divided by two.
T_{sw}	This is the time needed by the wireless interface to enter sleep mode T_{sleep} plus the time $T_{wake-up}$ needed to wakeup from sleep mode to an active mode (idle, receive or transmit).
C	The number of bytes used for the reservation phase (collision phase).
O_{total}	The total overhead in a frame that is introduced to transmit the actual data over the wireless link.
F	The size of a transmission frame.
TD	The total size available for a mobile to transmit data packets. This can be expressed with:
	$TD = F - O_{total}$
D	The size of a packet (uplink and downlink) used by a mobile. We assume that the whole frame is used, and that all mobiles have an equal share. The size of an uplink and a downlink packet is thus dependent on the number of mobiles using the frame. It can be expressed with:
	$D = TD / 2M$
M	The number of mobiles, each with uplink and downlink packets.

All properties can be expressed in bytes. When a property is related to time, then we use the virtual overhead

that expresses the number of bytes the wireless channel can transmit in that time.

In our analysis we will assume that each mobile has both uplink and downlink connections that both have similar bandwidth requirements. We further assume a packet length that allows a mobile to enter sleep mode. Thus, the packet length is greater than T_{sw} .

Analysis phase grouping

Figure 4 shows a typical phase grouping transmission frame with three mobiles, each using downlink and uplink packets.




Figure 4: Phase grouping transmission frame

In general we have M mobiles, each with uplink and downlink packets. The total overhead O_{total} can be expressed with:

$$O_{total} = O + TCS + M.O + O + C$$

or,

$$O_{total} = TCS + (M+2).O + C$$

We will now determine the total time a mobile can enter sleep mode. This time can be used to evaluate the energy a mobile consumes for its wireless interface.

We assume that a mobile is required to receive the whole downlink packet from the base station. Whether a mobile is able to transmit its uplink packet depends on the schedule made by the base station. In our analysis we evaluate the sleep period of a mobile that is scheduled as second in the uplink phase (e.g. Mobile 2 in Figure 4). We can then divide the uplink period in three phases: pre-uplink, uplink (in which the mobile transmits its data), and post-uplink.

When there is just one mobile, we do not have a pre-uplink phase. The mobile can only sleep in the contention phase. The total sleep time $T_{sleep-period}$ of the mobile in this situation is:

$$T_{sleep-period} (M=1) = O + C - T_{sw}$$

When there are more mobiles, then we have all phases. The pre-uplink sleep period is:

$$T_{sleep-pre} = (D + O) - T_{sw}$$

The post-uplink sleep period is the time in which the remaining $(M-2)$ data packets from the other mobiles are sent:

$$T_{sleep-post} = (D + O) \cdot (M-2) - T_{sw}$$

Together with the collision phase this gives a total sleep time for $M > 1$:

$$T_{sleep-period} (M > 1) = (D + O) - T_{sw} + (D + O) \cdot (M-2) + O + C - T_{sw}$$

Thus:

$$sleep-period \begin{cases} O + C - T_{sw} & M = 1 \\ (D + O) \cdot (M-1) + O + C - 2 T_{sw} & M > 1 \end{cases}$$

Analysis mobile grouping

We will now evaluate mobile grouping using the same assumptions as applied for phase grouping. Figure 5 gives an example of a mobile grouping transmission frame.




Figure 5: Mobile grouping transmission frame

In general we have M mobiles, each with uplink and downlink packets. The total overhead O_{total} can be expressed with:

$$O_{total} = O + TCS + O + 2M.O + O + C$$

or,

$$O_{total} = TCS + (2M+3).O + C$$

We can divide the uplink period in three phases: pre-uplink, uplink (in which the mobile transmits its data), and post-uplink.

When there is just one mobile, then we do not have a pre-uplink phase. The mobile can only sleep in the contention phase. The total sleep time $T_{sleep-period}$ of the mobile in this situation is:

$$T_{sleep-period} (M=1) = O + C - T_{sw}$$

When there are more mobiles, then we have all phases. The pre-uplink sleep period is:

$$T_{sleep-pre} = (2D + O) - T_{sw}$$

The post-uplink sleep period is the time in which the remaining $(M-2)$ data packets from the other mobiles are sent:

$$T_{sleep-post} = 2(D + O) \cdot (M-2) - T_{sw}$$

Together with the collision phase this gives a total sleep time for $M > 1$:

$$T_{sleep-period}(M > 1) = (2D + O) - T_{sw} + 2(D + O) \cdot (M-2) + O + C - T_{sw}$$

Thus:

$$sleep\text{-}period \begin{cases} O + C - T_{sw} & M = 1 \\ (2M-2)D - 2T_{sw} + (2M-2)O + C & M > 1 \end{cases}$$

Evaluation

We will now apply to these equations the characteristics of the WaveLAN modem together with some assumptions for some of the MAC protocol parameters. Note that the numbers mentioned are derived from an implementation with a WaveLAN modem of a MAC protocol using the mobile grouping strategy (see Section 0).

F 2544 bytes with a frame rate (f) of 100 Hz, and 5088 bytes with a frame rate of 50 Hz (assuming a transmission rate of 2 Mb/s)

TCS 53 bytes. (one ATM cell)

O 71 bytes. ($O_p = 37$ bytes, $O_{idle} = 22$ bytes. The internal time slots are 24 bytes, which results in an average synchronisation delay of 12 bytes)

C 53 bytes. (one ATM cell)

T_{sw} 73 bytes ($T_{sleep}=10$ bytes, $T_{wake-up}=63$ bytes)

The results are shown in Figure 6 and Figure 7. Figure 6 shows the total overhead O_{total} caused by the two mechanisms. As expected, phase grouping induces a smaller overhead than mobile grouping. When there are many mobiles using the frame (both in the uplink and in the downlink direction), then the overhead constitutes a significant part of the total available bandwidth (e.g. approx. 68% for mobile grouping and approx. 39% for phase grouping with 10 mobiles and a frame rate of 100 Hz). The figure also shows the percentage of overhead incurred with two different frame rates (and thus two different frame sizes). We evaluated a frame rate of 50 Hz

and of 100 Hz. A lower frame rate (larger frame size) decreases the overhead. This reduces the overhead that is required to transmit a certain amount of data, but increases the latency.




Figure 6: Overhead vs. number of mobiles as a function of frame rate.

When we look at the consequences for energy consumption, mobile grouping performs better. This is shown in Figure 7. The on time per mobile is shorter (or in other words, the sleep period per mobile is larger) when using mobile grouping compared to phase grouping. For example, with 10 mobiles and a frame rate of 100 Hz, the on time of the mobile's network interface is decreased from approx. 44% with phase grouping to approx. 24% with mobile grouping.




Figure 7: On time per mobile versus number of mobiles as a function of frame rate.

The figure shows that the on time is already significantly reduced with a small number of mobiles. The difference between mobile grouping and phase grouping is larger when the frame rate is decreased.

As the overhead increases with the number of mobiles using the wireless channel, mobile grouping seems particularly attractive for systems with a small cell size (e.g. pico-cellular with the size of an office-room). In these systems the number of mobiles in one cell will in general be small, and the available bandwidth high. Mobile

grouping strategy will then have a small overhead while allowing a large sleep period.

Note that the assumptions we have made are conservative for two reasons. First, we assumed that the whole frame is used. If this is not the case, then the sleep period can be larger. However, when using phase grouping, a mobile is in general forced to receive the whole downlink packet, and cannot enter sleep mode in that phase; whereas mobile grouping only needs to receive the TCS. Second, we assumed that the amount of uplink traffic is equal to the amount of downlink traffic. This might be true for voice applications (mobile phone), but is in general not true for applications running on a mobile computer. For these applications the downlink traffic in general will use more bandwidth. The disadvantage of having one large downlink packet (phase grouping) then becomes even more apparent.

Implementation

We have implemented a test-bed of the network interface that we used to experiment with the various techniques and mechanisms for e.g. error control and MAC protocol. Notice that the test-bed is designed to evaluate the *energy-efficiency* of designs, and is not designed for *low-power*. The actual implementation of the test-bed is therefore primarily designed to be *flexible*, and suitable for experimenting with various design alternatives. Because of this, the implementation test-bed used various flexible, but certainly not low-power components (i.e. we have used Xilinx FPGAs). Figure 8 shows a photograph of the network interface implementation.



Figure 8: Network interface implementation.

We use a WaveLAN modem as the physical layer. The WaveMODEM is a RF module that converts a serial transmit data stream from the host into Radio Frequency (RF) modulated signals. The raw data rate is 2 Mb/s. The modem provides the basic functionality to send and receive

frames of data. It does not include a MAC protocol, but provides signalling information like carrier sense.

Table 1 presents measurements of the power consumption of the testbed network interface. In it, we distinguish the power consumption of the wireless modem and the total power consumption (i.e. including the network interface).

Table 1: Power consumption testbed network interface.

Operating mode	Power modem [mW]	Total power [mW]
power off	0	285
sleep	35	320
idle	1325	1610
receive	1345	1630
transmit	1380	1665

In the analysis made we assumed that the transitions of the operating modes of the modem were instantaneous with respect to the power consumption. However, this is not true in practical implementations. In order to evaluate whether our assumptions were valid and adequate, we performed several measurements on our testbed. Figure 9 shows three transitions of the operating mode: sleep to idle, idle to sleep, and power off to sleep.

As can be seen in Figure 9, a sleep to idle transition consumes significantly more energy than in the idle steady state for approximately 796 μ s. On the other hand, this might be compensated by the energy savings during the idle to sleep transition. The figure also shows that the power-off to sleep mode transition already reaches the power consumption of the sleep steady state within 1700 μ s. This is thus well within the 200 ms that is required for the modem to reach its operational (sleep) state. These measurements thus show that the assumptions made for the analysis are approximately valid and adequate.




Figure 9: Modem operating state transitions.

V. CONCLUSIONS

In this paper we have analysed a TDMA MAC traffic scheduling strategy that is aimed for a high energy efficiency. The mobile grouping strategy reduces the number of operating mode transitions between transmitting, receiving, idle, and sleep, and minimises the on time of the mobile transceivers. The traffic scheduler schedules all traffic according to the QoS requirements and tries to minimise the number of transitions the mobile has to make. It schedules the traffic of a mobile such that all downlink and uplink connections are grouped into packets, taking into account the limitations imposed by the QoS of the connections. Traditionally, the MAC scheduler uses a frame that has an uplink phase in which mobiles have to transmit, and a downlink phase in which mobiles can receive data.

We have shown that with our scheduling strategy the mobile can sleep for a significant longer time than with the traditional scheduling strategies, and can thus save a considerable amount of energy. The disadvantage is that the channel efficiency is lower. However, we believe that for battery-powered mobile multimedia computing devices, performance sufficiency (using a QoS framework) and energy efficiency will become the predominant requirements for wireless communication.

With small numbers of mobiles, mobile grouping has a small overhead while allowing a large sleep period. The current trend in mobile multimedia computing is to have ever smaller transmission areas (pico-cellular systems). This not only saves energy because the transmitters can be low powered, it also provides a high aggregate bandwidth since it needs to be shared with only several mobiles. Mobile grouping is particularly suited for these small area systems, because the number of mobiles is relatively small, and there is sufficient bandwidth available.

We have implemented a highly adaptive network interface and a MAC protocol that is based on mobile grouping. It provides support for diverse traffic types and QoS while achieving a good energy efficiency of the wireless interface of the mobile. The scheduler of the base station is responsible for providing the connections on the wireless link the required QoS and tries to minimise the amount of energy spend by the mobile. Most of the resulting energy waste comes from the relatively long transition times between the various operating modes of current wireless radio's. Minimising these transition times in future radio designs will be beneficial and will further reduce the energy consumption significantly.

REFERENCES

- [1] Akyildiz I.F., McNair J., Martorell L.C., Puigjaner R., Yesha Y.: "Medium Access Control protocols for multimedia traffic in wireless networks", *IEEE Network*, pp.39-47, July/August 1999.
- [2] Borriss, M. "QoS support in ATM and selected protocol implementations", *technical report TU Dresden, IBDR*, <http://www.inf.tu-dresden.de/~mb14/atm.html>, Oct. 1995.
- [3] Chen, et al. "Comparison of MAC Protocols for Wireless Local Networks Based on Battery Power Consumption", *IEEE Infocom'98*, San Francisco, USA, pp. 150-157, March 1998.
- [4] Chockalingam, A., Zorzi, M.: "Energy consumption performance of a class of access protocols for mobile data networks", *VTC'98*, Ottawa, Canada, May 1998.

- [5] Choi S., Shin K.G.: "A cellular wireless local area network with QoS guarantees for heterogeneous traffic", *Mobile networks and applications* 3, pp. 89-100, 1998.
- [6] Colombo G., Lenzini L., Mingozi E., Cornaglia B., Santaniello R.: "Performance evaluation of PRADOS: a scheduling algorithm for traffic integration in a wireless ATM network", *Proceedings of the fifth annual ACM/IEEE international conference on mobile computing and networking (MobiCom'99)*, pp. 143-150, August 1999.
- [7] ETSI: "High Performance Radio Local Area Network (HIPERLAN)", *draft standard ETS 300 652*, March 1996.
- [8] Havinga P.J.M., Smit G.J.M., Bos M.: "Energy efficient wireless ATM design", *ACM/Baltzer Journal on Mobile Networks and Applications (MONET), Special issue on Wireless Mobile ATM technologies, Vol. 5, No 2., 2000*.
- [9] Hettich A., Evans D., Du Y., Lott M., Fifield R.: "Fast uplink signalling for an ATM radio interface using energy burst with random access", *proceedings wmATM'99*, pp.167-176, June, 1999.
- [10] IEEE, "Wireless LAN medium access control (MAC) and physical layer (PHY) Spec." P802.11D5, *Draft Standard IEEE 802.11*, May 1996.
- [11] Lettieri P., Schurgers C., Srivastava M.B.: "Adaptive link layer strategies for energy efficient wireless networking", ACM WINET.
- [12] Lettieri, P., Srivastava, M.B.: "Adaptive Frame Length Control for Improving Wireless Link Throughput, Range, and Energy Efficiency", *IEEE Infocom'98*, San Francisco, USA, pp. 307-314, March 1998.
- [13] Makrakis D.M., Mander R.S., Orozco-Barbosa L., Papantoni-Kazakos P.: "A spread-slotted random-access protocol with multi-priority for personal and mobile communication networks carrying integrated traffic", *Mobile Networks and Applications* 2, pp.325-331, 1997.
- [14] Mangione-Smith, B.: "Low power communications protocols: paging and beyond", Low power symposium 1995, <http://www.icsl.ucla.edu/~billms/Publications/pagingprotocol.pdf>.
- [15] McKenna D.: "Mobile platform benchmarks, a methodology for evaluating mobile computing devices", *Transmeta Corporation*, January 2000, <http://www.transmeta.com>.
- [16] Moorman, J.R., Lockwood J.W.: "Multiclass priority fair queuing for hybrid wired/wireless quality of service support", *Proceedings of the second ACM international workshop on Wireless Mobile Multimedia (WoWMoM'99)*, pp. 43-50, August 1999.
- [17] Shakkottai S., Srikant R.: "Scheduling real-time traffic with deadlines over a wireless channel", *Proceedings of the second ACM international workshop on Wireless Mobile Multimedia (WoWMoM'99)*, pp. 35-42, August 1999.
- [18] Sivalingam, K.M., Chen J.C., Agrawal, P., Srivastava, M.B.: "Design and analysis of low-power access protocols for wireless and mobile ATM networks", *Journal on special topics in mobile networking and applications (MONET)*, June 1998.
- [19] Sivalingam, K.M., Srivastava, M.B. Agrawal, P.: "Low power link and access protocols for wireless multimedia networks", *Proceedings IEEE Vehicular Technology Conference*, Phoenix, AZ, pp. 1331-1335, May 1997.
- [20] Stemm, M. et al.: "Reducing power consumption of network interfaces for hand-held devices", *Proceedings MoMuc-3*, 1996.
- [21] Su W., Gerla M.: "Bandwidth allocation strategies for wireless ATM networks using predictive reservation", *IEEE Globecom '97*, 1997.
- [22] WaveMODEM 2.4 GHz Data Manual, Release 2, AT&T 1995.
- [23] Woesner H., Ebert J., Schläger M., Wolisz A.: "Power-saving mechanisms in emerging standards for wireless LANs: The MAC level perspective", *IEEE Personal Communications*, Vol. 5, No. 3, June 1998.
- [24] Zorzi, M., Rao, R. R.: "Error control and energy consumption in communications for nomadic computing", *IEEE transactions on computers*, Vol. 46, pp. 279-289, March 1997.