Energy-efficient wireless networking for multimedia applications

(Wireless Communications and Mobile Computing, Wiley, 2001: 1:165-184)

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 identify the most prominent problems of wireless multimedia networking and present several state-of-the-art solutions with a focus on energy efficiency. Three key problems in networked wireless multimedia systems are 1) the need to maintain a minimum quality of service over time-varying channels, 2) to operate with limited energy resources, and 3) to operate in a heterogeneous environment. We identify two main principles to solve these problems. The first principle is that energy efficiency should involve all layers of the system. Second, Quality of Service is an essential mechanism for mobile multimedia systems not only to give users an adequate level of service, but also as a tool to achieve an energy-efficient system. Due to the dynamic wireless environment, adaptability of the system will be a key issue in achieving this.*

***Keywords** – energy efficiency, wireless networking, mobile computing, quality of service.*

1 Introduction

Advances in technology enable portable computers to be equipped with wireless interfaces, allowing networked communication even while on the move. Whereas today's notebook computers and personal digital assistants (PDAs) are self contained (*introvert*) tomorrow's networked mobile computers are part of a greater computing infrastructure (*extrovert*). Key problems are that these portable wireless network devices need to handle multimedia traffic in a dynamic and heterogeneous wireless environment, and the need to operate with limited energy resources.

Wireless communication is much more difficult to achieve than wired communication because the surrounding environment interacts with the signal, blocking signal paths and introducing noise and echoes. As a result wireless connections have a lower quality than wired connections: lower bandwidth, less connection stability, higher error rates, and, moreover, a highly varying quality. They need to be able to operate in environments that may change drastically – in short term as well as in long term – in available resources and available services. These factors can in turn increase communication latency due to retransmissions, can give largely varying throughput, and incur a high energy consumption.

Wireless networking is a broad area, and has many applications ranging from voice communication (cellular phones) to high performance multimedia networking. In this paper we are somewhat biased towards multimedia traffic, as it is expected that the new generation of wireless networks will carry diverse types of multimedia traffic.

1.1 Key problems of wireless multimedia networking

Three key problems in wireless multimedia networking are 1) to operate with limited energy resources, 2) the need to maintain quality of service (throughput, delay, bit error rate, etc) over time-varying channels, and 3) to operate in a heterogeneous environment.

Energy-efficiency – Portable wireless devices have severe constraints on the size, the energy consumption, and the communication bandwidth. Moreover, it is expected that these devices will be multimedia oriented, and need to handle many different classes of data traffic over a limited bandwidth wireless connection, including delay sensitive, real-time traffic such as speech and video. More extensive and continuous use of network services will only aggravate this problem since communication consumes relatively much energy. Unfortunately, the rate at which battery performance improves (in terms of available energy per unit size or weight) is fairly slow, despite the great interest generated by the booming wireless business. Aside from major breakthroughs it is doubtful that significant reduction of battery size and weight can be expected in the near future. The energy consumption these devices need for communication and computation will limit the *functionality* of the mobiles.

The way out is *energy efficiency:* doing more work with the same amount of energy. The art of low-power design used to be a narrow speciality in analog circuit design. Nowadays, it is appearing in many layers of a system.

Quality of Service – A Quality of Service model provides the basis for modern high-bandwidth and real-time multimedia applications like teleteaching and video conferencing. The notion of QoS service originally stems from communication, but because of its potential in the allocation of all scarce resources, it has found its way into other domains, e.g. operating systems [25].

Heterogeneity – In contrast to most stationary computers, mobile computers encounter heterogeneous network connections. As they leave the range of one network transceiver they switch to another. In different places they may experience different network qualities. There may be places where they can access multiple transceivers, or even may concurrently use wired access. The interface may also need to change access protocols for different networks, for example when switching from wireless LAN coverage in an office to cellular coverage in a city. This heterogeneity makes mobile computing more complex than traditional networking.

These three problem-areas that are characteristic for future mobile networking are strongly correlated. [23]. Wireless network protocols typically address network performance metrics such as throughput, efficiency, fairness and packet delay. In this paper we address the additional issue of *energy efficiency* of the wireless network protocols [1]. Considerations of energy efficiency are fundamentally influenced by the trade-off between energy consumption and achievable Quality of Service. The dynamic communication and application environment is an extra challenge, which might be solved using QoS principles. The aim in mobile multimedia networking is to meet the required QoS, while minimising the required amount of energy. To deal with the dynamic variations in networking and computing resources gracefully, both the mobile computing environment and the applications that operate in such an environment need to adapt their behaviour depending on the available resources including the batteries. Current research on several aspects of wireless networks (like error control, frame-length, access scheduling) indicate that continually adapting to the current condition of the wireless link have a substantial impact on the energy-efficiency of the system [8][10][33].

1.2 *Principles of energy-efficient wireless networking*

In this paper we will discuss a variety of energy reduction approaches that can be used for building an energy-efficient mobile system, and show the relationship with *multimedia* and the *dynamic environment*. In this paper the following main principles are identified:

1. *Involve all layers*. Energy efficiency is an issue involving all layers of the system, its physical layer, its communication protocol stack, its system architecture (Section 2.1), its operating system, and the entire network (Section 4).
2. *Quality of Service* is an essential mechanism for mobile multimedia systems not only to give users an adequate level of service, but also as a tool to achieve an energy efficient system.

QoS support in wireless networks involves several considerations beyond those addressed in earlier work on conventional wireline networks. In traditional networks, based on fixed terminals and high-quality/high-capacity links it is feasible to provide 'hard' QoS guarantees to users. However, in a mobile environment, mobility and the need for efficient resource utilisation require the use of a 'soft' QoS model [44]. The minimum QoS requirements for multimedia applications has a wide dynamic range depending on the user's quality expectations, application usage models, and application' tolerance to degradation. Users and applications require a certain QoS level. The system then operates in such a way that it will try

[1] In general, saving energy for the base station is not really an issue, as it is part of the fixed infrastructure and typically obtains energy from a mains outlet. However, since the current trend is to have ever smaller cell sizes, and the complexity of the base station is increasing, this issue might become more important in the future mainly because of economical and thermal reasons.

to satisfy these requirements, but never gives more quality than required and necessary. Due to the dynamic wireless environment, *adaptability* of the system will be a key issue in achieving this. This implicates several mechanisms that can be used to attain a high energy efficiency, e.g.:

- *Avoid useless activity*. This is the main driving force of for instance dynamic power management (Section 2.2), link layer protocols (Section 5) and adaptive error control (Section 6). Useless activity can be caused by various factors at all levels of the system (e.g. being unnecessary in a high power operational mode, applying error control to error-resilient data, trying to transmit a video frame that is already too old). If the operations would adapt to the required QoS and current environment, then energy-efficiency can be improved.
- *Scheduled operations*. This extents power management in the sense that communication is scheduled at appropriate time such that the differences in power states are exploited as much as possible (Section 4.2, and Section 5). This has a strong relation with the QoS model since timing constraints of multimedia connections are likely to be the limiting factors in the potential energy reduction.
- *Reduce the amount of data*. This is quite obvious and is applicable in all layers of the system. This relates to the trade-off between communication and computation (Section 3). Examples are adaptive error control that adapts its error coding according to the current channel conditions and the required quality, video transmission systems that adapt the quality to the expectations of the users, the available resources, and the channel conditions. Again, QoS can be used to determine whether it is really necessary to produce the data.

1.3 Outline

In this paper we give an overview of various aspects of energy efficient wireless networking with a focus on the lower layers of network protocol stack. Although some research has been done in this area, a lot of research issues remain open. Because it is not possible to go into depth, the intention of this paper is primarily to give insight in the field of energy efficient networking.

In Section 2 we provide the fundamentals of power management. Then, in Section 3 we review the main sources of energy consumption induced by the wireless channel. In Section 4 we provide an overview of mechanisms to reduce the energy consumption needed for communication in the network protocol stack, the operating system, and by decomposition. In Section 5 and Section 6 we will delve a bit deeper into the two most developed areas to reduce energy consumption: the *MAC Layer*, and *error-control*.

2 Power management

Traditionally, energy efficiency has been focussed on low-power techniques for VLSI design. As the issue of energy efficiency becomes more pervasive, the battle to use the bare minimum of energy will be fought on multiple fronts: semiconductor technology, circuit design, design automation tools, system architecture, operating system, and application design. Energy awareness is now appearing in the mainstream digital design community affecting all aspects of the design process. Eventually, the concern for low-power design will expand from devices to modules to entire systems, including application software and even to the user.

2.1 Low power system design

Most components are currently fabricated using CMOS technology. Main reason for this bias is that CMOS technology is cost efficient and inherently consuming less power than other technologies. The dominant factor of energy consumption (85 to 90%) in CMOS is *dynamic*. A first order approximation of the dynamic energy consumption of CMOS circuitry is given by the formu la:

$$P_d = C_{eff} V^2 f \qquad (1)$$

where P_d is the power in Watts, C_{eff} is the effective switch capacitance, V is the supply voltage, and f is the frequency of operations. The power dissipation arises from the charging and discharging of the circuit node capacitance found on the output of every logic gate. C_{eff} combines two factors C, the capacitance being charged/discharged, and the *activity weighting* α, which is the probability that a transition occurs.

$$C_{eff} = a\, C \qquad (2)$$

At lower levels energy consumption can thus be decreased by reducing the supply voltage, reducing the capacitive load and by reducing the switching frequency.

In general, a designer tries to make a system to be optimal for a certain application and environment. The designer has to select a particular algorithm, design or use an architecture that can be used for it, and determine various parameters such as supply voltage and clock frequency. However, energy efficiency in mobile systems is not only a one-time design problem that needs to be solved during the design phase. In a mobile system, power management extents the notion of hardware/software co-design, since we have to face a *much more dynamic application and communication environment*. When the system is operational, frequent adaptations to the system are required to obtain an energy efficient system that can fulfil the requirements imposed in terms of a general QoS model. This multi-dimensional design space offers a large range of possible trade-offs.

2.2 Dynamic power management

The essential characteristic of energy consumption for static CMOS circuits is that quiescent portions of a system dissipate a minimal amount of energy. Dynamic power management refers to the general class of techniques that manage the performance and throughput of a system based on its computational needs within the energy constraint [6]. Dynamic power management exploits periods of *idleness* caused by system under-utilisation. Especially in mobile systems, the utilisation is not constant and power management can be used effectively. It is common practice that designers focus on worst-case conditions, peak performance requirements and peak utilisation, which, however, is in practice only fully exploited during a small fraction of their operation.

Dynamic power management is based on deactivating functional units when they are not required. The main problems involved are the cost of shutting-down and restarting a module or component. Restarting induces an increase in latency (e.g. time to restore a saved CPU state, spin-up of a disk), and possibly also an increase in energy consumption (e.g. due to higher start-up current in disks). The two main questions involved are then: 1) when to shutdown, and 2) when to wake-up.

1. The time (and thus energy) that is required to determine when a module can be shut down (the so-called *inactivity threshold*) can be assigned statically or dynamically. In a *predictive* power management strategy the threshold is adapted according to the past history of active and idle intervals.
2. The other question is *when to wake-up*, where the typical policy is to wake up in response to a certain event such as user interaction or network activity. The problem with such a demand policy is that waking up takes time, and the extra latency is not always tolerable. Again, a predictive approach, where the system initiates a wakeup in advance of the predicted end of an idle interval, often works better.

Operating modes of a wireless interface

The wireless network interface of a mobile computer consumes a significant fraction of the total power [53]. Typically, the transceiver can be in five modes; in order of increasing energy consumption, these are off, sleep, idle, receive, and transmit (see Figure 1). 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 is still powered and ready to receive or transmit; in sleep mode, the transceiver circuitry is powered down, except sometimes for a small amount of circuitry listening for incoming transmissions [35]. The difference in the amount of energy consumed in these modes is significant.

Examples:

> *1) The power consumption of a WaveLAN modem when transmitting is typical 1675 mW, 1425 mW when receiving, and 80 mW when in sleep mode [56]. Increasing the sleep time period of the radio thus significantly improves the energy efficiency of the wireless network. Also important to notice is that the transition times between the operating modes can be quite high. In WaveLAN a transition time from sleep to idle takes 250 ms, and has during that period already the power consumption of the idle state [21]. Then, before the payload will be transmitted another 254 ms is required.*

2) A Bluetooth radio (Ericsson PBA 313 01/2 [2]) has similar characteristics: in sleep mode it consumes 100 mA, in idle 25 mA, in receive mode 52 mA, and in transmit mode 44 mA. From sleep to idle requires 110 ms, and from idle to transmit or receive mode typical 104 ms.

2.3 Energy efficiency

We define the *energy efficiency e* as the energy dissipation that is essentially needed to perform a certain function, divided by the actually used total energy dissipation.

$$e = \frac{\text{Essential energy dissipation for a certain function}}{\text{Actually used total energy dissipation}} \qquad (3)$$

The function to be performed can be very broad: it can be a limited function like a multiply-add operation, but it can also be the complete functionality of a network protocol.

Let us for example consider a medium access control (MAC) protocol that controls access to a wireless channel. The essential energy dissipation is the energy dissipation needed to transfer a certain amount of bits over the wireless channel, and the total actually used energy dissipation also includes the overhead involved in additional packet headers, error control, etc., but also 'physical' overhead induced by e.g. a frequency hopping scheme.

Note that the energy efficiency of a certain function is independent of the actual implementation, and thus independent of the issue whether an implementation is low-power. Low-power is generally closely related to the hardware, whereas energy-efficiency relates to the algorithm using the hardware. Thus, it is possible to have two implementations of a certain function that are built with different building blocks, of which one that has been build with power-hungry components has a high energy efficiency, but dissipates more energy than the other implementation which has a lower energy efficiency, but is built with low-power components.

3 Energy consumption in mobile systems

Several researchers have studied the power consumption pattern of mobile computers. However, because they studied different platforms, their results are not always in agreement. Laptops designers use several techniques to reduce energy consumption, primarily by turning devices off after a period of no use, or by lowering the clock frequency. Lorch reported that the energy use of a typical laptop computer is dominated by the backlight of the display, the disk and the processor [36]. Ikeda et al. observed that the contribution of the CPU and memory to power consumption has been on the rise the last few years [27]. Stemm et al. [53] concluded that the network interface consumes at least the same amount of energy as the rest of the system (i.e. a Newton PDA). Further, the fraction of energy consumed for networking by these mobiles, is only likely to increase as mobiles evolve towards a thin client network computer, and the communication traffic will increase. Another source of energy consumption is due to the fact that many high-performance network protocols require that all network access be through the operating system, which adds significant overhead to both the transmission path (typically a system call and data copy) and the receive path (typically an interrupt, a system call, and a data copy). This not only causes performance problems, but also incurs a significant energy consumption. Intelligent network interfaces can relieve this problem to some extend. To address the performance problem, several *user-level communication architectures* have been developed that remove the operating system from the critical communication path [7].

To make the wireless interfaces more energy efficient, algorithms embodying energy efficient protocols must be distributed across two or more wireless end-points [32]. This implies that the focus should be on the layers of the network stack through which the mobiles interact.

Even though it is difficult to compare these results because the measurements are made for different architectures, operating systems, communication interfaces, and benchmarks, there is a common pattern: there is no primary source of energy consumption, and the energy consumed for communication increases. The energy consumption is distributed over several devices and for several operations. The conclusion is that implementing an energy efficient system should involve all the functions in the system at all layers.

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.

WHAT WILL YOU BUILD? | **sales@docketalarm.com** | **1-866-77-FASTCASE**

fastcase®
Smarter legal research.