

Low power system design techniques for mobile computers

Paul J.M. Havinga, Gerard J.M. Smit

University of Twente, department of Computer Science
P.O. Box 217, 7500 AE Enschede, the Netherlands
e-mail: {havinga, smit}@cs.utwente.nl

Abstract

Portable products such as pagers, cordless and digital cellular telephones, personal audio equipment, and laptop computers are being used increasingly. Because these applications are battery powered, reducing power consumption is vital.

In this report we first give the properties of low power design and techniques to exploit them on the hardware level such as: minimize capacitance, avoid wasteful activity, and reduce voltage and frequency. We will then elaborate on low power system-design techniques in which the main themes are to avoid wasteful activity at system level and to exploit locality of reference.

Finally we review energy reduction techniques in the design of a wireless communication system, including system decomposition, communication and MAC protocols, and low power short range networks.

1 Introduction

The requirement of portability of hand-held computers and portable devices places severe restrictions on size and power consumption. Even though battery technology is improving continuously and processors and displays are rapidly improving in terms of power consumption, battery life and battery weight are issues that will have a marked influence on how hand-held computers can be used. These devices often require real-time processing capabilities, and thus demand high throughput. Power consumption is becoming the limiting factor in the amount of functionality that can be placed in these devices. More extensive and continuous use of network services will only aggravate this problem since communication consumes relatively much energy. Research is needed to provide intelligent policies for careful management of the power consumption while still providing the appearance of continuous connections to system services and applications.

The Moby Dick project

The technologies of PDA, digital cellular phone and smart card, when combined and integrated well, have the potential of replacing all of the things people have to carry around with them by one small device, the Pocket Companion. It is a small portable computer and wireless communications device that can replace cash, cheque book, passport, keys, diary, phone, pager, maps and possibly briefcases as well. The combination of an intelligent information system and a location system engenders many new types of applications, such as admission control, digital chequebook, paging, and an automatic diary that keeps track of where you were and with whom.

The Moby Dick project [Moby Dick 95] is a joint european project (Esprit Long Term Research 20422) to develop and define the architecture of a new generation of mobile hand-held com-

puters. The design challenges lie primarily in the creation of a single architecture that allows the integration of security functions, externally offered services, personality, and communication. Research issues include: security, energy consumption and communication, hybrid networks, data consistency, and environment awareness.

To support multimedia functionality for the intended applications of the Pocket Companion the system needs to have real-time properties. This however, does *not* imply maximal performance. The main theme is: *enough performance for minimal energy consumption*. This is in contrast to current research in computer systems which aims at the *highest performance*, and where energy consumption is of minor concern.

Background

Several researchers have studied the power consumption pattern of mobile computers. However, because they studied different platforms, their results are not always in line, and sometimes even conflicting. Lorch reported that the energy use of a typical laptop computer is dominated by the backlight of the display, the disk and the processor [Lorch 95]. Stemm et al. concluded that the network interface consumes at least the same amount of energy as the rest of the system (i.e. a Newton PDA) [Stemm 96]. If the computer is able to receive messages from the network even when it is 'off', the energy consumption increases dramatically. Ikeda et al. observed that the contribution of the CPU and memory to power consumption has been on the rise the last few years¹ [Ikeda 94]. Laptops use several techniques to reduce this energy consumption, primarily by turning them off after a period of no use, or by lowering the clock frequency. Some researchers proposed to replace the hard disk by a flash RAM.

Note that there is an inherent trade-off between energy consumption and performance, since low-power techniques have associated disadvantages. For instance, decreasing the CPU frequency can raise response time, and spinning down the disk causes a subsequent disk access to have a high latency.

Outline of the paper

With the increasing integration levels, energy consumption has become one of the critical design parameters. Consequently, much effort has to be put in achieving lower dissipation at all levels of the design process. While low-power components and subsystems are essential building blocks for portable systems, little effort has been directed towards dedicated low-power hardware architectures by considering the system as a whole. A system wide architecture is beneficial because there are dependencies between subsystems, e.g. optimization of one subsystem may have consequences for the energy consumption of other modules. In this paper we will discuss a variety of energy reduction approaches that can be used for building an energy efficient system.

Energy reduction techniques can be applied in all design levels of the system. First of all, we have to use components that use the latest developments in low power technology. Furthermore, as the most effective design decisions derive from the architectural and system level, a cautious design at these levels can reduce the power consumption considerable. However, it is not just a problem of the power-conscious hardware designer, but also involves careful design of the operating system and application programs. Furthermore, because the applications have direct knowledge of how the user is using the system, this knowledge must be penetrated into the power management of the system.

We first explore sources of energy consumption and show the basic techniques used to reduce the power dissipation. Then we give an overview of energy saving mechanisms at the system

1. The IBM Thinkpad laptops have shown an increase in the fraction of power consumption in the memory and CPU from 18% to 40% from 1992 to 1995.

and architectural level. Finally, we will show as an example the techniques used in the Moby Dick project in order to reduce energy consumption for communication at the architectural and system level.

2 Properties of low power design

Throughout this paper, we discuss ‘power consumption’ and methods for reducing it. Although they may not explicitly say so, most designers are actually concerned with reducing energy consumption. This is because batteries have a finite supply of energy (as opposed to power, although batteries also put limits on peak power consumption as well). Energy is the time integral of power; if power consumption is a constant, energy consumption is simply power multiplied by the time during which it is consumed. Reducing power consumption only saves energy if the time required to accomplish the task does not increase too much. A processor that consumes more power than a competitor may or not may not consume more energy for a certain program. For example, even if processor A’s power consumption is twice that of processor B, A’s energy consumption could actually be less if it can execute the same program more than twice as quickly as B.

2.1 Design flow

The design flow of a system constitutes of various levels of abstraction. When a system is designed with the emphasis on power optimization as a performance goal, then the design must embody optimization at all levels of the design flow. In general there are three levels on which energy reduction can be incorporated. The *system level*, the *architecture level*, and the *technological level*. For example, at the system level inactive modules may be turned off to save power. At the architectural level, parallel hardware may be used to reduce global interconnect and allow a reduction in supply voltage without degrading system throughput. At the technological level several optimisations can be applied at the gate level.

The system and architecture have to be designed targeted to the possible reduction of energy consumption at the gate level. An important aspect of the design flow is the relation and feedback between the levels. Figure 1 shows the general design flow of a system with some examples of how energy reduction can be obtained.

Given a design specification, a designer is faced with several different choices on different levels of abstraction. The designer has to select a particular algorithm, design or use an architecture that can be used for it, and determines various parameters such as supply voltage and clock frequency. This multi-dimensional design space offers a large range of possible trade-offs. The most influence on the properties of a design is obtainable at the highest levels. Therefore the most effective design decisions derive from choosing and optimizing architectures and algorithms at the highest levels. It has been demonstrated by several researchers [Sheng 92] that system and architecture level design decisions can have dramatic impact on power consumption. However, when designing a system it is a problem to predict the consequences and effectiveness of design decisions because implementation details can only be accurately modelled or estimated at the technological level and not at the higher levels of abstraction.

2.2 CMOS component model

Most components are fabricated using CMOS technology. The sources of energy consumption on a CMOS chip can be classified as static and dynamic power dissipation. *Static* energy consumption is caused by short circuit currents (P_{sc}), bias (P_b) and leakage currents (P_l). *Dynamic* energy consumption (P_d) is caused by the actual effort of the circuit to switch.

$$P = P_d + P_{sc} + P_b + P_l \quad (1)$$

<i>abstraction level</i>	<i>examples</i>
system	compression method energy manager scheduling medium access protocols system partitioning
architecture	communication error control parallel hardware hierarchical memories compiler
technological	asynchronous design clock frequency control reducing voltage reduce on-chip routing

Figure 1: General design flow and related examples for energy reduction

The contributions of this static consumption are mostly determined at the circuit level. During the transition on the input of a CMOS gate both p and n channel devices may conduct simultaneously, briefly establishing a short from the supply voltage to ground. This effect causes a power dissipation of approx. 10 to 15%. Also, lower operating voltages as being used nowadays, tend to reduce the short circuit component. While statically-biased gates are usually found in a few specialized circuits such as PLAs, their use has been dramatically reduced in CMOS design [Burd 95]. Leakage currents also dissipate static energy, but are also insignificant in most designs (less than 1%).

In general we can say that careful design of gates generally makes their power dissipation typically a small fraction of the dynamic power dissipation, and hence will be omitted in further analysis.

The dominant component of energy consumption (85 to 90%) is CMOS is therefore *dynamic*. A first order approximation of the dynamic power consumption of CMOS circuitry is given by the formula:

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

where P_d is the power in Watts, C_{eff} is the effective switch capacitance in Farads, V is the supply voltage in Volts, and f is the frequency of operations in Hertz [Lapsley 94]. The power dissipation arises from the charging and discharging of the circuit node capacitances found on the output of every logic gate. Every low-to-high logic transition in a digital circuit incurs a voltage change ΔV , drawing energy from the power supply. C_{eff} combines two factors C , the capacitance being charged/discharged, and the activity weighting α , which is the corresponding probability that a transition occurs.

$$C_{eff} = \alpha C \quad (3)$$

A designer at the technological and architectural level can try to minimize the variables in these equations to minimize the overall energy consumption. However, as will be shown in the next sections, power minimization is often a subtle process of trade-offs.

3 Reducing power at the technological and architectural level

The equations 1 and 2 suggest that there are essentially four ways to reduce power:

- reduce the capacitive load C_{eff}
- reduce the supply voltage V ,
- reduce the switching frequency f ,
- reduce the activity α

3.1 Minimize capacitance

Energy consumption in CMOS circuitry is proportional to capacitance. Therefore a technique that can be used to reduce energy consumption is to *minimize the capacitance*. This can not only be reached at the technological level, but much profit can be gained by an architecture that exploits locality of reference and regularity. However, note that a sole reduction in chip area, that typically translates into reduced capacitances, could lead to an inefficient design. For example, a power efficient architecture that occupies a larger area can reduce the overall energy consumption, e.g. by exploiting locality in a parallel implementation.

Connections to external components, such as external memory, typically have much greater capacitance than connections to on-chip resources. As a result, accessing external memory can increase energy consumption. So, a way to reduce capacitance is to reduce external accesses and optimize the system by using on-chip resources like caches and registers. Furthermore, use few external outputs, and have them switch as infrequently as possible.

Routing capacitance is the main cause of the limitation in clock frequency. Circuits that are able to run faster can do so because of a lower routing capacitance. Consequently, they dissipate less power at a given clock frequency. So, energy reduction can be reached by optimizing the clock frequency of the design even if the resulting performance is far in excess of the requirements [Xilinx 95].

3.2 Reduce voltage and frequency

One of the most effective ways of energy reduction of a circuit at the technological level is to reduce the supply voltage, because the energy consumption drops quadratically with the supply voltage. For example, reducing a supply voltage from 5.0 to 3.3 volts (a 44% reduction) reduces power consumption by about 56%. As a result, most processor vendors now have three volt versions. The problem that then arises is that lower supply voltages will cause a reduction in performance. In some cases, low voltage versions are actually five volt parts that happen to run at the lower voltage. In such cases the system clock must typically be reduced to ensure correct operation. Therefore any such voltage reduction must be balanced against any performance drop. To compensate and maintain the same throughput, extra hardware can be added. This is successful up to the point where the extra control, clocking and routing circuitry adds too much overhead [Rabaey 94]. In other cases, vendors have introduced ‘true’ three volt versions of their processors that run at the same speed as their five volt counterparts.

The variables *voltage and frequency* have a trade-off between delay and energy consumption. Reducing clock frequency f alone does not reduce energy, since to do the same work the system must run longer. As the voltage is reduced, the delay increases. A common approach to power reduction is to first increase the performance of the module - for example by adding parallel hardware -, and then reduce the voltage as much as possible so that the required performance is still reached (figure 2). Therefore, a major theme in many power optimization

techniques is to optimize the speed and lower the critical path, so that the voltage can be reduced. However, these techniques often translate in larger area requirements, hence there is a new trade-off between area and power.




Figure 2: Impact of voltage scaling and performance to total power consumption

Weiser et al. [Weiser 94] have proposed a system in which the clock frequency and operating voltage is varied dynamically under control of the operating system while still allowing the processor to meet its task completion deadlines. They point out that in order to operate properly at a lower voltage, the clock rate must be simultaneously reduced.

3.3 Avoid wasteful activity

The activity weighting α of equation 3 can be minimized by *avoiding wasteful activity*. There are several techniques to achieve this. Because CMOS power consumption is proportional to the clock frequency, dynamically turning off the clock to unused logic or peripherals is an obvious way to reduce power consumption [Larri 96], [Intel486SX]. Control can be done at the hardware level or it can be managed by the operating system or the application. Some processors and hardware devices have *sleep or idle modes*. Typically they turn off the clock to all but certain sections to reduce power consumption. While asleep, the device does no work. A wake-up event wakes the device from the sleep mode. Devices may require different amounts of time to wake up from different sleep modes. For example, many 'deep sleep' modes shut down on-chip oscillators used for clock generation. A problem is that these oscillators may require microseconds or sometimes even milliseconds to stabilize after being enabled. So, it is only profitable to go into deep sleep mode when the device is expected to sleep for a relatively long time.

The technique of dynamically turning off the clock can also be applied to the design of synchronous finite state machines (FSM). For example [Kogst 97] uses gated clocks in FSM designs to disable the state transition of so called self-loops.

Energy consumption is proportional to the frequency at which signals change state from 0 to 1 or vice-versa and to the capacitance on the signal line. This is true for every signal path in a system, whether it is a clock signal, a data pin, or an address line. This implies that power consumption can be reduced by carefully minimizing the number of transitions. A correct choice of the *number representation* can have a large impact on the switching activity. For example, program counters in processors generally use a binary code. On average, two bits are changed for each state transition. A Gray code, in which typically a single bit changes, can give interesting energy savings. However, according to [Piguet 96], a Gray code incrementer

requires more transistors to implement than a ripple carry incrementer. Therefore a combination can be used in which only the most frequently changing LSB bits use a Gray code.

Another way to avoid wasteful activity is by applying an *asynchronous design methodology*. CMOS is a good technology for low power as gates only dissipate energy when they are switching. Normally this should correspond to the gate doing useful work, but unfortunately in a synchronous circuit this is not always the case. Many gates switch because they are connected to the clock, not because they have new inputs to process. The biggest gate of all is the clock driver that must distribute a clock signal evenly to all parts of a circuit, and it must switch all the time to provide the timing reference even if only a small part of the chip has something useful to do. A synchronous circuit therefore wastes power when particular blocks of logic are not utilized, for example, to a floating point unit when integer arithmetic is being performed.

Asynchronous circuits though are inherently data driven and are only active when performing useful work. Parts of an asynchronous circuit that receives less data will automatically operate at a lower average frequency. Unfortunately, extra logic is required for synchronization, so asynchronous circuits are larger than synchronous circuits.

Reversible logic [Merkle 93] or adiabatic logic tries to reduce energy consumption by not erasing information. Today's computers erase a bit of information every time they perform a logic operation. These logic operations are therefore called 'irreversible'. We can improve the efficiency of erasing information with conventional methods, such as used in large cache systems. An alternative is to use logic operations that do not erase information. These are called reversible logic operations, and in principle they can dissipate arbitrarily little heat. To achieve a completely reversible system (which erases no bits at all) is very difficult.

4 Low power system level design

In the previous section we have explored sources of energy consumption and showed the low level design techniques used to reduce the power dissipation. In this section we will concentrate on these techniques at *system level* and the relevance for low power system design.

The two main themes that can be used for energy reduction at system level are:

- avoid wasteful activity, and
- exploit locality of reference.

4.1 Avoid wasteful activity

In this section we will show several approaches that can be applied at the system level to avoid wasteful activity.

1. *Scheduling and caching*

In a system *scheduling* is needed when multiple functional units need to access the same object. Scheduling is used by the operating system to provide each unit a share of the object in time. Scheduling is applied at several parts of a system for processor time, communication, disk access, etc. Currently scheduling is performed on criteria like priority, latency, time requirements etc. Power consumption is in general only a minor criterion for scheduling, despite the fact that much energy could be saved.

Caching can be used to provide the temporary buffer that is needed for scheduling, but can also enable locality of reference. Locality of reference enables the partitioning of the memory into smaller memories. Smaller memories not only consume less energy due to a reduced switching capacitance, but moreover, this architecture can utilize a pipelined structure that

greatly reduces the critical path.

We will now show several possible mechanisms in which an energy aware scheduling and caching system can be beneficial.

- *Processor time scheduling*

Most systems spend only a fraction of the time performing useful computation. The rest of the time is spent idling. The operating systems energy manager should track the periods of computation, so that when an idle period is entered, it can immediately power off major parts of the system that are no longer needed [Burd 95]. Since all power-down approaches incur some overhead, the task of an energy aware scheduler can be to collect requests for computation and compact the active time-slots into bursts of computation.




Figure 3: Power consumption in time of a typical processor system.

Weiser et al. [Weiser 94] have proposed a system that reduces the cycle time of a processor for power saving, primarily by allowing the processor to use a lower voltage. For background and high latency tolerable tasks, the supply voltage can be reduced so that just enough throughput is delivered, which minimizes energy consumption. By detecting the idle time of the processor, they can adjust the speed of the processor while still allowing the processor to meet its task completion deadlines. Suppose a task has a deadline of 100 ms, but it will only take 50 ms of CPU time when running at full speed to complete. A normal system would run at full speed for 50 ms, and the idle for 50 ms in which the CPU can be stopped. Compare this to a system that runs the task at half speed, so that it completes just before its deadline. If it can also reduce the voltage by half, then the task will consume a quarter of the energy of the normal system. This is because the same number of cycles are executed in both systems, but the modified system reduces energy use by reducing the operating voltage.

They classified idle periods into 'hard' and 'soft' events. Obviously, running slower should not allow requests for a disk block to be postponed. However, it is reasonable to slow down the response to a keystroke, such that processing of one keystroke finishes just before the next. Another approach is to classify jobs or processes into classes like background, periodic and foreground. With this sort of classification the processor can run at a lower speed when executing low priority background tasks only.

- *File system*

In the operating system's file system a scheduler can try to collect disk operations in a cache and postpone low priority disk I/O only until the hard drive is running already or has enough data.

- *Communication*

The limited bandwidth of current wireless networks may cause needless energy consump-

tion. The medium access protocols of a wireless system can be adapted and tuned for low energy consumption. A base station that divides the available bandwidth equally among 10 mobiles causes them to consume 10 times as much power (100 times as much power total!) compared to a base station that uses a TDMA protocol to coordinate delivery of data to receivers. An example of an energy aware MAC protocol is LPMAC [Mangione-Smith 96]. It uses peer-to-peer wireless networking to reduce power consumption across the entire network. One host or base station is responsible for traffic scheduling and tries to minimize power consumption by minimizing the number of state-transitions. In this way a mobile is allowed to doze (and power off the receiver) as long as the network interface is reactivated at schedule time to receive the data at full speed.

In the higher level protocols of a communication system scheduling is used to control the transmission of messages. In a situation with varying and multiple network connectivity it may be wise to prefetch some information or postpone the actual transmission until a more power economic network is available. For example an application can schedule times to turn on the processor when it is connected to a wired network so that the application can download information from the network when it consumes less energy or does not need its batteries.

2. Energy manager

Power down of unused modules is a commonly employed approach for energy reduction. The division of the system into modules must be such that the modules must provide a clustered functionality. For example, locality of reference can be detected and exploited during memory assignment to induce an efficient and effective power down of large blocks of memory.

To take advantage of low-power states of devices, either the operating system needs to direct (part of) the device to turn off (or down) when it is predicted that the net savings in power will be worth the time and energy overhead of turning off and restarting, or the modules use a demand- or data-driven computation to automatically eliminate switching activity of unused modules. The device or system will enter the sleeping state when it is idle or when the user indicates to do so.

In order to achieve this, changes must be made to current designs for hardware, drivers, firmware, operating system, and applications. One of the key aspects is to move power management policy decisions and coordination of operations into the operating system. The operating system will control the power states of devices in the system and share this information with applications and users. This knowledge can be used and integrated in the Quality of Service model of the system.

Applications play the most critical role in the user's experience of a power-managed system. In traditional power-managed systems, the hardware attempts to provide automatic power management in a way that is transparent to the applications and users. This results in some legendary user problems such as screens going blank during video or slide-show presentations, annoying delays while disks spin up unexpectedly, and low battery life because of inappropriate device usage. Because the applications have direct knowledge of how the user is using the system to perform some function, this knowledge must be penetrated into the power management decision-making in the system in order to prevent these kinds of user problems.

Obviously, careless application's use of the processor and hard disk drastically affects battery life time. For example, performing non-essential background tasks in the idle loop prevents the processor from entering a low power state (see for example [Lorch 96]). So, it is not suffi-

cient to have the system to be low power, but the applications running on the system have to be written energy aware as well.

3. Code and algorithm transformation

As much of the power consumed by a processor is due to the fetching of instructions from memory, high code density can reduce energy consumption. However, this only works well when the execution cycle is not (much) longer. Today, the cost function in most compilers is either speed or code size. An energy aware compiler has to make a trade-off between size and speed in favour of energy reduction. The energy consumed by a processor depends on the previous state of the system and the current inputs. Thus, it is dependent on instruction choice and instruction ordering. Reordering of instructions can reduce the switching activity and thus overall energy consumption. However, it was found not to have a great impact [Tiwari 94].

At the algorithm level functional pipelining, retiming, algebraic transformations and loop transformations can be used [Mehra 96]. The system essential power dissipation can be estimated by a weighted sum of the number of operations in the algorithm that has to be performed [Chandrakasan 95]. The weights used for the different operations should reflect the respective capacitance switched. The size and the complexity of an algorithm (e.g. operation counts, word length) determine the activity. Operand reduction includes common sub-expression elimination, dead code elimination etc. Strength reduction can be applied to replace energy consuming operations by a combination of simpler operations (for example by replacing multiplications into shift and add operations). Drawbacks of this approach are that it introduces extra overhead for registers and control, and that it may increase the critical path.

4.2 Exploit locality of reference

The implementation dependent part of the power consumption of a system is strongly related to a number of properties that a given system or algorithm may have [Rabaey 95]. The component that contributes a significant amount of the total energy consumption is the *interconnect*. Experiments have demonstrated that in designs, about 10 to 40% of the total power may be dissipated in buses, multiplexers and drivers. This amount can increase dramatically for systems with multiple chips due to large off-chip bus capacitance. The power consumption of the interconnect is highly dependent on algorithm and architecture-level design decisions. Two properties of algorithms are important for reducing interconnect power consumption: locality and regularity.

Locality relates to the degree to which a system or algorithm has natural isolated clusters of operation or storage with a few interconnections between them. Partitioning the system or algorithm into spatially local clusters ensures that the majority of the data transfers take place within the clusters and relatively few between clusters. The result is that the local buses are shorter and more frequently used than the longer highly capacitive global buses. Locality of reference can be used to partition memories. Current high level synthesis tools are targeted to area minimization. For power reduction, however, it is better to minimize the number of accesses to long global buses and have the local buses be accessed more frequently. In a direct implementation targeted at area optimization, hardware sharing between operations might occur, destroying the locality of computation. An architecture and implementation should preserve the locality and partition and implement it such that hardware sharing is limited. The increase in the number of functional units does not necessarily translate into a corresponding increase in the overall area and energy consumption since (1) localization of interconnect allows a more compact layout and (2) fewer (access to) multiplexers and buffers are

needed.

Regularity in an algorithm refers to the repeated occurrence of computational patterns. Common patterns enable the design of less complex architecture and therefore simpler interconnect structure (buses, multiplexers, buffers) and less control hardware. These techniques have been exploited by several researchers [e.g. Mehra 96 and Rabaey 95], but mainly in the DSP domain where a large set of applications inherently have a high degree of regularity.

We will now show two mechanisms that exploit locality of reference to reduce energy consumption.

1. Application specific modules

Localization reduces the communication overhead in processors and allows the use of minimum sized transistors, which results in drastic reductions of capacitance. Pipelining and caching are examples of localization. Another way to reduce data traffic is to integrate a processor in the memory, as for example proposed by Patterson in intelligent RAM [Patterson 96, McGaughy 96].

At system level locality can be applied to divide the functionality of the system into dedicated modules. When the system is decomposed out of application-specific coprocessors the data traffic can be reduced, for instance because unnecessary data copies are removed. For example, in a system where a stream of video data is to be displayed on a screen, the data can be copied directly to the screen memory, without going through the main processor.

Furthermore, processors often have to perform tasks for which they are not ideally suited. Although they can perform such tasks, they may still take considerably longer, and might be more energy demanding, than a custom hardware implementation. Application-specific integrated circuits (ASICs) or dedicated processors placed around a standard processor can offer an alternative approach. A system designer can use the processor for portions of algorithms for which it is well suited, and craft an application-specific coprocessor (e.g. custom hardware) for other tasks. This is a good example of the difference between power and energy: although the application-specific coprocessor may actually consume more power than the processor, it may be able to accomplish the same task in far less time, resulting in a net energy savings.

By careful repartitioning a system, not only the power consumption can be reduced but the performance is actually improved as well [Mangione-Smith 96].

2. Hierarchical memory systems

Hierarchical memory systems can be used in a processor system to reduce energy consumption. The basic idea is to store a frequently executed piece of code or frequently used data in a small memory close to or in the processor (a cache). As most of the time only a small memory is read, the energy consumption is reduced.

Memory considerations must also be taken into account in the design of any system. By employing an on-chip cache significant power reductions together with a performance increase can be gained.

Apart from caching data and instructions at the hardware level, caching is also applied in the filesystem of an operating system. The larger the cache, the better performance. Energy consumption is reduced because data is kept locally, and thus there is less data traffic. Furthermore, the energy consumption is reduced because less disk and network activity is required.

The compiler can have impact on power consumption by reducing the number of instructions with memory operands. The most energy can be saved by a proper utilization of registers

[Tiwari 94]. It was also noted that writes consumes more energy, because a processor with a write-through cache (like the Intel 486) always causes an off-chip memory operation

5 Energy reduction techniques applied in communication

The wireless network interface of a mobile computer consumes a significant fraction of the total power [Stemm 96]. Measurements show that on typical applications like a web-browser or e-mail, the energy consumed when the interface is on and idle is more than the cost of receiving packets. This is because the interface is generally longer idle than actually receiving packets. Furthermore, switching between states (i.e. off, idle, receiving, transmitting) consumes time and energy.

There are a number of techniques that can be applied for the system and architecture design of a wireless communication system that can be used to reduce power consumption. A careful design of *all* network layers is required. In this section we will elaborate on the techniques used in the Moby Dick project for reducing energy consumption in communication. There are several ways to achieve this: e.g. by system decomposition, by using hybrid networking with low power short range networks, and by applying power aware MAC protocols.

5.1 System decomposition

In normal systems much of the network protocol stack is implemented on the main processor. Thus, the network interface and the main processor must always be on for the network to be active. Because almost all data is transported through the processor, performance and energy consumption is a significant problem.

In a communication system locality of reference can be exploited by decomposition of the network protocol stack and cautious management of the data flow. This can reduce the energy consumption for several reasons:

- First, when the system is constructed out of independent components that implement different layers of the communication stack, unnecessary data copies between successive layers of the protocol stack are eliminated. This eliminates wasteful data transfers over the global (!) bus, and thus saves much dissipation in buses, multiplexers and drivers.
- Secondly, dedicated hardware can do basic signal processing and can move merely the necessary data directly to its destination, thus keeping data copies off of the system bus. Moreover, this dedicated hardware might do its tasks much more energy efficient than a general purpose processor.
- Finally, a communications processor can be applied to handle most of the lower levels of the protocol stack, thereby allowing the main processor to sleep for extended periods of time without affecting system performance or functionality.

This decomposition can also be applied beyond the system level of the portable: in our approach certain functions of the system can be migrated from the portable system to a remote server that has plenty of energy resources. This remote server handles those functions that can not be handled efficiently on the portable machine. For example, a base station could handle parts of the network protocol stack in lieu of the mobile. The remote server has a private dedicated communication with the mobile so that the mobile units can use an internal, light weight, protocol to communicate with the base station rather than TCP/IP or UDP. The net result is saving in code and energy. In such a system it is also efficient to adapt the protocols for the specific environment it is used in. For example, wireless networks have a much higher *error rate* than the normal wired networks. In the presence of a high packet error rate, some network protocols (such as TCP) may overreact to packet losses, mistaking them for

congestion. This leads to backing off to a lower transfer rate which increases the energy consumption because it leads to a longer transfer time. Any protocol that leaves a mobile receiver idle unnecessarily wastes energy. The limitations of TCP can be overcome by a more adequate congestion control during packet errors [Rizzo 97]. Buffering of data on a base station can be used to perform only local retransmissions that are caused by errors in the wireless network.

In order to save energy a normal mode of operation of the mobile will be a sleep or power down mode. To support full connectivity while being in a deep power down mode the network protocols need to be modified. Store-and-forward schemes for wireless networks, such as the IEEE 802.11 proposed sleep mode, not only allow a network interface to enter a sleep mode but can also perform local retransmissions not involving the higher network protocol layers. However, such schemes have the disadvantage of requiring a third party, e.g. a base station, to act as a buffering interface.

In the higher level protocols of a communication system caching and scheduling is used to control the transmission of messages. In a situation with varying and multiple network connectivity it may be wise to prefetch some information or postpone the actual transmission until the quality of the connection is better, or until another, more power economic, network is available. An application can for example schedule the times to turn on the processor when it is connected to a wired network so that the application can download information from the network when it consumes less energy or does not need its batteries.

5.2 Low power short range networks

Portable computers need to be able to move seamlessly from one communication medium to another, for example from a GSM network to an in-door network, without rebooting or restarting applications. Applications require that networks are able to determine that the mobile has moved from one network to another network with a possible different QoS. The network that is most appropriate in a certain location at a certain time depends on the user requirements, network bandwidth, communication costs, energy consumption etc. The system and the applications might adapt to the *cost* of communication (e.g. measured in terms of ampère-hours or telephone bills).

Over short distances, typically of up to five metres, high-speed, low-energy communication is possible. Private houses, office buildings and public buildings can be fitted with 'micro-cellular' networks with a small antenna in every room at regular intervals, so that a Pocket Companion never has to communicate over a great distance — thus saving energy — and so that the bandwidth available in the aether does not have to be shared with large numbers of other devices — thus providing high aggregate bandwidth. Over large distances (kilometres rather than metres), the Pocket Companion can make use of the standard infrastructures for digital telephony (such as GSM).

The Moby Dick project does research on low-power short range wireless ATM networks based on near-field RF coupling [Linnenbank 96]. We have demonstrated a wireless link that delivers a bandwidth of 1 Mbps per cell. Cells have the size of a single office room.

A nano-cellular system has a number of advantages:

1. Because the distance between base stations and mobiles is small, a low transmission power is sufficient.
2. A small cell size can be used to locate mobiles and/or people in a building. A condition is that the location boundaries are well-defined. This can be obtained when the transmissions does not pass through walls (very high-frequency RF or IR) or have a rapid spatial decay of field strength (such as near-field RF coupling). Knowing where people are is the key principle for building location-aware applications.

- Small cells imply that users can utilise the full capacity of distinct cells, thereby attaining a high bandwidth density (Mbps/m²). The total aggregate bandwidth of an entire office building will be the number of cells times the bandwidth per cell.

5.3 Power aware MAC protocol

The structure of current wireless networks may cause needless energy consumption. In a wireless system the medium access protocols can be adapted and tuned for low energy consumption. A power aware TDMA protocol coordinates the delivery of data to receivers. The basic objective is that the protocol tries to minimize all actions of the network interface, i.e. minimize 'on-time' of the transmitter as well as the receiver. A base station is responsible for traffic scheduling. Mobiles with scheduled traffic are indicated in a list, which allows mobiles without traffic to rapidly reduce power. As switching between states (i.e. off, idle, receiving, transmitting) consumes time and energy, the number of state-transitions have to be minimized. By scheduling bulk data transfers, an inactive terminal is allowed to doze and power off the receiver.

A base station dictates a frame structure within its range. A frame consists of a number of data-cells and a traffic control cell¹. The traffic control is transmitted by a base station and contains the information about the subsequent data-cells, including when the next traffic control cell will be transmitted.




Figure 4: example of a TDMA frame structure

The approach described above leads to a number of low-power mechanisms, each with their particular advantages.

- Mobiles with scheduled traffic are indicated in the traffic control cell, which allows mobiles without traffic to rapidly reduce power. This mechanism requires that all mobiles and the base station are synchronized to be able to receive the traffic control cell in time.
- By explicitly scheduling all bulk data transfers, a terminal is allowed to doze (and power off the receiver) as long as the network interface is reactivated at the scheduled time to transceive the data at full speed.
- The overhead and energy consumption involved in the frame mechanism with traffic control, can be reduced when the frame size can be adapted to the situation. For example in case of a room in which only one mobile communicates, the frame size can be increased. There is a trade-off between frame size and the latency. When a low latency is required the frame size can be adapted accordingly.
- Some applications require a distribution of information to a group of users. A multicast or broadcast mechanism can reduce energy consumption since the information is sent only once, and the base station - with plenty of energy - disseminates the information. Notice that the performance (total aggregated throughput) is increased as well. The network and MAC protocols need to be adapted for these mechanisms.

1. Note that the frame structure influences the latency. This is an example of a trade-off between energy consumption and QoS.

5. Data traffic and energy consumption is reduced when the protocol allows mobile-to-mobile communication. Such a communication without a base station essentially halves the required communication effort.

6 Conclusions

More and more attention will be focused on low power design techniques as there will become an increasing numbers of portable, battery powered systems. System designers can decrease energy consumption at several levels of abstraction. At technological and architectural level energy consumption can be decreased by reducing the supply voltage, reducing the capacitive load and by reducing the switching frequency. Much profit can be gained by avoiding wasteful activity at both the architectural as system level. At system level, they can take advantage of power management features where available, as well as decomposed system architectures and programming techniques for reducing power consumption.

Remarkably, it appears that some energy preserving techniques not only lead to a reduced energy consumption, but also to more performance. For example, optimized code runs faster, is smaller, and therefore also consumes less energy. Using a cache in a system not only improves performance, but, - although requiring more space - uses less energy since the data is kept locally. The approach of using application specific coprocessors is not only more efficient in terms of energy consumption, but has also a performance increase because the specific processors can do their task more efficient than a general purpose processor. Energy efficient asynchronous systems also have the potential of a performance increase, because the speed is no longer dictated by a clock, but is as fast as the flow of data.

However, some trade-offs need to be made. These techniques often lead to less speed, that only can be improved by adding more hardware. Most energy efficient systems use more area, not only to implement the new data flow or storage, but also to implement the control part. Furthermore, energy efficient systems can be more complex. Another consequence is that although the application specific coprocessor approach is more efficient than a general purpose processor, it is less flexible. Furthermore, the latency from the user's perspective might be increased, because a system in sleep has to be wakened up. For instance, spinning down the disk causes the subsequent disk access to have a high latency.

Applications play a critical role in the user's experience of a power-managed system. Therefore, the application and operating system must allow a user to have influence on the power management.

Any consumption of resources by one application might affect the others, and as resources run out, all applications are affected. Since communication bandwidth, energy consumption and application behaviour are closely linked, we believe that a QoS framework can be a sound basis for integrated management of the resources.

7 References

- [Abnous 96] Abnous A, Rabaey J.: "Ultra-Low-Power Domain-Specific Multimedia Processors," Proceedings of the IEEE VLSI Signal Processing Workshop, San Francisco, October 1996.
- [Burd 95] Burd T.D., Brodersen R.W.: "Energy efficient CMOS microprocessor design", proc. 28 th. annual HICSS Conference, Jan. 1995, vol. I, pp 288-297.
- [Chandrakasan 95] Chandrakasan A.P., et al.: "Optimizing power using transformations", Transactions on CAD, Jan. 1995.

- [Ikeda 94] Ikeda T.: "ThinkPad Low-Power Evolution", IEEE Symposium on Low Power Electronics, October 1994.
- [Intel486SX] information can be browsed on: <http://134.134.214.1/design/intarch/prodbref/272713.htm>
- [Koegst 97] Koegst, M, et al.: "Low power design of FSMs by state assignment and disabling self-loops", Proceedings Euromicro 97, pp 323-330, September 1997.
- [Lapsley 94] Lapsley, P: "Low power programmable DSP chips: features and system design strategies", Proceedings of the International Conference on Signal Processing, Applications and Technology, 1994.
- [Larri 96] Larri G.: "ARM810: Dancing to the Beat of a Different Drum", Hot Chips 8: A Symposium on High-Performance Chips, Stanford, August 1996.
- [Linnenbank 96] Linnenbank, G.R.J. et al.: "A request-TDMA multiple-access scheme for wireless multimedia networks", Proceedings Third Workshop on Mobile Multimedia Communications (MoMuC-3), 1996.
- [Lorch 95] Lorch, J.R.: "A complete picture of the energy consumption of a portable computer", Masters thesis, Computer Science, University of California at Berkeley, 1995
- [Lorch 96] Lorch, J.R., Smith, A.J.: "Reducing power consumption by improving processor time management in a single user operating system", proceedings of 2nd ACM international conference on mobile computing and networking, Rye, November 1996.
- [Mangione-Smith 96] Mangione-Smith, W., et al.: "A low power architecture for wireless multimedia systems: lessons learned from building a power hog", proceedings of the international symposium on low power electronics and design, Monterey, August, 1996
- [McGaughy 96] McGaughy, B: "Low Power Design Techniques and IRAM", March 20, 1996, information can be browsed on http://rely.eecs.berkeley.edu:8080/researchers/brucemcg/iram_hw2.html
- [Mehra 96] Mehra R., Rabaey J.: "Exploiting regularity for low power design", Proc. of the International Conference on Computer-Aided Design, 1996
- [Merkle 93] Merkle, R.C.: "Reversible Electronic Logic Using Switches", Nanotechnology, Volume 4, pp 21 - 40, 1993 (see also: <http://nano.xerox.com/nanotech/electroTextOnly.html>)
- [Moby Dick 95] Mullender S.J., Corsini P., Hartvigsen G. "Moby Dick - The Mobile Digital Companion", LTR 20422, Annex I - Project Programme, December 1995 (see also <http://www.cs.utwente.nl/~havinga/pp.html>)
- [Patterson 96] "A Case for Intelligent DRAM: IRAM", Hot Chips 8 A Symposium on High-Performance Chips, information can be browsed on: <http://iram.cs.berkeley.edu/publications.html>
- [Piguet 96] Piguet, C, et al.: "Low-power embedded microprocessor design", proceeding Euromicro-22, pp. 600-605, September 1996.
- [Rabaey 94] Rabaey J. et al.: "Low Power Design of Memory Intensive Functions Case Study: Vector Quantization", IEEE VLSI Signal Processing Conference, 1994.
- [Rabaey 95] Rabaey J., Guerra L., Mehra R.: "Design guidance in the Power Dimension", Proc. of the ICASSP, 1995.
- [Rizzo 97] L.Rizzo: "Effective Erasure Codes for Reliable Computer Communication Protocols", ACM Computer Communication Review, Vol. 27- 2, pp 24-36, April 97
- [Sheng 92] Sheng S., Chandrakasan C., Brodersen R.W.: "A portable multimedia terminal", IEEE Communications Magazine, December 1992, pp 64-75
- [Stemm 96] Stemm, M, et al.: "Reducing power consumption of network interfaces in hand-held devices", proceedings mobile multimedia computing MoMuc-3, Princeton, Sept 1996.
- [Tiwari 94] Tiwari V. et al.: "Compilation Techniques for Low Energy: An Overview", IEEE Symposium on Low Power Electronics, October 1994.

- [Weiser 94] Weiser, M, et al.: "Scheduling for reduced CPU energy", proceedings of the first USE-NIX Symposium on operating systems design and implementation", pp. 13-23, November 1994.
- [Xilinx 95] "Minimizing power consumption in FPGA designs", XCELL 19, page 34, 1995.