Marian Verhelst
Wim Dehaene

ACSP
Analog Circuits And Signal Processing

Energy Scalable Radio Design

for Pulsed UWB Communication and Ranging

Springer

Energy Scalable Radio Design

ANALOG CIRCUITS AND SIGNAL PROCESSING SERIES

Consulting Editor*: Mohammed Ismail.** ***Ohio State University

For other titles published in this series, go to
http://www.springer.com/series/7381

Marian Verhelst • Wim Dehaene

Energy Scalable Radio Design

for Pulsed UWB Communication and Ranging

Springer

Dr. Marian Verhelst
Katholieke Universiteit Leuven
Dept. Electrotechnical
Engineering
Div. Microelectronics &
Sensors (MICAS)
Kasteelpark Arenberg 10
3001 Leuven
Heverlee
Belgium
marian.verhelst@gmail.com

Prof. Wim Dehaene
Katholieke Universiteit Leuven
Dept. Electrotechnical
Engineering
Div. Microelectronics &
Sensors (MICAS)
Kasteelpark Arenberg 10
3001 Leuven
Heverlee
Belgium
wim.dehaene@esat.kuleuven.be

ISBN 978-90-481-2693-4 e-ISBN 978-90-481-2694-1
DOI 10.1007/978-90-481-2694-1
Springer Dordrecht Heidelberg London New York

Library of Congress Control Number: 2009926168

Cover design: eStudio Calamar S.L.

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

To Thomas.
Marian

Preface

The book you are holding is the result of 5 years of cooperation between a PhD researcher and her advisor. It was a rather exceptional cooperation, one where the technical and strategic thinking flowed quite easily in the same direction yet remained very complimentary. This was probably caused by a similar mindset in two individuals with a different background and experience. Combine this with an infinite eagerness to learn, solve, and create, and you can move mountains.

The result is a nonclassic research outcome. It is characterized more by width than by depth. We handle the design of a low-power pulsed UWB based wireless communication and ranging system in all of its aspects and abstractions levels. We have tried to find the right balance between generality and a concrete application. Therefore, the techniques and methods we describe are not limited to the application domain of UWB or sensor networks. Yet, we did not want to write a pure methodology book decoupled from any application or circuit at the risk of being generic and possibly irrelevant. What you hold in your hands is thus a journey starting from the optimal selection of the communication technique, going over optimized acquisition and synchronization algorithms via optimally flexible architectures, ending in a verified silicon implementation. We spent a lot of effort in defining what the optimal degree of flexibility of a system is, and, for a given application, how this can be applied to find the best architecture. Apart from the architectures and circuits themselves, this is the most important part of the book.

We think that the future of semiconductor system on a chip (SoC) design lies in recombining the different levels of design abstraction in a deliberate manner. Although a full SoC cannot be designed at transistor level, the future will be in the hands of those who, at least in the conceptual phase, have the widest span of abstraction levels in their thinking and conceiving. We hope that this book will be instructive in demonstrating how these kinds of problems can be tackled.

The book you are holding is the outcome of a research trajectory we personally enjoyed very much. We learned a lot and gained valuable insight. We profited from the fact that a mind once stretched by a new idea, never regains its original dimensions [O.W. Holmes]. We dare to hope, dear reader, that you will enjoy the learning experience in this book as much as we enjoyed the research that led to its creation.

Marian Verhelst, Wim Dehaene
Leuven, March 2009

Table of contents

Chapter 1

Introduction and Motivation

1.1 Dreaming of a Smart Environment

The exponential increasing integration density, as dictated by Moore's law [Moo65, Intb], together with recent advantages in Micro ElectroMechanical Systems and wireless communications, allow electronics to be so small and powerful that they can be integrated into any physical object (Figure 1.1). This observation opens up a whole new world of opportunities: clothes and furniture can be made smart, machine and construction monitoring can be automated, and accurate tracking of personnel and goods becomes feasible. The environment becomes sensitive and responsive to the presence of people and sympathetic to their needs. Distributed networks of small, intelligent, embedded devices provide us with information, communication and entertainment, and promise to make our future life easier and our surroundings more secure and pleasant. Different buzz words have been used to refer to this brave new world: "smart dust" [Kah99], "ubiquitous computing" [Win91], "pervasive computing" [Sat01], and "ambient intelligence" [Aar03, Phi] popped up in scientific and later also in popular literature.

The basic unit, which allows all these promising applications, is a microsensor node. Several of these microsensor nodes configure themselves into a wireless distributed network for collecting, processing, and disseminating wide ranges of complex data [Rag02, Aky02b]. The sensor node is typically made up out of a limited set of components, containing one or more sensors, a limited storage and control system, an energy subsystem, including the energy supply and energy management, and a wireless communication unit.

1.2 Limited Energy Resources and the Energy Gap

Moore's law does not apply to batteries. Unlike solid-state electronics, which continue to benefit from scaling at the device level and continued miniaturization of the package, batteries depend on electro-chemical reactions within the cell, which do not follow similar scaling laws [Cha07]. Figure 1.2 illustrates this by comparing advances in battery, hard disk, and microprocessor technology. While the energy density of integrated circuits experienced a spectacular increase over the last decade, the energy density of batteries

M. Verhelst and W. Dehaene, *Energy Scalable Radio Design: for Pulsed UWB Communication and Ranging*, Analog Circuits and Signal Processing, © Springer Science+Business Media B.V. 2009

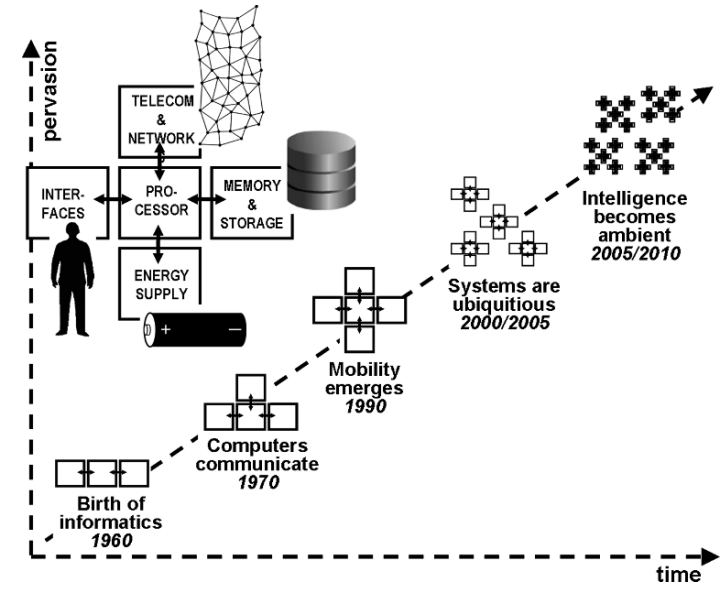

Figure 1.1: *An (expected) evolution of computing from 1960–2010. Adapted from [Wal07]*

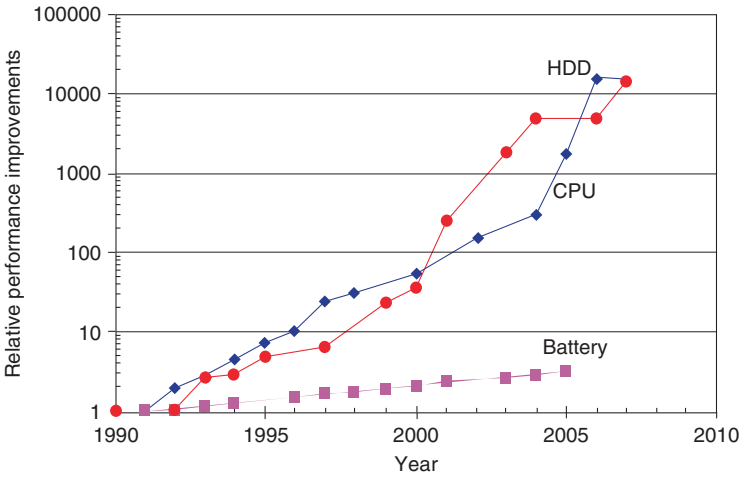

Figure 1.2: *Relative improvement in the energy density of lithium ion batteries vs. the areal density of hard disk drives and the number of transistors in Intel microprocessors. Source [Cha07] ©2007 IEEE*

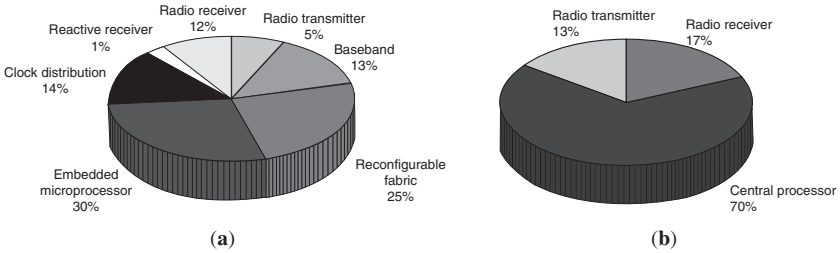

Figure 1.3: *Energy breakdown of a sensor node developed in academia: PicoNode (**a**) and a commercial mote: MICAz, Crossbow Technology (assume 10% Rx and Tx duty cycle) (**b**) Source [Rab02] ©2002 IEEE, resp. [Teca]*

only doubled every 5–20 years (depending on the particular chemistry) [Pow94]. Recently, new developments in energy scavenging were presented, where energy is harvested from the environment through seismic, photo-voltaic, or thermal conversion [Rou03]. In the future, batteries might be replaced by these energy scavenging units. However, the energy produced by these new techniques is still very limited.

As a result, there exists a continuously growing energy gap between the available energy and the energy needs of battery- powered devices. A sensor node, hence, has to operate on an extremely frugal energy budget. Smart energy management together with an ultralow-energy consumption of all the different sensor node components is crucial for the lifetime of the sensor node. Figure 1.3 shows the energy breakdown of an academic (the PicoNode, UC Berkeley, [Rab02]) and a commercial sensor node (the MICAz mote, Crossbow Technology, [Teca]). The plot shows that, in both examples, a relatively large part of the energy consumption (30%) goes to the wireless link physical layer (receiver, transmitter, and baseband). It is that part of the sensor node where this work will focus on: an energy-efficient physical layer wireless link.

1.3 Strategies to Bridge the Energy Gap

The increasing energy gap between a sensor node's energy supply and its energy needs asks for energy saving strategies at all levels during design and at run time. This requires a completely new way of designing circuits and systems. The following sections sketch different strategies to bridge this energy gap, and some related work in the different research domains.

1.3.1 Power- and Energy-Oriented System-to-Circuit Design

The traditional digital design flow for very large-scale integration (VLSI) systems, introduced in the 1980s (see Chap. 2), focuses primarily on optimizing speed to implement circuits for computationally intensive tasks. The design's power consumption, together with its area, was only a secondary parameter and was not optimized as long as it stayed within reasonable bounds. During the 1990s, however, the increasing demand for portable

and nomadic devices elevated power consumption (Sect. 1.3.1.1) and, more recently, energy consumption (Sect. 1.3.1.2) to be primary design objectives.

1.3.1.1 Power Analysis and Optimization at Design Time

The subject of low-power design quickly attracted wide attention from the research and industrial community. Different power saving techniques were established at all levels of design. At physical, circuit, and logic level, the effect of place and route optimization, transistor sizing, clock and power gating, and combined supply and threshold voltage scaling on the design's power consumption were studied [Cha95a, Ali94, Kao02]. Hardware duplication or pipelining at architectural level allows to lower the supply further, resulting in additional power savings [Cha92b]. Other architectural power optimization techniques include input reordering, resource sharing, topology selection, and the minimization of glitching activity [Cal92, Cha95a, Cha95b]. Transformations at algorithmic level [Meh94, Cha95b] and algorithm selection [Pot00] will, finally, also heavily alter the system's power consumption. These power minimization techniques at the higher design levels, however, require tools that can accurately predict power consumption early in the design process [Meh94, Lan96b, Rab91]. The combination of this power estimation and power minimization capabilities finally resulted in various high-level low-power synthesis systems: Flamel [Tri87], SAW [Wal89], SPAID [Har89], HYPER [Cha92a, Rab91], and CATHEDRAL [Fra93]. These tools allow a quick design space exploration, in which several "functionally equivalent" implementation alternatives are analyzed to identify an optimal solution [Moh02]. They, however, all synthesize a given, fixed algorithm while minimizing the resulting power consumption for a fixed throughput, and do not cover optimizations at system level.

1.3.1.2 Energy Analysis and Optimization at Design Time

More recently, the research community shifted its focus to energy consumption rather than pure power consumption. The energy consumption of a system is its power consumption multiplied by the time over which it is active. In many applications, especially for energy-limited devices like sensor nodes, the energy to execute a certain operation is much more important than its power consumption during this action. Figure 1.4 illustrates this graphically. The power consumption of exactly the same operation (e.g., a 16 bit multiplication) is plotted twice, once when it is optimized to finish in time T and once for a reduced delay requirement of $3T$. The slower implementation, option 2, has

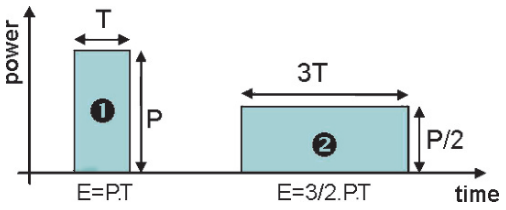

Figure 1.4: *Comparing energy and power consumption*

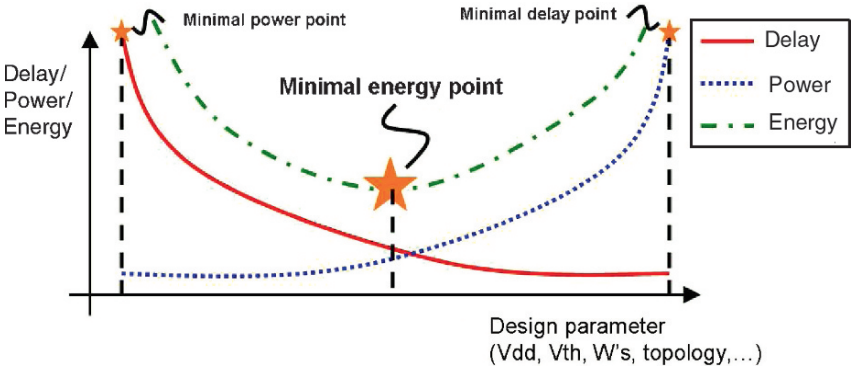

Figure 1.5: *Trade-off between delay and power consumption*

more margin to implement various power saving techniques and benefits from a lower power consumption $P/2$. This example, however, shows that the solution offering the lowest power consumption is not necessarily the one consuming the lowest energy per operation (E). In energy-limited designs, which do not have a tight throughput constraint, energy but not power consumption should be the primary design objective.

Designing for energy-efficiency, hence, requires carefully balancing the system's performance (delay) against its power consumption. As discussed in Sect. 1.3.1.1, several design parameters at different levels of abstraction can be tuned to optimize the system's power consumption for a fixed throughput. Repeating this for various delay requirements results in power–performance trade-off curves, as plotted in Figure 1.5. The best trade-off between power and performance is the minimal energy point. If the system is subject to delay or throughput constraints, the design might have to settle for a more energy-consuming solution to satisfy these constraints.

Several energy optimization design frameworks are reported in literature [Gon97a, Zyu02, Mar04, Wan01]. While the majority of them focus on energy optimization on the circuit and microarchitectural level, some work has been done on the algorithm to architecture mapping on specific reconfigurable platforms [Wan01, Zha00, Bon05]. Very little is, however, done on energy-optimization of ASICs at the higher levels of design abstraction [Ben00], though it is well known that the most significant energy savings can be obtained there. Figure 1.6, which will be worked out into more detail in Chap. 2, illustrates this graphically.

This research gap can be explained by the difficulty of performing energy optimization at these highest levels of design: at system and algorithmic level. Not only is it very hard to model the energy consumption at these levels of design, but the optimizations at these levels also require a crossdisciplinary view. In the context of the design of wireless transceivers, covered in this work, decisions at this level include the selection of the air interface, the communication algorithms, the transceiver architecture, etc. Studying the impact of these decisions on the system's performance and its energy consumption involves both digital and analog implementation theory, as well as communication

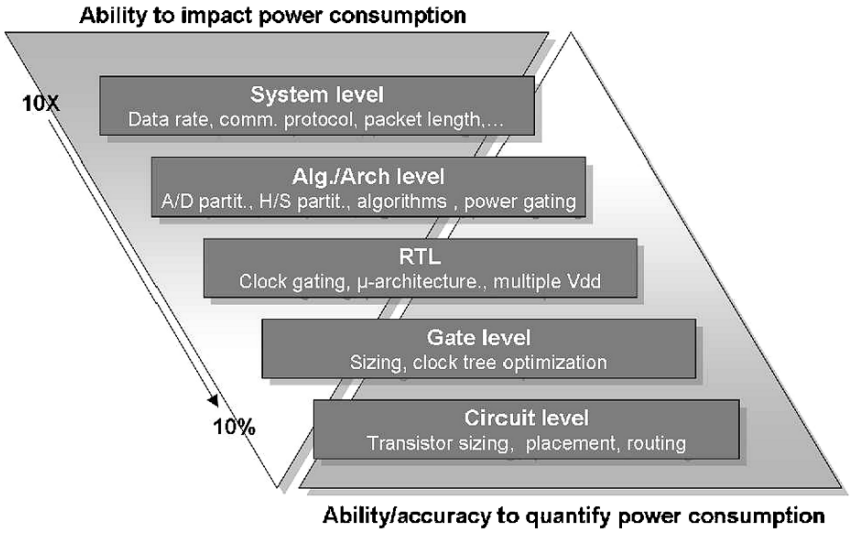

Figure 1.6: *Ability to impact energy consumption, resp. to model energy consumption over the different design abstraction layers*

theory. Moreover, crosslayer optimization will be necessary, since high-level system aspects (modulation type, preamble lengths, etc.) have to be considered in parallel with lower level architectural and circuit implications. This crosslayer approach is necessary in energy-limited designs, but it is in flat contradiction with the strict separation between the different abstraction layers of the classical design flow. It is precisely this new vision on energy-driven system-to-circuit design which forms the foundation of this book. History reveals that all efficient design strategies are strongly tied to one particular target application area [Man88]. The proposed design strategy will, in the next chapters, be refined and applied in the context of wireless communication. The optimization will be restricted to the physical layer of the wireless link. Interesting results on energy minimization at the MAC, network, and upper layers of the protocol stack in the context of wireless sensor networks can, for instance, be found in [Aky02b, dS01]. Optimizations at these levels include a.o. the development of energy-efficient routing protocols, addressing methods, collisions avoidance, wake-up schemes, etc. Also, here, a crosslayer approach is often proposed, where different layers share their network-status information at run time to minimize the energy consumption of the complete system under the current Quality of Service and energy constraints [Con04, Bou06a].

1.3.2 Energy-Efficient Wireless Communication and IR-UWB

1.3.2.1 Energy-per-Useful-Bit

A very important first step in the design of an energy-efficient wireless communication system is the selection of the air interface and the communication algorithms, including

modulation, data rate, carrier frequency, and synchronization techniques. Recently, the energy-per-useful-bit (EPUB) metric was introduced [Amm06], which allows a meaningful comparison of the different alternatives in the context of energy-optimal design. The EPUB includes the energy consumption of both the transmitter and receiver per communicated data bit, and amortizes the energy consumption during the synchronization preamble over the number of data bits in the packet [Amm06]. The metric, hence, jointly considers the communication theoretical aspects (modulation efficiency, data rate, bit-error-rate performance, synchronization, etc.) as well as the implementation aspects (power consumption) of the different alternatives. This is exactly what is needed in an energy-driven physical layer design flow.

1.3.2.2 Impulse Radio UltraWideband Communication

As will be shown further (Chap. 3), impulse radio ultrawideband (IR-UWB) [Win98, Por03, Siw04] is an excellent air interface candidate for communication in sensor networks. IR-UWB is a wireless communication technique based on the transmission of ultralow-power pulses. The very short duration of the pulses in time, maximum a few nanoseconds, results in a very wide spectrum in the frequency domain (Figure 1.7b). This is essentially the mathematical dual of conventional narrowband systems, where sinusoidal signals are narrow in the frequency domain and "wide" over time, as illustrated in Figure 1.7a. The low spectral density of IR-UWB signals allows to operate in already allocated frequency bands, without interfering with other users of these bands.

The advantages of this communication technique in the context of wireless sensor networks are manifold [Por03]: spreads the pulse energy over a wide frequency range, makes the wireless link more robust against narrowband interference and against spectral notches due to multipath destructive interference. In the time domain, this can be explained by the very short duration of the wideband pulses, which minimizes the effect of interpulse

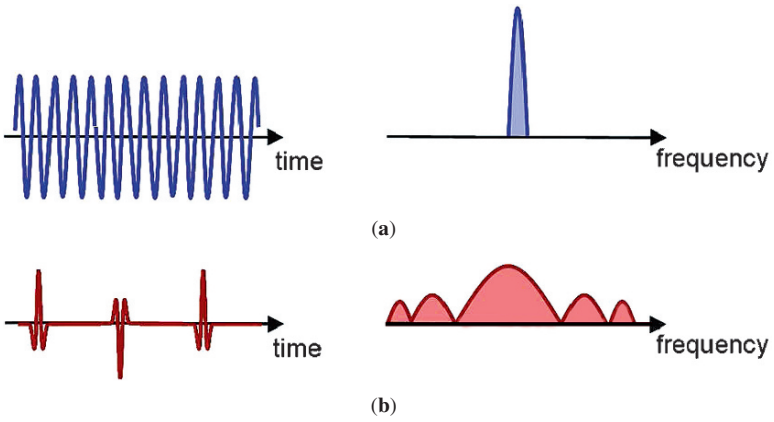

Figure 1.7: *Time and frequency domain representation of traditional narrowband signals (**a**) and impulse-based UWB signals (**b**)*

interference. The very short duration of the pulses also allows accurate ranging based on time-of-arrival measurements. Last but not least, IR-UWB will prove to offer a very low energy consumption per useful bit (Chap. 3). Especially, the energy consumption of the IR-UWB transmitter is extremely low: the UWB transmitter benefits from a low complexity, low transmit power level, loose linearity requirements, and its ability to operate in a duty-cycled way [Wen07, Ryc07b, Kim04, Ryc05, Mar03]. As a consequence, the link will be asymmetric, with a receive chain which consumes significantly more than the transmit chain. Also, this property is desirable in many sensor network applications. Often, the transmit power requirements are much more stringent than the receive requirements, since most of the information typically transits from an energy-limited sensor device to a, in terms of energy, more relaxed master node.

The design of a low-energy IR-UWB receiver, however, remains challenging. This receiver has the difficult task to synchronize to the tiny pulses drown in noise, to stay synchronized, to decode the data, and to do accurate time-of-arrival measurements. Since its energy consumption will heavily dominate the EPUB of the wireless link, energy should be minimized at all levels of design. The energy-driven design of such an IR-UWB receiver forms the core of this book. Related work on existing state-of-the-art IR-UWB receivers will be covered extensively in Chaps. 3 and 7.

1.3.3 Runtime Energy Scalability

Previous sections focused on the optimization of the energy consumption of the physical layer link at design time. Additional energy savings can, however, be obtained by intelligently managing the system at run time.

The wireless sensor network environment is characterized by fast- varying channel conditions, application requirements, and energy availability. Wireless transceivers operating in this environment should be designed to still function under worst-case conditions. This would, however, result in a system which is over-conservative and which consumes much more energy than needed during the majority of the time. It is better to have a system which is able to dynamically scale down its performance depending on the channel and application requirements. Such dynamic energy management [Ben00, Bou06a] results in significant additional energy savings and allows to continuously adapt the transceiver system to the current circumstances. Possible examples are the dynamic adaptation of the transmit power, receive gain, data rate, amount of channel compensation, etc.

This form of runtime energy management is only possible when the transceiver system can be reconfigured at run time. This can be realized in different ways. The most straightforward solution is to multiplex various alternative implementations of the transceiver on one die. At run time, one of them is selected for use based on the current operating conditions. The area penalty of this approach would, however, be unacceptably high. A more elegant solution to create the necessary trade-off mechanisms in the transceiver system is to built several “flexibility knobs” into the design. The selection of appropriate transceiver parameters to make flexible is, however, not straightforward. Flexibility never comes for free, as illustrated in Figure 1.8. The introduction of every flexibility knob brings an un-

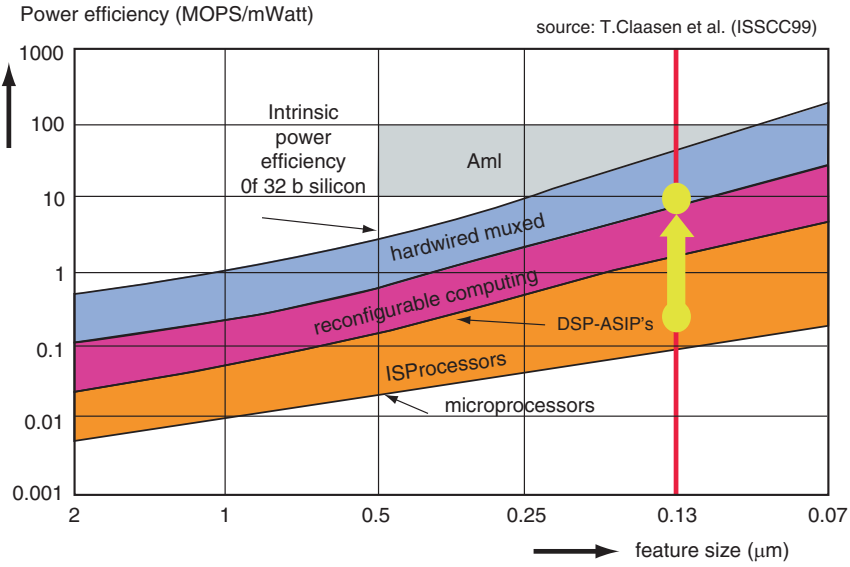

Figure 1.8: *Power-efficiency in function of flexibility of the design. Source [Cla99] ©1999 IEEE*

avoidable energy penalty, which will be larger when the range over which the parameter can vary increases. Only when this energy penalty is smaller than the expected runtime energy savings due to this additional flexibility, the knob should be included in the design. For all flexibility knobs the optimal flexibility range, which results in the lowest energy consumption at run time, should hence be determined at design time.

A large body of work exists on runtime energy management strategies for a given, flexible system (See [Ben00] for a survey). Very little, however, has been done on the design of these flexible systems: which flexibility knobs to build into the system and over which range they should be able to vary? This requires a careful investigation of the power–performance–flexibility trade-off at every design abstraction layer. Only by meticulously balancing the benefits of making the design more flexible vs. the penalty in terms of power consumption under various channel and application scenarios, the optimal degree of flexibility can be determined. An important hurdle that has to be cleared when tackling this problem is, "how to measure the flexibility of a design." Only a few flexibility measures were presented in literature, and most of them take a system's view and lack the connection to the resulting implementation [Hau02, Com04]. Up to now, no good all-round quantitative measure of flexibility has been found, which can be used in these power–performance–flexibility trade-offs. We believe such a general measure is neither possible nor meaningful. Only in the context of a specific application domain, flexibility can be measured and compared.

1.4 Book Scope and Organizational Overview

Previous discussions revealed the necessity of a consistent energy-driven design strategy at all levels of system design for wireless communication applications. This flow should not only strive to an optimal balance between the system's performance and its power consumption but also to simultaneously study the effect of the introduction of runtime flexibility on these two parameters. At design time, this power–performance–flexibility trade-off should be explored to derive the required degree of flexibility of the system and come to the most energy-efficient solution. This approach requires crosslayer thinking, since it involves combining communication theoretical aspects as well as implementation aspects.

The goal of the work reported in this book is dual; next to the discussion of such an energy-driven system-to-circuit design strategy, an energy-efficient physical layer wireless link for communication in sensor networks is designed. To this end, the proposed energy-driven design strategy is applied to the design of this wireless communication system. As shown in Figure 1.9, the scope of the design is narrowed along the design flow. Starting with system-level studies covering the complete wireless link, the work gradually focuses on the design of the flexible receiver back-end chip. However, whenever a design decision is taken or a power–performance trade-off is explored, the repercussions on the energy-efficiency of the complete wireless link are always considered.

The remainder of the book is structured as follows:

Chapter 2 starts by summarizing the classical performance-driven top–down design flow, its history, and its advantages. Next, the drawbacks of this flow in the context of energy-limited design are described. To overcome these problems, several adaptations to the classical top–down flow are presented: the introduction of an algorithmic/architectural

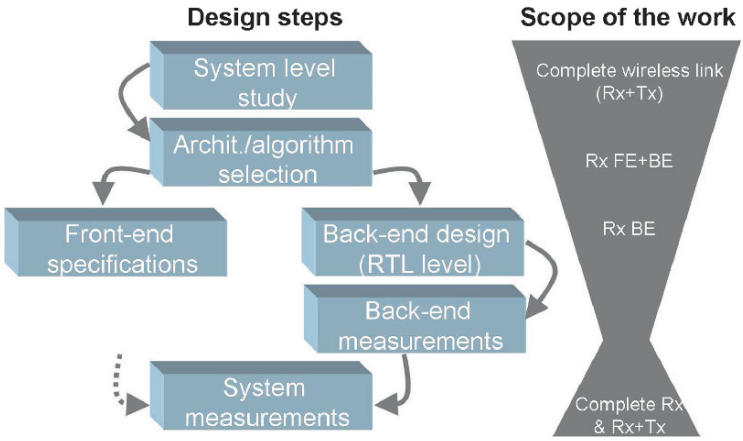

Figure 1.9: *The scope of the different steps taken in the reported design of an energy-efficient physical layer wireless link*

design space exploration together with an energy-driven, crosslayer design strategy to optimally balance power and performance. Finally, the importance of considering runtime flexibility at design time is discussed.

Chapter 3 covers the system-level design step of the energy-efficient wireless physical layer link. The system-level specifications are derived starting form two different target applications. Subsequently, various air interface candidates are compared in terms of EPUB, and IR-UWB is selected. The chapter ends by a short overview of IR-UWB communication, its basic principles, its history, and its advantages.

Chapter 4 describes an elaborate energy-oriented algorithmic/architectural design space exploration on the design of an energy-efficient IR-UWB receiver. Several coarse receiver architecture/algorithm pairs are codesigned and compared to come to the most energy-optimal solution.

Chapter 5 covers the classical algorithmic/architectural- level design step. The chapter starts with an extensive study of the necessary algorithms for acquisition, data reception and ranging. Subsequently, the final receiver architecture is refined and all important front-end building block specifications are derived.

Chapter 6 focuses on the RT level design of the digital receiver back-end. The chapter starts by introducing the novel design concept of nested *FLEXmodules*, which is the key to a flexible implementation without large power penalties. Small, nested, reconfigurable units allow to easily shut down unused blocks, slow down and gate clocks whenever possible, and distribute the control. Secondly, a multidimensional measure of flexibility based on this concept allows to quantify and weigh the flexibility of strategic units in the design. Two case-studies are worked out in detail, where this concept is applied to find the required degree of flexibility of a particular subblock. Finally, the complete implemented digital back-end is described.

Chapter 7 presents the results from measuring the implemented back-end chip in an isolated way, as well as integrated in two different full IR-UWB receiver systems. This integration resulted in two of the first complete IR-UWB receivers reported in literature, including both back-end and front-end, implementing all necessary algorithms for communication, time-of-arrival measurements, and synchronization. The measurements demonstrate the implemented flexibility, the resulting energy savings, and the system's best-in-class energy-efficiency. An elaborate comparison with other state-of-the-art IR-UWB receivers is made.

Finally, *Chapter 8* summarizes and concludes this book.

Chapter 2

Adaptation of Classical Design Flow for Energy-Driven System-to-Circuit Design

2.1 Introduction

This chapter starts by sketching the classical (digital) top–down design flow and its advantages; separate layers of abstraction allow to gradually refine a design. This strategy increases the design efficiency and allows the designer to tackle the design of large systems. Unfortunately, it is exactly this strict separation between the different abstraction layers which forms a problem in present-day energy-limited designs. The design of such energy-optimized systems requires a careful balancing of the system's performance and its power consumption. This is only possible by considering high-level system aspects in parallel with lower level architectural and circuit implications. A new kind of energy-driven crosslayer design strategy is needed. The second part of this chapter will propose several adaptations of the classical design flow. Yet, the basic concept of this design flow, the different design abstraction layers, is not abandoned, since it is indispensable in the design of large complex systems.

The insertion of an early design space exploration (DSE) step in the design process and the introduction of crosslayer design techniques are two of the adaptations proposed in this chapter. Additionally, the role of flexibility in the system under design will be discussed. By introducing flexibility knobs at design time, the system can dynamically adapt to changing requirements at run time, which results in additional energy savings. The presented design strategy should, hence, optimally balance the system's power consumption, performance, and flexibility to minimize the overall energy consumption at run time.

The proposed design strategy will target the design of energy-efficient wireless communication circuits. These systems will typically consist of both an analog and a digital part. It is important that for every design decision, its influence on the complete mixed-signal system is considered. This chapter ends by showing how this is reflected in the design flow.

M. Verhelst and W. Dehaene, *Energy Scalable Radio Design: for Pulsed UWB Communication and Ranging*, Analog Circuits and Signal Processing, © Springer Science+Business Media B.V. 2009

2.2 Classical (Digital) Top–Down Design Flow: Gajski–Kuhn

2.2.1 Moore's Law and the Design Productivity Gap

Already, in 1965, Intel cofounder Gordon E. Moore made the observation that the number of transistors that can be inexpensively placed on an integrated circuit is increasing exponentially [Moo65]. In 1975, he altered his projection, originally a doubling of the chip's transistor count every year, to a doubling approximately every 2 years [Wikc, Intb]. This became later known as "Moore's law." Surprisingly enough, this law still holds after more than 30 years of chip design. Figure 2.1 shows Moore's law together with the transistor count on Intel's processor generations. The reality, hence, follows the prediction extremely closely.

This exponential increase has steadily and reliably led to increasing performance and energy-efficiency. On one hand, this is all very exciting for chip designers and semiconductor companies, since it allows them to double the chip's computing power and complexity every 2 year. On the other hand, it becomes harder and harder to design these chips. The design of a chip with a certain silicon area will require more resources and larger design teams.

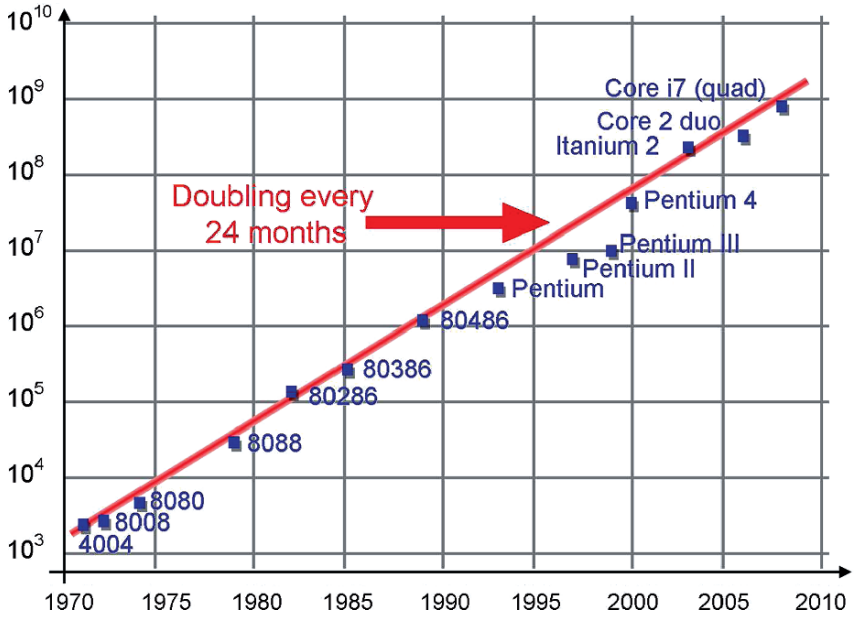

Figure 2.1: *Moore's law together with growth of transistor counts for Intel processors. Data from [Wikd]*