

Writing a Paper



Paper Title

Author 1

Author 2

Author 3

Author mailing address line 1

Author mailing address line 2

Author email addresses

ABSTRACT

System level power management has become a key technique to render modern wireless communication devices economically viable. Despite their relatively large impact on the system's energy consumption, power management for radios has been limited to shutdown-based schemes, while processors have benefited from superior techniques based on Dynamic Voltage Scaling (DVS). However, similar scaling approaches that trade off energy versus performance are also available for radios. To utilize these in the radio power management, existing packet scheduling policies have to be thoroughly rethought to make them energy-aware, essentially opening a whole new set of challenges the same way the introduction of DVS did to CPU task scheduling. We use one specific scaling technique, Dynamic Modulation Scaling (DMS), as a vehicle to outline these challenges, and to introduce the intricacies caused by the non-preemptive nature of packet scheduling and the time-varying wireless channel.

I. INTRODUCTION

1. System Power Management

In commercial and research efforts alike, recent trends have shifted the emphasis in system design away from ever-increasing throughput alone. Instead, energy and power consumption are becoming ever more formidable and important constraints to address [2][22]. Indeed, technology integration and miniaturization have intensified cooling and packaging challenges. In addition, the energy supply itself is often limited to build-in batteries, as untethered operation is desired for mobility and portability. A class of systems where this observation holds especially true is that of personal communication devices, which are also becoming increasingly prevalent.

To reduce power spending without sacrificing performance, system designers have drawn on the concept of 'energy or power awareness'. The key idea is that the system only delivers the performance that is strictly required, thereby avoiding superfluous power consumption. Power and energy awareness are therefore often paraphrased as 'the right power/energy at the right time and

the right place'. When designing according to this concept, there are two aspects to consider: the technique itself that introduces the awareness, and the power management strategy that exploits it.

The oldest and most straightforward technique is to shut down unused parts of the system [3]. Shutdown-based power management has been explored for hard disks, displays [17], and communication modules [37], amongst others. For processors, it has been incorporated in the kernel of Real-Time Operating Systems (RTOS) [33]. However, a breakthrough came with the development of Dynamic Voltage Scaling (DVS), a technique for digital circuits that is more effective than shutdown [5]. The reason is the convex nature of the power-speed curve, which comes about by varying the operating voltage. Numerous DVS-based power management strategies have been proposed to harness this potential [4][9][10][11][15][19][20][31][39][40].

Unfortunately, energy-awareness brought forth by DVS is restricted to digital circuits, such as processors, which can leverage DVS in their task-scheduling engine. Voltage scaling not being applicable, is it nevertheless possible to find a similar technique for other parts of the system? Or is shutdown the best we can do?

2. Power Management in Communication Subsystems

The search for superior power management techniques is especially relevant for the communication subsystem, and in this paper, we introduce and discuss a powerful scaling-based approach that can be viewed as the counterpart of DVS.

The importance of radio power management arises from the fact that communication is often the dominant power hog in the complete system [25]. While the inherent energy consumption of digital circuits is rapidly decreasing due to Moore's law and ingenious design techniques, the power that is radiated to carry information does not follow this trend. This observation is particularly true for wireless RF (radio-frequency) devices, where the radiated power depends on the transmit distance and is physically constrained by Maxwell's laws. With the increasing proliferation of cellular phones with Internet browsing

capabilities, Bluetooth-connected PDA's, and wireless LANs, radio power management will undoubtedly gain in importance.

Radios typically possess a number of control knobs, which give rise to a convex power-performance curve. As such, these knobs can be used to introduce energy awareness in a similar fashion the operating voltage did in DVS. Possible radio control knobs are the modulation level, the error coding, or combinations of both. We refer to the resulting techniques as Dynamic Modulation Scaling (DMS), Dynamic Code Scaling (DCS), and Dynamic Modulation-Code Scaling (DMCS), respectively. To provide energy awareness, they need to be integrated into the system power management.

The aim of this paper is to provide insight into such scaling-based radio power management strategies, the associated challenges and possible solutions. Just as people created energy aware versions of RTOS task scheduling policies, we investigate energy-aware packet scheduling as part of these power management strategies. As a vehicle for this exploration, we focus on Dynamic Modulation Scaling (DMS), which we first introduced in [28]. It is a readily accessible control knob in a number of existing communication systems.

II. RELATED WORK

DMS for radios can be viewed as the counterpart of DVS for digital circuits, as both exploit the presence of a convex power-speed curve via the notion of scaling. We can therefore find inspiration in the way task scheduling was extended to make it energy-aware, when developing energy-aware packet-scheduling policies. Nevertheless, a direct correspondence is impossible due to the inherent differences between radios and digital circuits, which we will detail in section III.4. Voltage scaling was first included in self-timed circuits [20] and later extended toward synchronous ones [11], where a buffer is used to smooth the load and steer the adaptation. The initial DVS work in operation system research was in the context of a workstation-like environment [9][39], with average throughput as the performance metric. Task scheduling under real-time constraints, i.e., the presence of deadlines, is considered in [4][10][15][19][31][40].

In the realm of communications, a shutdown approach is proposed in [37], leveraging the time-varying nature of the wireless channel. Transmissions are deferred to times where the channel can support energy efficient data transmission, while taking into account various delay constraints. A scheduling technique based on code scaling (DCS) is described in [23]. It finds an energy-efficient schedule for a set of packets that have to be transmitted

before an overall deadline, and is similar to the work presented in [40] on processor scheduling. The basic concept of modulation scaling (DMS), we proposed it first in [28]. Here, we build upon this concept to describe a complete framework for dynamic power management of the radio communication subsystem in wireless embedded systems.

We explore radio power management through DMS, which essentially relies on the ability of the radio to change its modulation on the fly. This is indeed practically feasible, and appropriate hardware architectures have been developed [6]. Originally, these architectures were used for a different purpose, namely adapting the modulation to maximize the system throughput [1][36][38]. In this case, the appropriate choice of the modulation level only depends on the current condition of the wireless channel, and does not involve any scheduling decisions. On the other hand, as we will illustrate in this paper, energy awareness is strongly related to scheduling, where current decisions and future events are tightly interwoven. Although utilizing the same radio control knob, DMS and throughput maximization not only serve a different purpose, but also require completely different adaptation policies. This change in mindset is similar to the one that occurred in the CPU world when the design goal switched from throughput maximization to energy efficiency.

III. DYNAMIC MODULATION SCALING (DMS)

When considering energy-aware techniques beyond shutdown, such as DMS, we require some more insight into the operation of radios. In this section, we introduce the basics of modulation scaling.

1. Energy and Delay Breakdown

The first step in understanding DMS is investigating how energy and throughput depend on the modulation level in a digital wireless communication system. In order to transmit information, bits are coded into channel symbols, which correspond to different waveforms [24]. The number of possible waveforms determines how many bits are coded into one symbol, which is given by the modulation level b , expressed in number of bits per symbol. The average time to transmit one bit over the channel is the inverse of the average bit rate R_b , and given by (1), where R_s is the symbol rate (number of symbols that are transmitted per second).

$$T_{bit} = \frac{1}{R_b} = \frac{1}{b \cdot R_s} \quad (1)$$

The energy associated with the transmission of one bit can be expressed as (2). The power needed to generate the information carrying electro-magnetic waves, is delivered

Elements of a Paper

- Title
- Abstract
- Introduction
- Related Work
- Technical Sections
- Results Sections
- Conclusions (and Future Work)
- Acknowledgments
- References
- Appendices

What are the most
important parts of a
paper?

<http://infolab.stanford.edu/~widom/paper-writing.html>

<http://homes.cs.washington.edu/~mernst/advice/write-technical-paper.html>

Basics of Technical Papers

- A paper is different from a technical report
- The most important parts of a paper
 - Abstract
 - Introduction
- Abstract:
 - Encourage the reader to read the full paper
- Introduction:
 - A reviewer typically makes the decision to accept or reject the paper based on the introduction and reads the rest of the paper looking for evidence to support this decision
 - A casual reader will continue on if the introduction is captivating

Writing Style

- The flow is the most important
 - Sketch it out
 - Treat it like giving a talk in terms of flow
 - Flesh this out before writing any finished sentences
- General rules
 - Appendices for the gory details, e.g. proof or algorithm description, pseudo code
 - Define all ‘variables’ (acronyms, unusual words, definitions)
 - Explain acronyms first time they appear (abstract does not count)
- I write like I talk, but not everyone likes this style
- Consider forward references
 - You can keep the flow without going into too much detail right away
 - Prepares readers for what is to come: important for reviewers
 - Very useful in the introduction

Title

- Long and descriptive
 - E.g., “*Linear-Time External Multipass Sorting with Approximation Guarantees*”
 - Be careful not to go overboard
- Short and sweet
 - E.g., “*Approximate External Sort*”
 - Often not specific enough
- Middle-of-the-road length, plus a cute name that sticks in people's minds:
 - E.g., “*Floosh: A Linear-Time Algorithm for External Sort*”
 - This is often a good compromise between specific and descriptive

Abstract

- Definition:
 - Concise summary of a completed research project or paper
- Goal:
 - Make the reader want to learn more about your research, read your paper, or attend your presentation
 - Even more important for online publication – ‘sell your work’
- Length:
 - Typically, around 200 to 300 words
 - Limit often specified

Long Distance Wireless Sensor Networks: simulation vs reality

Marco Zennaro
The Abdus Salam
International
Centre for Theoretical Physics
Trieste, Italy, and
KTH
Stockholm, Sweden
mzennaro@ictp.it

Antoine Bagula
University of Cape Town
Cape Town, South Africa
bagula@cs.uct.ac.za

David Gascon, Alberto
Bielsa Noveleta
Libelium
Zaragoza, Spain
d.gascon@libelium.com,
a.bielsa@libelium.com

ABSTRACT

Wireless sensor networks allow unprecedented abilities to observe and understand large-scale, real-world phenomena at a fine spatial-temporal resolution. Their application in Developing Countries is even more interesting: they can help solve problems that affect communities. One of the limitations of current wireless sensors is the communication range with most devices having 100 meters as maximum limit. In contrast, many applications require long-range wireless sensor network where nodes are separated by large distances giving the advantage of being able to monitor a vast geographic area. In this paper we will present the results of simulations and of experiments carried out using off-the-shelf equipment over distances ranging from 300m to 12km. The results show that long distance wireless sensor networks (LDWSN) are possible and that the quality of these links is high.

Categories and Subject Descriptors

H.4 [Information Systems Applications]: Miscellaneous;
D.2.8 [Software Engineering]: Metrics—complexity measures, performance measures

General Terms

Wireless sensor networks, long distance wireless, ICT4D

Keywords

Wireless sensor networks, long distance wireless, ICT4D, simulation

1. INTRODUCTION

In practical applications that require sensor monitoring over long distances such as farming or water quality mon-

itoring in developing regions, the short wireless range provided by WSNs may become a limiting factor in terms of both cost (multi-hop routing over long distances may require many sensors) and coverage (the short range sensors can cover only a few hundred of meters). This paper revisits the problem of long distance wireless sensor network deployment by (1) assessing the relevance of using simulation in planning long distance links and (2) proposing a long distance wireless sensor deployment as case study. The main contributions of our paper are twofold. First, using the Radio Mobile simulation software, we evaluate the accuracy of using a simulation package that builds around real maps to preplan long distance wireless sensor links. Secondly, we present a case study of a long distance wireless sensor network deployment using the Waspnote [1] technology with experiments conducted in harsh conditions.

2. SIMULATION OF THE LINKS

To check if radio links were feasible, we decided to use Radio Mobile [2], a free tool for the design and simulation of wireless systems. It predicts the performance of a radio link by using information about the equipment and a digital map of the area. Radio Mobile uses a digital terrain elevation model for the calculation of coverage, indicating received signal strength at various points along the path. It automatically builds a profile between two points in the digital map showing the coverage area and first Fresnel zone. During the simulation, it checks for line of sight and calculates the Path Loss, including losses due to obstacles. It is based on the ITIS (Longley-Rice) propagation model. Digital elevation maps (DEM) are available for free from several sources, and are available for most of the world [3].

2.1 Candidate locations

To test the feasibility of long wireless sensor links, it is necessary to find a location with an unobstructed line-of-sight between two sites. For our experiments we selected the Los Monegros Desert near Huesca, Spain. Los Monegros is located within the provinces of Zaragoza and Huesca. The lack of human activity ensured an interference-free environment. We did not carry out a site survey when selecting the candidate locations.

We selected ten spots in the area, which allowed us to establish six links with distances ranging from 300m to 12km. We considered both links with line of sight (LOS) and those

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

NSF-10 San Francisco, California USA
Copyright 2010 ACM 978-1-4503-0193-0 ...\$10.00.

Other Considerations for the Abstract

- It takes time to write a good abstract
- The abstract is a unique piece of text: do not copy content from elsewhere in the paper
- Style
 - Explain acronym first time it appears or use the full description
 - No details: no quotations, references, etc.
 - Include relevant search terms
- Keywords (below the abstract)
 - Less important now with advanced search engines
 - Used for assignment of review committee – be strategic

Structure of the Abstract

1. **Motivation:** Why do we care about the problem? This section sets the stage and specifies the area your work fits into.
 2. **Problem Statement:** What problem are you trying to solve? What is the scope of your work?
 3. **Approach:** What did you actually do to get your results?
 4. **Results:** What did you learn, create, or invent? What's the answer? Put the result there, in numbers.
 5. **Implications:** What are the larger implications of your findings, especially for the problem or gap identified?
- As a rule of thumb, each part should be one sentence, but feel free to be creative

Example 1

Motivation

Problem

Approach

Results & Implications

Soft-input soft-output algorithms are the principal components of the iterative decoding used in turbo codes and other ‘turbo’ feedback schemes. To enable an efficient implementation, especially on energy constrained platforms such as portable devices, it is crucial to reduce the computational complexity to a minimum. We propose an enhancement to the MAX-LOG-MAP algorithm by adding a traceback operation similar to that used in the Viterbi algorithm, and devise a new efficient way to initialize the start state of the traceback. This enhancement is effective for each decoding iteration, and provides saving on top of existing techniques such as early termination and memory optimizations. It reduces the computational complexity by an additional 15%, without incurring any performance penalty.

Example 2

Motivation
Problem
Approach
Results & Implications

In their quest to study the intricate underwater ecosystems, oceanographic scientists need spatially-rich data that is collected within the moving reference frame created by the oceans' currents. To this end, we are developing a system of networked underwater ocean explorers that drift freely with the currents. However, effective interpretation of collected sensor data hinges on knowing the positions of these drifters while submerged. Due to their uncontrollable motion, positions have to be tracked in time. Unlike in static networks, position estimation therefore represents a recurring cost, which should be minimized for reasons of limited on-board energy supplies. In this paper, we present a low-overhead scheme that is able to trade accuracy for energy savings by cleverly selecting the specific links that are required for the self-localization algorithm. Our technique is also agile enough to quickly adapt to on-demand changes in accuracy requirements, and cuts the position estimation costs by 40% or more under various ocean current dynamics.

Example 3

Results & Implications

Approach

Problem

Motivation

System-level power management has become a key technique to render modern wireless communication devices economically viable.

Despite their relatively large impact on the system energy consumption, power management for radios has been limited to shutdown-based schemes, while processors have benefited from superior techniques based on dynamic voltage scaling (DVS).

However, similar scaling approaches that trade-off energy versus performance are also available for radios. To utilize these in radio power management, existing packet scheduling policies have to be thoroughly rethought to make them energy-aware, essentially opening a whole new set of challenges the same way the introduction of DVS did to CPU task scheduling. We use one specific scaling technique, dynamic modulation scaling (DMS), as a vehicle to outline these challenges, and to introduce the intricacies by the non-preemptive nature of packet scheduling and the time-varying wireless channel.

Introduction

- Five main elements
 - What is the problem?
 - Why is it interesting and important?
 - Why is it hard? (E.g., why do naive approaches fail?)
 - Why hasn't it been solved before? (Or, what's wrong with previous proposed solutions? How does mine differ?)
 - What are the key components of my approach and results?
Also include any specific limitations.
- Write about a paragraph for each main elements

Introduction (cont.)

- You need to say within the first page or two (single column) or first page (double column) what you are doing (the abstract does not count)
- You can be more informal in the introduction (*my personal preference*)
- Forward references are very useful in the introduction
- (*my preference*) Do not end introduction with

The rest of this paper is organized as follows. We first discuss motivation and related work in Section 2. In Section 3 we describe our hardware measurement configuration, benchmarking, and the resulting Mica2 energy model. In Section 4 we describe the design of TOSSIM and our modifications to it. Section 5 describes the implementation of PowerTOSSIM in detail, and Section 6 presents experiments to validate its accuracy and scalability. Section 7 describes future work and concludes.

1.1 System Power Management

In commercial and research efforts alike, recent trends have shifted the emphasis in system design away from ever-increasing throughput alone. Instead, energy and power consumption are becoming ever more formidable and important constraints to address [Benini and de Micheli 1997; Pedram 2001]. Indeed, technology integration and miniaturization have intensified cooling and packaging challenges. In addition, the energy supply itself is often limited to build-in batteries, as untethered operation is desired for mobility and portability. A class of systems where this observation holds especially true is that of personal communication devices, which are also becoming increasingly prevalent.

To reduce power spending without sacrificing performance, system designers have drawn on the concept of “energy or power awareness.” The key idea is that the system only delivers the performance that is strictly required, thereby avoiding superfluous power consumption. Power and energy awareness are therefore often paraphrased as “the right power/energy at the right time and the right place.” When designing according to this concept, there are two aspects to consider: the technique itself that introduces the awareness, and the power management strategy that exploits it.

- Low energy design is important for personal communication devices.
- We can achieve this through adaptation, which enables energy awareness.

The oldest and most straightforward technique is to shut down unused parts of the system [Benini et al. 1999]. Shutdown-based power management has been explored for hard disks, displays [Lorch and Smith 1998], and communication modules [Wang and Mandayam 2001], among others. For processors, it has been incorporated in the kernel of real-time operating systems (RTOS) [Srivastava et al. 1996]. However, a breakthrough came with the development of dynamic voltage scaling (DVS), a technique for digital circuits that is more effective than shutdown [Chandrakasan et al. 1992]. The reason is the convex nature of the power–speed curve, which comes about by varying the operating voltage. Numerous DVS-based power management strategies have been proposed to harness this potential [Burd et al. 2000; Govil et al. 1995; Gruian 2001; Gutnik and Chandrakasan 1997; Krishna and Lee 2000; Manzak and Chakrabarty 2000; Nielsen et al. 1994; Shin and Choi 1999; Weiser et al. 1994; Yao et al. 1995].

Unfortunately, energy awareness brought forth by DVS is restricted to digital circuits, such as processors, which can leverage DVS in their task-scheduling engine. Voltage scaling not being applicable; is it nevertheless possible to find a similar technique for other parts of the system? Or is shutdown the best we can do?

- This adaptation appears in two forms, and slowdown, which is also called scaling, is more effective than shutdown.
- However, scaling has only been shown for processors and not for other parts of an embedded system.

1.2 Power Management in Communication Subsystems

The search for superior power management techniques is especially relevant for the communication subsystem, and in this paper, we introduce and discuss a powerful scaling-based approach that can be viewed as the counterpart of DVS.

The importance of radio power management arises from the fact that communication is often the dominant power hog in the complete system [Raghunathan et al. 2002]. While the inherent energy consumption of digital circuits is rapidly decreasing due to Moore's law and ingenious design techniques, the power that is radiated to carry information does not follow this trend. This observation is particularly true for wireless RF (radio-frequency) devices, where the radiated power depends on the transmit distance and is physically constrained by Maxwell's laws. With the increasing proliferation of cellular phones with Internet browsing capabilities, Bluetooth-connected PDAs, and wireless LANs, radio power management will undoubtedly gain in importance.

- To remedy this issue, we will create scaling for the communication subsystem.
- This is especially important since communications is a major power hog of the entire system.

Radios typically possess a number of control knobs, which gives rise to a convex power-performance curve. As such, these knobs can be used to introduce energy awareness in a similar fashion as the operating voltage in DVS. Possible radio control knobs are the modulation level, the error coding, or combinations, or both. We refer to the resulting techniques as dynamic modulation scaling (DMS), dynamic code scaling (DCS), and dynamic modulation-code scaling (DMCS). To provide energy awareness, they need to be integrated into the system power management.

The aim of this paper is to provide insight into such scaling-based radio power management strategies, the associated challenges, and possible solutions. Just as people created energy-aware versions of RTOS task-scheduling policies, we investigate energy-aware packet scheduling as part of these power management strategies. As a vehicle for this exploration, we focus on dynamic modulation scaling (DMS), which we first introduced in Schurgers et al. [2001a]. It is a readily accessible control knob in a number of existing communication systems.

- Fortunately, radios allow you to control some settings, thereby enabling adaptation and scaling.
- But to really leverage this ability, we need to design new power management strategies and scheduling techniques.

Paper Title

Author 1

Author 2

Author 3

Author mailing address line 1

Author mailing address line 2

Author email addresses

ABSTRACT

System level power management has become a key technique to render modern wireless communication devices economically viable. Despite their relatively large impact on the system's energy consumption, power management for radios has been limited to shutdown-based schemes, while processors have benefited from superior techniques based on Dynamic Voltage Scaling (DVS). However, similar scaling approaches that trade off energy versus performance are also available for radios. To utilize these in the radio power management, existing packet scheduling policies have to be thoroughly rethought to make them energy-aware, essentially opening a whole new set of challenges the same way the introduction of DVS did to CPU task scheduling. We use one specific scaling technique, Dynamic Modulation Scaling (DMS), as a vehicle to outline these challenges, and to introduce the intricacies caused by the non-preemptive nature of packet scheduling and the time-varying wireless channel.

I. INTRODUCTION

1. System Power Management

In commercial and research efforts alike, recent trends have shifted the emphasis in system design away from ever-increasing throughput alone. Instead, energy and power consumption are becoming ever more formidable and important constraints to address [2][22]. Indeed, technology integration and miniaturization have intensified cooling and packaging challenges. In addition, the energy supply itself is often limited to build-in batteries, as untethered operation is desired for mobility and portability. A class of systems where this observation holds especially true is that of personal communication devices, which are also becoming increasingly prevalent.

To reduce power spending without sacrificing performance, system designers have drawn on the concept of 'energy or power awareness'. The key idea is that the system only delivers the performance that is strictly required, thereby avoiding superfluous power consumption. Power and energy awareness are therefore often paraphrased as 'the right power/energy at the right time and

the right place'. When designing according to this concept, there are two aspects to consider: the technique itself that introduces the awareness, and the power management strategy that exploits it.

The oldest and most straightforward technique is to shut down unused parts of the system [3]. Shutdown-based power management has been explored for hard disks, displays [17], and communication modules [37], amongst others. For processors, it has been incorporated in the kernel of Real-Time Operating Systems (RTOS) [33]. However, a breakthrough came with the development of Dynamic Voltage Scaling (DVS), a technique for digital circuits that is more effective than shutdown [5]. The reason is the convex nature of the power-speed curve, which comes about by varying the operating voltage. Numerous DVS-based power management strategies have been proposed to harness this potential [4][9][10][11][15][19][20][31][39][40].

Unfortunately, energy-awareness brought forth by DVS is restricted to digital circuits, such as processors, which can leverage DVS in their task-scheduling engine. Voltage scaling not being applicable, is it nevertheless possible to find a similar technique for other parts of the system? Or is shutdown the best we can do?

2. Power Management in Communication Subsystems

The search for superior power management techniques is especially relevant for the communication subsystem, and in this paper, we introduce and discuss a powerful scaling-based approach that can be viewed as the counterpart of DVS.

The importance of radio power management arises from the fact that communication is often the dominant power hog in the complete system [25]. While the inherent energy consumption of digital circuits is rapidly decreasing due to Moore's law and ingenious design techniques, the power that is radiated to carry information does not follow this trend. This observation is particularly true for wireless RF (radio-frequency) devices, where the radiated power depends on the transmit distance and is physically constrained by Maxwell's laws. With the increasing proliferation of cellular phones with Internet browsing

capabilities, Bluetooth-connected PDA's, and wireless LANs, radio power management will undoubtedly gain in importance.

Radios typically possess a number of control knobs, which give rise to a convex power-performance curve. As such, these knobs can be used to introduce energy awareness in a similar fashion the operating voltage did in DVS. Possible radio control knobs are the modulation level, the error coding, or combinations of both. We refer to the resulting techniques as Dynamic Modulation Scaling (DMS), Dynamic Code Scaling (DCS), and Dynamic Modulation-Code Scaling (DMCS), respectively. To provide energy awareness, they need to be integrated into the system power management.

The aim of this paper is to provide insight into such scaling-based radio power management strategies, the associated challenges and possible solutions. Just as people created energy aware versions of RTOS task scheduling policies, we investigate energy-aware packet scheduling as part of these power management strategies. As a vehicle for this exploration, we focus on Dynamic Modulation Scaling (DMS), which we first introduced in [28]. It is a readily accessible control knob in a number of existing communication systems.

II. RELATED WORK

DMS for radios can be viewed as the counterpart of DVS for digital circuits, as both exploit the presence of a convex power-speed curve via the notion of scaling. We can therefore find inspiration in the way task scheduling was extended to make it energy-aware, when developing energy-aware packet-scheduling policies. Nevertheless, a direct correspondence is impossible due to the inherent differences between radios and digital circuits, which we will detail in section III.4. Voltage scaling was first included in self-timed circuits [20] and later extended toward synchronous ones [11], where a buffer is used to smooth the load and steer the adaptation. The initial DVS work in operation system research was in the context of a workstation-like environment [9][39], with average throughput as the performance metric. Task scheduling under real-time constraints, i.e., the presence of deadlines, is considered in [4][10][15][19][31][40].

In the realm of communications, a shutdown approach is proposed in [37], leveraging the time-varying nature of the wireless channel. Transmissions are deferred to times where the channel can support energy efficient data transmission, while taking into account various delay constraints. A scheduling technique based on code scaling (DCS) is described in [23]. It finds an energy-efficient schedule for a set of packets that have to be transmitted

before an overall deadline, and is similar to the work presented in [40] on processor scheduling. The basic concept of modulation scaling (DMS), we proposed it first in [28]. Here, we build upon this concept to describe a complete framework for dynamic power management of the radio communication subsystem in wireless embedded systems.

We explore radio power management through DMS, which essentially relies on the ability of the radio to change its modulation on the fly. This is indeed practically feasible, and appropriate hardware architectures have been developed [6]. Originally, these architectures were used for a different purpose, namely adapting the modulation to maximize the system throughput [1][36][38]. In this case, the appropriate choice of the modulation level only depends on the current condition of the wireless channel, and does not involve any scheduling decisions. On the other hand, as we will illustrate in this paper, energy awareness is strongly related to scheduling, where current decisions and future events are tightly interwoven. Although utilizing the same radio control knob, DMS and throughput maximization not only serve a different purpose, but also require completely different adaptation policies. This change in mindset is similar to the one that occurred in the CPU world when the design goal switched from throughput maximization to energy efficiency.

III. DYNAMIC MODULATION SCALING (DMS)

When considering energy-aware techniques beyond shutdown, such as DMS, we require some more insight into the operation of radios. In this section, we introduce the basics of modulation scaling.

1. Energy and Delay Breakdown

The first step in understanding DMS is investigating how energy and throughput depend on the modulation level in a digital wireless communication system. In order to transmit information, bits are coded into channel symbols, which correspond to different waveforms [24]. The number of possible waveforms determines how many bits are coded into one symbol, which is given by the modulation level b , expressed in number of bits per symbol. The average time to transmit one bit over the channel is the inverse of the average bit rate R_b , and given by (1), where R_s is the symbol rate (number of symbols that are transmitted per second).

$$T_{\text{bit}} = \frac{1}{R_b} = \frac{1}{b \cdot R_s} \quad (1)$$

The energy associated with the transmission of one bit can be expressed as (2). The power needed to generate the information carrying electro-magnetic waves, is delivered

Related Work Section

- This section can be at two different places in the paper
1. Beginning
 - If it can be short yet detailed enough
 - If it's critical to take a strong defensive stance about previous work right away
 - Part of introduction or as a separate section after the introduction
 2. End
 - If it can be summarized quickly early on or used as part of the flow of the introduction
 - Put a forward reference
 - Just before the conclusions

Body

- **Guideline #1**
 - A clear important technical contribution should have been articulated by the time the reader is a quarter of the way through the paper
- **Guideline #2**
 - Every section of the paper should tell a linear story. It should look forward to the next step, and provide linkages.
 - There should be no significant interruptions, those can go in the appendices

Conclusions

- Fairly important: reader can skip ahead to this
- Don't copy text from the paper
- Include future work, but I don't always do this