



UNIVERSITÀ DEGLI STUDI DI PADOVA

DIPARTIMENTO DI INGEGNERIA

DELL'INFORMAZIONE

CORSO DI LAUREA IN INGEGNERIA DELLE
TELECOMUNICAZIONI

Green Radios Systems

Laureando:

Fabio *Penzo*

Relatore:

Prof. Nevio *Benvenuto*

Anno accademico 2009/2010

*Alle persone che mi hanno sostenuto e soprattutto a quelle che non l'hanno
potuto fare.*

Green Radios Systems

Contents

1	Overview	4
2	Introduction	5
3	Design step to enable flexibility	8
3.1	Control step to exploit flexibility	9
4	The need for reconfigurable radio architectures	10
4.1	Energy-scalable reconfigurable radio architectures	13
4.2	SDR-Based digital baseband engines	14
5	Reconfigurable analog front-ends	19
5.1	Reconfigurability for energy scalability design	20
6	Energy-aware cross-layer radio managment	23
6.1	State-of-the art energy managment techniques	25
6.2	System-level cross-layer performance-energy optimization . . .	27

7	Instantiation in concrete use cases	30
7.1	Uplink data transmission over a wireless link	30
7.2	Uplink multiuser video transmission over a WLAN	32
8	Toward energy-efficient spectrum sharing and seamless connectivity	34
9	An example : Energy-efficient packet transmission over a wireless link	38
10	Conclusions	43

1 Overview


This article will provide an essential overview of green radios systems (GRS) implementations scenarios. The basic needs and design step for GRS will be shown in the first sections, emphasizing the methods to control the energy-performance tradeoff of the system. Section 7 will provide some example of GRS use in concrete cases, and in Section 9 several practical rules for energy-efficient packet transmission will be illustrated.

2 Introduction

Wireless communication has witnessed a continuous trend toward application diversification, leading to a significant growth in usage and services. With the proliferation of devices like SmartPhones, PDAs, laptops and multimedia applications, user demand is progressively shifting from simple data-rate increase to complex and heterogeneous quality of service (QoS) requirements. Mobile markets in developing countries are currently experiencing explosive growth in mobile usage. In China for example, there will be 10-20 times the number of subscribers in the UK, leading to proportionately higher power consumption reflecting the size of the country. Carbon footprint will become an increasingly important constraint in future system design. With these devices being battery-powered, the performance requirements are coupled with severe constraints on energy efficiency. There exists a continuously growing gap between the available energy, resulting from battery technology evolution, and the steeply increasing energy requirements of emerging radio systems (Figure 1). Moreover, given the worldwide growth in the number of mobile subscribers, the move to higher data rate mobile broadband and the increasing contribution of information technology to the overall energy consumption of the world, there is a need on environmental grounds to reduce the energy requirements of radio access networks. Climate change and the need to reduce energy consumption are increasingly acknowledged. A typical mobile phone network may consume approximately 40 MW, even excluding

the power consumed by the users handsets. Mobile communications thus contributes a significant proportion of the total energy consumed by the information technology industry, which is growing significantly year on year. From an operator perspective, reduced energy consumption translates directly to lower operating expenditure.


Technology scaling, platform improvements, and circuit design progress are not sufficient for bridging this energy gap. A clear need for holistic system-level strategies exists. In parallel, the combination of the increasing need for functional flexibility in communication systems and the exploding cost of system-on-chip (SoC) design will make implementation of wireless standards on multipurpose reconfigurable radios the only viable option in the coming years. Given the energy gap discussed above, a major challenge is to enable low energy reconfigurable radio implementations, suited for handheld multimedia terminals and competitive with fixed hardware implementations. To make such terminals a reality, a two-step approach is advocated. First, effective energy scalability is enabled in the design of the radio baseband and front-end. Second, the scalability is exploited to achieve low-power operation by a cross-layer controller that follows at run time the dynamics in the application requirements and propagation conditions (Figure 2).



[FIG1] The energy gap between required and available energy is growing.

3 Design step to enable flexibility

To enable the translation of functional flexibility into energy scalability, the reconfigurable radio (considering algorithms, architectures, components, and circuits) should first be designed accordingly. For the reconfigurable digital baseband engine, one has to carefully trade off flexibility and energy efficiency: flexibility should only be introduced where its impact on the total average power is sufficiently low or where it offers a broad range of control options that can be exploited effectively later in the control step (targeted flexibility). For the reconfigurable analog front-end, architectures and circuits should be designed for a broad range of requirements in carrier frequency, channel bandwidth, and noise performance with minimal penalty in power consumption, while also offering energy scalability.



[FIG2] Reconfigurable and energy-scalable radio solutions achieve low-energy operation through cross-layer joint QoS and energy management.


3.1 Control step to exploit flexibility

To exploit flexibility for saving energy, it is mandatory to control the reconfigurable radio system as function of the operation conditions. The key observation here is that wireless communications systems typically face very dynamic conditions (in terms of propagation environment and of application requirements). By carefully adapting the system to these dynamics at run time, capitalizing on the energy scalability discussed above, much energy can be saved compared to conventional design. This problem has to be addressed from a cross-layer perspective, as measuring performance requires taking into account the characteristics of the protocol stack, whereas optimizing energy expenditure assumes detailed knowledge of the low-level radio hardware. Besides the energy constraint, spectrum is also becoming a major resource bottleneck, due to the accelerated deployment of broadband personal communication and the continuously increasing demand for higher data rates. New paradigms for efficiently exploiting the spectrum are clearly needed. A current trend is the evolution toward dynamic and open access to spectrum, motivated by underutilization of many licensed frequency bands. This has led to the concept of cognitive radio. Such CognitiveRadios(CR) will extend the flexibility enabled by new reconfigurable radio architectures by adding intelligent control solutions. Additional requirements will make energy-efficient implementations even more challenging: such CRs will have to couple energy savings and efficient spectrum usage. The design framework introduced above is clearly paving the way toward this challenging goal.

4 The need for reconfigurable radio architectures

Wireless communications are routinely used today for a large variety of applications including voice, data transfer, Internet access, audio, and video streaming, etc. Pushed by the insatiable demand for bandwidth and pulled by the steady improvement of semiconductor technology (Moore's law), the performance offered by wireless standards is due to improve with the years, seemingly without bounds, as illustrated in Figure 3. The observed and predicted evolution of the main classes of wireless standards is indicated as thick arrows (from bottom to top): wireless personal area networks (WPANs), wireless local area networks (WLANs), wireless metropolitan area networks (WMANs), and cellular networks. Future communications systems will have to seamlessly and opportunistically integrate these multiple radio technologies. From a user perspective, it is indeed very attractive to have a single handheld device that can support a large number of these wireless standards enabling ubiquitous the connectivity through seamless horizontal and vertical handovers. The concept of software-defined radio (SDR) in which a common hardware platform is used for different standards supports this trend. A closer inspection of standards specifications reveals that a large variety of bit rates, modulation formats, physical bandwidths, and carrier frequencies has to be supported. In addition, several parallel transceivers and antennas are required to offer, for instance, multiple-input/multiple-output

capabilities. In parallel to the increasing need for functional flexibility in radio systems, progress in CMOS technology opens opportunities to implement flexible platforms, as new transistors can be made ever faster and smaller. Yet, the newer technology nodes urge the need for high-volume, multipurpose devices, as the costs related to SoC design and manufacturing are rising exponentially. The combination of the increasing need for functional flexibility and these huge costs related to SoC design will make implementation of wireless standards on SDR platforms the only viable option in the future. In light of the aforementioned requirements (functional flexibility and platform programmability), one of the major challenges of SDR clearly resides in energy consumption.



[FIG3] The variety of wireless standards.


4.1 Energy-scalable reconfigurable radio architectures

In this section, I will provide an overview of reconfigurable mobile terminal designs. Will be considered both the SDR-based digital baseband engine and reconfigurable analog front-end circuits. In both cases, will be introduced techniques to enable energy scalability.

4.2 SDR-Based digital baseband engines

When considering reconfigurable digital baseband engines, one has to carefully trade off flexibility and energy efficiency. To achieve an acceptable energy efficiency for mobile devices (100 MOPS/mW on average), flexibility should only be introduced in the platform where its impact on the total energy cost is marginal or where it enables control options that can be effectively exploited in the control step (targeted flexibility). Indeed, flexibility offered by fine-grained adaptive algorithms and implementations may, at the system level, be more energy efficient than fixed nonadaptive hardware solutions. This comes from the fact that these flexible solutions have the potential to continuously adapt to the environment and application dynamics for saving energy. Different processor architecture styles have been proposed for SDR modems. Most of them are designed keeping in mind the specific features of physical (PHY) layer signal processing algorithms in wireless standards. By combining these architecture styles, one can obtain a trade-off among energy efficiency, performance, and mapping productivity. Fine-grained reconfigurable arrays (FGAs) and coarse-grained reconfigurable arrays (CGAs) are mainly motivated by the data flow dominance in baseband algorithms. The main bottleneck of the first class is the high interconnect cost that hampers scalability in deeper submicron technology (65 nm and below) and yields significant energy overhead. The second class improves on this, proposing fewer but more complex functional units. Importantly, C-based compilers for CGA architectures recently became available. Very long in-


struction word (VLIW) architectures and vector architectures, also referred to as single-instruction multiple-data (SIMD) architectures, are exploiting the high data-level parallelism of typical algorithms for wireless communications. The first class benefits from well-established compiler technology, giving it an edge in programmer productivity, while the second class improves on energy efficiency. Their combination, achieves a trade-off between productivity and energy efficiency and is very well fitted for SDR implementations. Next to these solutions for modulation/demodulation processing in the PHY layer, the platform has to perform the other tasks required in a wireless modem, which can have very different characteristics. The medium access control (MAC) processing is, for instance, mainly control-dominated. On the other hand, packet detection and time synchronization have a significantly higher duty cycle. The required sub-functions of the wireless modem should be designed according to their nature (i.e., control or data processing) and flexibility/energy efficiency requirements. This calls for heterogeneous multiprocessor system-on-chip (MPSoC) platforms. The platform template is depicted in Figure 4. The different cores, connected by a segmented bus, are listed below.



[FIG4] Example of opportunistically partitioned SDR platform.

- Digital front-end tiles achieve packet detection with limited need in flexibility and code-mapping productivity ($<5\%$ of the total application code) but very high energy-efficiency requirements. They are heavily exploiting subword parallelism and techniques to reduce the instruction and data memory overhead.
- Baseband engines handle the burst processing based on the combination of the aforementioned architectures (CGA processor coupled with vector processing), which provides an appropriate trade-off between performance, energy efficiency, and flexibility.
- Forward error correction (FEC) processors are implemented on an application-specific architecture for advanced FEC algorithms (turbo and LDPC-codes). These are indeed hard to map to typical baseband processors (mainly because of the extensive interleaving requirements). They also require an effective solution for the severe data-access bottleneck.
- A platform controller is responsible for the MAC functionality and

the PHY macro-pipeline scheduling. This typically runs on an RISC processor. When combined with adequate power-management techniques, this partitioning enables targeted flexibility with minimal impact on the average power consumption and standby time. Further energy gains are achievable by also applying the energy-scalable design paradigm when implementing the baseband processing algorithms onto the baseband engine. Energy-scalable algorithms enable trading off system performance and implementing energy cost (i.e., execution energy). A linear controller is designed to adapt the rate at which the equalizer coefficients are updated as a function of the observed equalization error. As a result, the computation load is automatically adapted to the propagation conditions, saving up to 60% of the computation without hampering performance. Also, other groups have considered the use of adaptive algorithms to reduce the power consumption. For instance, introduces dynamic algorithm transformations for low-power signal processing. The focus is, however, on the algorithm itself and not on the optimal matching with the necessary run-time controller that requires a global approach.




[FIG5] Conceptual view of a state-of-the-art reconfigurable front-end.

5 Reconfigurable analog front-ends

A high-level block diagram of a state-of-the-art reconfigurable analog front-end is shown in Figure 5. As for all radio frequency (RF) transceivers, it is required to carefully analyze frequency planning, filtering, and the radio budget for noise and nonlinearity. This leads almost automatically to the so-called reconfigurable zero-IF architecture because of the cost constraints inherently present in radio transceiver design. It is also illustrated in Figure 5 that the reconfigurable frontend needs to include more than the RF integrated circuit (RFIC) alone. Indeed, when considering the different blocking levels the receivers need to cope with, RF filtering is needed between the antenna and the low-noise amplifier (LNA) on the RFIC. This is already the case in single-mode receivers, which often have to cope with an out-of-band blocking level (typically 0 dBm), which has to be attenuated by an antenna filter before it can be amplified by the LNA (typically 20 dB). SDRs will suffer even more from these blocking requirements, since inside one terminal several radios can be active at the same time. A transmission by, for instance, a WLAN radio will severely block the reception of a cellular communication, and, consequently, the blocking requirements of an SDR receiver will be even tougher than those of a single-standard receiver. That is why the architecture includes a reconfigurable antenna filter, implemented for obvious reasons not on silicon but on a high-quality thin-film substrate, and co-designed with the RFIC (right part of Figure 5).

5.1 Reconfigurability for energy scalability design




[FIG6] Baseband-channel select filter current consumption and input referred noise level (IRN) as a function of cutoff frequency for four different power modes.

An example of reconfigurability for energy scalability is a flexible power amplifier (PA) that was introduced in [2]. An IEEE 802.11a-compliant direct-conversion transmitter was considered. As depicted in the left part of Figure 7, it consists of the cascade of the following stages: a differential I/Q mixer upconverting and combining the in-phase (I) and quadrature (Q) baseband signals into a single RF signal; a PA driver stage; and an external PA stage. Both the driver and external PA are made flexible in terms of performance (output power and linearity) versus dc power consumption. The digital lines to control this flexibility are represented in the left part of Figure 7. Four techniques to change the circuit characteristics have been applied (right part of Figure 7); namely, tunable degeneration resistance, tunable load resistance, tunable bias current, or voltage and tunable supply voltage. As shown by measurements provided in [2], the proposed circuit flexibility at the PA


level enables significant trade-offs between the transmitter performance metrics, namely, the output power and the signal-to-noise-and-distortion ratio (SINAD) at the transmitter (capturing the linearity), and the corresponding energy cost measured by the dc power consumption. Based on appropriate pruning at designand calibration-time, assuming that the average path loss is known [2], one can then carry out the optimization of the trade-off between the effective link performance metric; namely, the link SINAD and the transmitter power consumption P_{dc} . The resulting tradeoffs are reproduced in Figure 8 for average path losses ranging from 60 to 90 dB. There remain around 10 well-distributed working points per trade-off curve, making the energy-scalable transmitter usable for system-level energy management. With a path loss of 80 dB, for instance, when configured to transmit at 36 Mbps (16-QAM, code rate 3/4), the transmitter energy consumption is 25 nJ/bit while it scales down to 16.6 nJ/bit (33% improvement) when configured for 24-Mbps transmission (16- QAM, code rate 1/2).

Another example of reconfigurability for energy scalability was introduced in [3] by adding significant digital control to a baseband-channel select filter. This filter is able to change its component values according to the instantaneous requirements and hence enables a large trade-off range in bandwidth, noise level, and power consumption. For a fixed capacitor size, power consumption in the filter scales linearly with cutoff frequency, while the integrated input-referred noise level stays fixed. Consumed power and generated noise can be traded linearly by digitally changing the capacitor sizes used, as

illustrated by the performance in the four power modes shown in Figure 6.




[FIG7] (a) Flexible OFDM transmitter and (b) simplified circuit diagram.



[FIG8] Power consumption of the transmitter front-end versus link SINAD trade-off for various path losses.

6 Energy-aware cross-layer radio management

In this section, we start by providing an overview of state-of the-art energy-saving techniques for wireless communication systems, emphasizing the need for a system-level approach. Next, we propose a framework for energy-aware cross-layer radio management, enabling the exploitation of the flexibility introduced in the previous section for saving energy. As pointed out in the beginning of this article, the key idea is that significant energy savings are possible by continuously adapting the energy-scalable reconfigurable radio to the actual environment conditions and performance constraints at run time (Figure 2), exploiting cross-layer synergies. This corresponds to a major paradigm shift system design to adaptively provide the just-required application performance while minimizing energy consumption. This will be illustrated through different realistic case studies.



[FIG9] Top-down and bottom-up approaches can be contradictory.

6.1 State-of-the art energy management techniques

The need for improving the energy efficiency of wireless communication devices has already triggered a lot of research at various levels, from circuit to communication theory and networking protocols. The energy management problem, in its most general formulation, consists of dynamically controlling the system to minimize the average energy consumption under a performance constraint. Existing research can be classified into two categories:

- Top-down approaches. Approaches that are intrinsically utilization and hardware-aware but communication-unaware are categorized as top-down. The communicating device is treated as any electronics circuit and general-purpose techniques like dynamic power management and energy-aware design are applied. The first technique is defined as dynamically reconfiguring an electronic system to provide the requested performance levels with a minimum number of active components and the minimum loads on those components. The second technique can be defined as designing systems that present a desirable energy–performance behavior for energy management.
- Bottom-up approaches. Approaches that are intrinsically communication-aware but hardware-unaware are categorized as bottom-up. They rely on the fundamentals of information and communication theory to derive energy-aware transmission techniques and communication algorithms. We find here, for instance, the transmission scaling techniques,

which exploit the fundamental trade-off that exists between transmission rate/power and energy [4]. Network power management techniques also fall in this category, targeting the minimization of the transmission power under QoS constraints.

Top-down and bottom-up approaches can easily result in a fundamental contradiction when applied independently. A good example is the conflict between transmission scaling at the PHY layer (bottom-up) and sleeping schemes at the MAC layer (top-down). Scaling tends to minimize transmission energy consumption by transmitting with the lowest power over the longest feasible duration, whereas sleeping tends to minimize the duty cycle of the radio circuitry by transmitting as fast as possible (Figure 9). Clearly, the two techniques are contradictory when it comes to defining the optimal transmit rate and power allocation. The above reveals the need for a holistic cross-layer approach to carry out system-level energy management, jointly considering techniques to optimize the energy consumption in the different layers. This conclusion is strengthened by the simple observation that user-relevant performance metrics can only be measured on top of the whole protocol stack, close to the user, while energy consumption is mainly conditioned by the setting and policies within the lowest protocol layers, close to the hardware. The topic of cross-layer optimization has recently gained a lot of interest. However, existing approaches still miss the advocated holistic view, hampering real-life deployment and sometimes leading to undesired side effects.

6.2 System-level cross-layer performance-energy optimization

We now introduce an energy-aware cross-layer management framework dedicated to wireless systems. The proposed approach tries to handle cross-layer interactions in order to minimize energy consumption under performance constraint. A key contribution to tackle optimization complexity is the so-called Pareto-based design-time/run-time principle [5], which was recently applied to the context of wireless communications. The general idea is to shift most of the complexity to design time by carrying out a systematic exploration leading to control solutions that result in minimum complexity at run time. The proposed approach can be summarized as follows.

- **SCALABLE ARCHITECTURE DEFINITION.** A first step is to enable run-time controllable parameters (or knobs) that significantly impact performance and energy consumption at system level. The control dimension settings in real implementations are discrete, interdependent, and can have a nonlinear influence.
- **DISCRETIZATION OF RUN-TIME DYNAMICS.** The second important design-time step is the discretization of the system run-time dynamics (e.g., environment conditions, application requirements) through the definition of relevant external variables that can be tracked at run time.
- **PROBLEM MODELING AND PARTITIONING.** The system energy and

performance behavior are modeled in function of these configuration parameters and external variables. A hierarchy of power-management-specific abstraction layers is then introduced to partition the global optimization problem into subproblems that can be solved locally. A valid abstraction layer can be characterized by intermediate metrics that can be optimized locally without significantly hampering the optimality of the global optimization. It is hence possible to derive the optimal points along these local metrics (according to the Pareto multiobjective optimality criterion) and systematically prune away the nonoptimal configurations. This enables a drastic complexity reduction of the next optimization steps.

- **CONTROL POLICY DEFINITION.** The three first steps result in sufficient system characterization to conduct the global optimization. Given the Pareto-optimal configuration points for each abstraction layer in the considered partitioning, the fourth step then consists of deriving the policy that will be used at run-time to determine the optimal configuration for the whole system. This requires determining 1) how to communicate the local information and 2) how to combine that information to obtain a global optimum.
- **RUN-TIME OPERATION.** The run-time control algorithm tracks the external variables, selects current scenario, and runs the optimization policy to select the optimal configuration points of each layer in order


to provide the just required performance at minimum energy consumption. This results in low complexity overhead thanks to the design time/run-time partitioning, and especially the availability of a set of potential operating points that is fully characterized up front and that is organized in monotonous Pareto curves.

7 Instantiation in concrete use cases

7.1 Uplink data transmission over a wireless link

The first case study considers the uplink transmission of data traffic over an 802.11a wireless link. An energy-efficient run-time radio link controller is derived based on the framework introduced above. The goal is to achieve the minimum energy consumption by jointly adapting the system configuration to the actual application requirements (i.e., average rate requirement) and environment parameters (i.e., channel state information). This will be shown to yield a significant energy efficiency improvement. Given the scalable transmitter introduced above, we can rely on increased performance/energy flexibility in the lower layers of a 54-Mbit/s IEEE 802.11a transceiver. The parameters that influence the system-level energy efficiency and performance can then be identified for the considered case. Next, a discrete channel model for an indoor WLAN link is derived, in order to discretize the environment dynamics. The system is then partitioned into abstraction layers according to the concepts introduced above, as illustrated in Figure 10, and the design-time characterization of the different layers is obtained using sets of Pareto-optimal points. This provides us the necessary information to derive the run-time policy. The resulting trade-off between the average net data rate and the energy efficiency on top of the MAC (i.e., on top of the three abstraction layers) is shown in Figure 11. When compared with traditional WLAN radio link control schemes, the proposed scheme can be shown to


improve the energy efficiency by up to 40%.




[FIG10] Hierarchical problem partitioning and link controller.

7.2 Uplink multiuser video transmission over a WLAN

The second case study consists of multiple users equipped with mobile terminals who want to send real-time video upstream to the access point of a WLAN. They transmit over a shared, slowly fading wireless channel (typical indoor scenario). The IEEE 802.11a is taken for the PHY layer [orthogonal frequency division multiplexing (OFDM)-based, 5-GHz band], and the QoS-enabled IEEE 802.11e for the MAC. Based on the proposed framework, a cross-layer run-time scheduler was developed [6]. The goal is to guarantee reliable and timely delivery of the video streams over the wireless links from the mobile terminals to the access point, while minimizing the overall energy consumption in the terminals. This is achieved by jointly optimizing the shutdown (top-down) and scaling (bottom-up) energy-management strategies. Based on appropriate design-time performance-energy modeling, the controller (located at the access point and relying on the so-called hybrid coordination function of the 802.11e MAC protocol) determines at run-time on a regular basis the allocation of channel resources to the different mobile terminals and the configuration of the MAC and PHY parameters in these devices, adapting to channel and traffic dynamics in the system [6]. In Figure 12, the energy gain of the proposed approach is compared with three transmission strategies (PHY: scaling only, MAC: sleeping only, Fixed: no energy management) with MPEG4 traffic, as a function of the number of users for an average aggregate load of 4Mbps. It can be seen that energy gains ranging between 3 to 10 can be obtained compared to traditional approaches.



[FIG11] Energy efficiency versus average data-rate tradeoff obtained based on the proposed radio link control policy and energy-scalable transmitter.



[FIG12] Energy gain per flow as a function of the number of users for a constant aggregate load.

8 Toward energy-efficient spectrum sharing and seamless connectivity

Besides the energy constraint, spectrum is also becoming a major bottleneck. The search for new paradigms for efficiently exploiting the spectrum has led to the concept of CR, first coined in [7], and much related to our approach. Such a CR is indeed defined as a radio that can autonomously change its transmission parameters based on interaction with the complex environment (radio scene, application and user requirements) in which it operates [8]. By radio scene, we mean, for instance, the changing level and nature of radio interference, available wireless access networks, etc. The following definition is also often encountered: a radio that co-exists with legacy wireless systems by using the same spectrum resources without significantly interfering with them. This is an approach targeting more specifically an efficient use of the spectrum, through interference avoidance and opportunistic re-use of frequency bands originally allocated for licensed operation. Standardization is currently ongoing in the IEEE 802.22 working group, which targets the definition of a cognitive wireless regional area network for operation in the television bands. The spectrum sensing and agile air interface requirements of CR definitely call for SDR-based implementations.

- The need to detect and/or generate virtually any kind of waveform in any band pushes the specification of the underlying reconfigurable radio to the limit.


- Spectrum sensing will be of paramount importance. It enables identification of spectral “holes” over a range of frequencies and determines the usage of that spectrum by other users, both primary and secondary users. These sensing functions have to detect and possibly “localize” communication signals without a priori knowledge. A challenge in waveform detection and spectrum sensing lies in the requirement to simultaneously sense and communicate. Indeed, the communication itself uses a relatively narrow portion of the spectrum while the sensing function has to cover a very broad frequency band, potentially several GHz.
- The large dynamic range due to signals, noise, and interference poses a significant challenge for low power terminal implementation: wideband digitization has been the “Holy Grail” of SDR for years and significant progress is needed in analog-to-digital converter technology. Successive approximation register and pipeline analog-to-digital converters are likely candidates for sufficiently high sample rates and effective number of bits [9] as typically required for SDR and CR (in the range of 108 to 109 samples/s and 8 to 14 bits).
- The transmission properties can also be tailored to make CR more spectrum efficient. In order to increase throughput, it may be needed to exploit discontinuous chunks of spectrum. An interesting approach is noncontiguous OFDM, which allows concatenating nonadjacent fre-

quency bands in a single multicarrier transmitter [10]. This technique is inspired from and can be combined with orthogonal frequency division multiple access.

- A cognitive controller functional block should intelligently configure the platform run-time controllable parameters depending on the context. It requires proper monitors (i.e., observables) of the radio environment (access network detection, channel state information, spectrum scanning information, interference metric) and of the user needs and profiles as inputs. Based on the analysis of these monitors, a system state (including both environment and user requirements) can be extracted and appropriate actions can be carried out, relying on a configuration database. As future wireless networks will be extremely heterogeneous and dynamic, this process should be combined with learning capabilities (e.g., extraction of channel or spectrum occupation statistics) to adapt the response to a given system state. These requirements will make energy-efficient implementations even more challenging: CR will have to couple energy savings with efficient spectrum usage. The framework introduced above definitely enables meeting parts of this requirement. Several extensions will however be needed.
- An important step in the design flow is to decide on the energy and quality metrics that are relevant to exploiting spectrum resources. In particular, these scarce resources will have to be presented in a way

that supports rich spectrum-sharing possibilities; i.e., sharing not only in time but also in frequency space or using modulation schemes such as code division multiple access or ultra wide-band, resulting in more complex decision patterns.

- Since future wireless communications systems are expected to be extremely complex and very dynamic in nature, it will become difficult to effectively predict at design time the optimal set of configuration points, as was the case for the run-time optimization schemes introduced in the previous section. As a result, we will need to embed online calibration techniques to scan the environment and learn how to adapt the run-time decision policy to the environment.
- Finally, further research on distributed run-time optimization approaches is required. Wireless communication is indeed evolving to self-deployment and self-maintenance. Distributed approaches are needed in this context to offer intelligent networking solutions and efficient use of the spectrum. The competition/cooperation trade-offs can be analyzed and optimized by game-theoretic techniques while the distributed power control (over the users and the sub-bands) can be optimized by means of tools from information theory [8].



9 An example : Energy-efficient packet transmission over a wireless link

Consider a wireless node whose lifetime is finite, equal to T , say. Suppose that packets arrive at the node in the time interval and must be transmitted to a receiver before T . In the figure, the arrival times of packets, t_i , are marked by crosses and interarrival epochs are denoted by d_i . Without loss of generality, we assume that the first packet is received at time 0. The node transmits the packets according to a schedule that determines the beginning and the duration of each packet transmission. We seek an answer to the question: How should the transmission schedule be chosen so that the total energy used to transmit the packets is minimized? Let $\varepsilon(q)$ denote the transmission energy per bit for the particular coding scheme that is being used, which has code rate $R = 1/q$ bits/transmission. Hence, q is the number of transmissions per bit. The following are the only assumptions we make.

- 1) $\varepsilon(q) \geq 0$
- 2) $\varepsilon(q)$ is monotonically decreasing in q .
- 3) $\varepsilon(q)$ is strictly convex in q .

Assumption 1) is obvious. We shall now demonstrate that assumptions 2) and 3) hold by considering the energy required to reliably transmit one bit of information over a wireless link. The following example assume a discrete

time additive white Gaussian noise (AWGN) channel model for the wireless link.


Consider an AWGN wireless channel with average signal power constraint P and noise power N . As is well known [11], the information theoretically optimal channel coding scheme, which employs randomly generated codes, achieves the channel capacity given by

$$C = \frac{1}{2} \log_2 \left(1 + \frac{P}{N} \right) \quad (1)$$

More precisely, given any $0 < \alpha < 1$, information can be reliably transmitted at rate $R = \alpha C$. To determine the energy per bit ε , note that $q = 1/R$ can be interpreted as the number of transmissions per bit. Substituting into (1), we obtain

$$\varepsilon(q) = qP = qN(2^{[2/(\alpha q)]} - 1)$$

It is easy to see that $\varepsilon(q)$ is monotonically decreasing and convex in q , and that as q approaches infinity the energy required to transmit a bit, $\varepsilon_\infty = (2/\alpha) \ln 2$. Figure 14 plots $\varepsilon(q)$ versus q for $N = 1$ and $\alpha = 0.99$. The range of q in the plot corresponds to SNR values from 20 dB down to 0.11 dB. This is a typical range of SNR values for a wireless link. In this range, $\varepsilon(q)$ can be decreased by a factor of 20 by increasing transmission time and correspondingly decreasing power.



[FIG14] Energy per bit versus transmission time with optimal coding

Now we determine the energy-optimal offline schedule for the above model of a finite number of packets to be transmitted in a given finite-time horizon. This offline optimal schedule provides a lower bound on energy that can be used for evaluating the performance of online algorithms. After briefly introducing the basic setup, a necessary condition for optimality is stated (Lemma 2). This motivates the definition of the specific schedule offline optimal-energy schedule(OOE) (Definition 1). OOE is shown to be feasible (Lemma 3), and energy-optimal (Theorem 1). Suppose that the arrival times $t_i, i = 1, \dots, M$ for the M packets arrive in the interval $[0, T)$ are known in advance, before $t = 0$. Assuming equal length packets each with B bits, the offline vector $\vec{\tau}$ so as to minimize $w(\vec{\tau}) = \sum_{i=1}^M w(\tau_i)$. The assumption that $w(\tau)$ decreases with τ trivially implies that it is suboptimal to have $\sum_i \tau_i < T$, for we could simply increase the transmission times of one or more packets and reduce $w(\vec{\tau})$. Hence, we only consider “non-idling” transmission sched-

ules where $\sum_i \tau_i = T$. It is also sufficient to consider FIFO schedules where packets are transmitted in the order they arrive. The FIFO and non-idling conditions combined with the causality constraint, i.e., that packet transmission cannot begin before the packet arrives, yield the following feasibility conditions.

[4, Lemma 1] : A non-idling FIFO schedule $\vec{\tau}$ is feasible if and only if

$$\sum_{i=1}^k \tau_i \geq \sum_{i=1}^k d_i$$

for $k = 1, 2, \dots, M - 1$, and $\sum_{i=1}^M \tau_i = \sum_{i=1}^M d_i$.

[4, Lemma 2] : A necessary condition for optimality is

$$\tau_i \geq \tau_{i+1}$$

for $i \in \{1, 2, \dots, M - 1\}$.

Given interarrival times d_i , $i \in \{1, 2, \dots, M\}$, let $k_0 = 0$, and define

$$m_1 = \max_{k \in \{1, 2, \dots, M\}} \left\{ \frac{1}{k} \sum_{i=1}^k d_i \right\}$$

and

$$k_1 = \max \left\{ k : \frac{1}{k} \sum_{i=1}^k d_i = m_1 \right\}.$$

For $j \geq 1$, let

$$m_{j+1} = \max_{k \in \{1, 2, \dots, M-k_j\}} \left\{ \frac{1}{k} \sum_{i=1}^k d_{k_j+i} \right\}$$

and

$$k_{j+1} = k_j + \max \left\{ k : \frac{1}{k} \sum_{i=1}^k d_{k_j+i} = m_{j+1} \right\}$$

where k varies between 1 and $M-k_j$. We proceed as above to obtain pairs (m_j, k_j) until $k_j = M$ for the first time. Let $J = \min \{j : k_j = M\}$. The pairs (m_j, k_j) , $j = 1, 2, \dots, J$, are used to define a schedule whose transmission times are denoted by $\vec{\tau}^*$, and Theorem 1 shows that $\vec{\tau}^*$ is the optimal offline schedule.

[4, *Definition 1*]: The vector of transmission times $\vec{\tau}^*$ given by $\tau_i^* = m_j$, if $k_{j-1} < i \leq k_j$ is called OOE.

Lemma3: The following holds for $\vec{\tau}^*$ of OOE.

1. It is feasible and $\sum_{i=1}^M \tau_i^* = T$.
2. It satisfies the condition stated in Lemma 2.

[4, *Theorem 1*]: The schedule OOE of Definition 1 is the optimum offline schedule.

10 Conclusions

The wireless standards scene and its evolution strengthens the need for functional flexibility in future radios. Multimode terminals supporting an increasingly large variety of standards (cellular, WLANs, WMANs, WPANs) are subject to a cost increase that is addressed by more flexible radio interfaces. Energy efficiency, however, is the main obstacle to successfully deploying such reconfigurable radios. To address this, it is essential to design energy-scalable SDRs, both for the radio front-end and the digital baseband platform. Complementing this, an essential ingredient is an intelligent controller that optimally exploits this scalability and the run-time dynamics to translate potential energy scalability to actual low-power operation. To realize this goal, we have introduced an energy-aware cross-layer radio management framework. It was instantiated in different case studies, showing the applicability of our approach in realistic setups. Results have shown that substantial gains can be achieved through effective cross-layer optimization and problem partitioning. Next, it was shown that SDRs will play a crucial role in enabling CRs, which will enable saving on both the scarce radio spectrum and battery lifetime. A key building block for the design of such CRs, i.e., the appropriate control intelligence to make the SDR platform cognitive, can be derived by incrementally building on the proposed framework. As a result, green (or environmentally friendly) reconfigurable radio systems will emerge, which offer a wide variety and ubiquitous availability of wireless

services, while overcoming energy and spectrum scarcity.

References

1. A. Dejonghe, B. Bougard, S. Pollin, J. Craninckx, A. Bourdoux, L. Van der Perre and F. Catthoo, "Green Reconfigurable Radio Systems", *IEEE Signal Processing Magazine*, vol. 24, pp. 90 - 101, May 2007.
2. B. Debaille, B. Bougard, and G. Lenoir, "Energy-scalable OFDM transmitter design", in *Proc. DAC 2006*, July 2006.
3. V. Giannini, J. Craninckx, B. Come, J. Compiet, S. D'Amico, and A. Baschiroto, "Flexible baseband low-pass filter and variable gain amplifier for software-defined radio front-end", in *Proc. European Solid-State Circuits Conf.*, Montreux, Switzerland, Sept. 2006.
4. E. Uysal-Biyikogly, B. Prabhakar, and A. El Gamal, "Energy-efficient packet transmission over a wireless link", *IEEE/ACM Trans. Networking*, vol. 10, pp. 487-499, Aug. 2002.
5. F. Thoen and F. Catthoor, "Modeling, Verification and Exploration of Task- Level Concurrency in Real-Time Embedded Systems", *MA: Kluwer*, 1999.
6. S. Pollin, R. Mangharam, B. Bougard, F. Catthoor, L. Van der Perre, I. Moerman, and R. Rajkumar, "MEERA: Cross-layer methodology for energy- efficient resource allocation for wireless networks", *IEEE Trans. Wireless Commun.*, to be published.

7. J. Mitola and G.Q. Maguire, “Cognitive radio: Making software radios more personal”, *IEEE Pers. Commun.*, vol. 6, no. 4, pp. 13–18, Aug. 1999.
8. S. Haykin, “Cognitive radio: Brain-empowered wireless communications”, *IEEE J. Select. Areas Commun.*, vol. 23, no. 2, pp. 201–220, May 2005.
9. B. Le, T.W. Rondeau, J.H. Reed, and C.W. Bostian, “Analog-to-digital converters”, *IEEE Signal Processing Mag.*, vol. 22, no. 6, pp. 69–77, Nov. 2005.
10. R. Rajbanshi, Q. Chen, A.M. Wyglinski, J.B. Evans, and G.J. Minden, “Comparative study of frequency agile data transmission schemes for cognitive radio transceivers”, in *Proc. First Int. Workshop Technology and Policy for Accessing Spectrum (TAPAS 2006)*, Boston, MA, Aug. 2006.
11. T. Cover and J. Thomas, *Elements of Information Theory*, New York: Wiley, 1991.