# GreenDisc: A HW/SW Energy Optimization Framework in Globally Distributed Computation

Marina Zapater , Jose L. Ayala , and Jose M. Moya

**Abstract.** In recent future, wireless sensor networks (WSNs) will experience a broad high-scale deployment (millions of nodes in the national area) with multiple information sources per node, and with very specific requirements for signal processing. In parallel, the broad range deployment of WSNs facilitates the definition and execution of ambitious studies, with a large input data set and high computational complexity. These computation resources, very often heterogeneous and driven on-demand, can only be satisfied by high-performance Data Centers (DCs). The high economical and environmental impact of the energy consumption in DCs requires aggressive energy optimization policies. These policies have been already detected but not successfully proposed.

In this context, this paper shows the following on-going research lines and obtained results. In the field of WSNs: energy optimization in the processing nodes from different abstraction levels, including reconfigurable application specific architectures, efficient customization of the memory hierarchy, energy-aware management of the wireless interface, and design automation for signal processing applications. In the field of DCs: energy-optimal workload assignment policies in heterogeneous DCs, resource management policies with energy consciousness, and efficient cooling mechanisms that will cooperate in the minimization of the electricity bill of the DCs that process the data provided by the WSNs.

### 1 Introduction

Wearable personal health systems for continuous monitoring are widely recognized to be a key enabling ICT technology for next-generation advanced citizencentric eHealth delivery solutions. Through enabling continuous biomedical monitoring and care, they hold the promise of improved personalization and quality of care, increased ability of prevention and early diagnosis, and enhanced patient autonomy, mobility and safety. Furthermore, wearable personal health systems

<sup>\*</sup> Research by Marina Zapater has been partly supported by a PICATA predoctoral fellowship of the Moncloa Campus of International Excellence (UCM-UPM).

can help the eHealth sector realize its potentials in terms of rapid sustained market growth, reduction of healthcare costs and avoidance of unnecessary cost to the public purse.

To provide the necessary accurate, integrated and long-term assessment and feedback, these wearable personal health systems must acquire, monitor and analyze a large number of physiological and metabolic parameters both during physical activity and rest. The number as well as the nature of the parameters of interest depends on the actual biomedical application/scenario and target population, e.g., fitness and illness prevention (healthy people or people at risk), rehabilitation (patients after an event) or disease management (acute- or chronically-ill patients).

It is largely accepted that Wireless Body Sensor Networks (WBSN) will be the underlying common architecture and technology of these personal health systems. More specifically, the WBSN will consist of a number of sensor nodes attached to the subject/patient body. Each WBSN node ensures the accurate sensing and capture of its target physiological data, its (pre-) processing and wireless communication to the wearable Personal Digital Assistant (PDA), which acts as the network coordinator and central data collector. This PDA will be responsible for the storage, organization, complementary analysis and fusion of the collected physiological and metabolic information, its user-friendly representation, and its dissemination to the relevant medical staff or central monitoring service through private and/or public wireless access networks.

The signal processing applications executed in the nodes of the wireless body sensor network require complex operations that should be adapted to the nature of the sensed signal. Therefore, it is needed an efficient application-specific architecture but, at the same time, flexible enough to provide the required performance to every simultaneous process in execution (multiple sensors that capture multiple sources of information per node).

The envisioned setting proposes a worldwide deployment of nodes for constant monitoring of biophysical and environmental variables, as already proposed in some ambitious European initiatives (FET Flagships 2013: Guardian Angels). The large amount of acquired data to be processed and stored, as well as the computing-demanding algorithms developed for the signal classification and the knowledge acquisition, requires the provision of powerful data centers distributed across the territory and communicated with the WBSNs. The incredible cost of operation and cooling of state-of-the-art data centers could be an obstacle for enterprises as the one proposed here or by the European Union, what motivates a research effort on cooling techniques and energy efficiency.

## 2 Power optimization in the processing nodes of the WBSNs

Ultra-low power application-specific processors. In the application context of WBSNs, we can find multiple active elements (sensors and processors) integrated in a single silicon die, as well as the constraint of portable electronics. These facts explain the energy limitations found in these systems. A typical architecture of a WBSN node is composed of: a microprocessor, a data memory hierarchy, an instruction memory subsystem, sensors, actuators, RF transducer and an energy source [5]. In case of WBSNs for bio-medical applications, there is an additional constraint for the energy dissipation. The heat produced by bio-medical systems has to be carefully controlled to avoid any damage in the skin and tissues of the area of placement. A recent work by [21] that studied the energy distribution in an ASIP running a heart beat detection algorithm has shown that the two main sources of power consumption are the memories (especially, the instruction memory) and the functional units (FU). Therefore, the optimization of these two elements will benefit the efficiency of the proposed system.

Design and control of a banked instruction memory is an efficient mechanism to reduce the leakage power [11]. Most of the traditional approaches assume that the configuration of the memory in banks is performed statically. However, a few approaches also allow a dynamic selection [6], increasing the complexity of the control logic.

Historically, the Functional Units have been monolithic elements with a static behavior. Recently, several units with a dynamic changing function have appeared in the literature. For example, reconfigurable functional units [8] are functional units based on look-up tables capable of modifying the operation upon selection of the control signals. The morphable functional units, already existing in the industry [13], and the mutable functional units [18] are elements that implement several functionalities with a slight increment in the logic area. Similarly, the variable-latency functional units present an improved performance with a limited overhead [2]. However, most of these works only evaluate the area/performance trade-off, but do not analyze the impact on power consumption that the multiple alternatives can have.

**Power optimization in the radio interface.** In WSNs, the radio represents a considerable portion of the power consumption in the node. Two different approaches can be considered in order to optimize this fact: reducing the volume of information to be transmitted, or reducing the overhead of the communication protocol.

Nowadays, a novel technique for information compression has emerged: the compressed sensing. This technique reduces drastically the sampling frequency under the Shannon limit. The algorithm is based in two premises: the information dispersion exhibited by some signals (what allows the compression), and the availability of a dictionary to represent the signal with a very limited number of samples. This technique can be efficiently used in signals that own the mentioned characteristics, like images, voice, bio-medical signals, etc.

The optimization of the MAC layer in IEEE 802.15.4 (the "de facto" standard adopted for wireless communication in low-cost sensor networks) is becoming very popular nowadays in the research community. Different energy minimization techniques have been proposed, where the radio link is turned into a low-power state when the data transmission or reception is not required [17]. Also, other approaches adapt the duty cycle to the network needs [15], or reduce the idlelistening cycle required by the MAC layer [12].

**Design automation of applications for the processing node.** The development and optimization of applications for signal processing in the processing nodes is a time-consuming task, with a huge manual effort. An optimal implementation of the algorithm must adapt the application to the platform in order to exploit the available resources, reduce the computing overhead and memory use, and hence the energy consumption. The use of a high-level model of the system can help on performing a quick and energy-efficient mapping of the application specified in a design language to the node.

There are several approaches described in the literature that propose models of WSNs, but most of them only characterize in terms of energy the microprocessor, the memory, or the radio interface. Other approaches consider an holistic view of the system [3], but lack of an analytical model and topologies similar to the ones found in WBSNs.

Finally, the model created for the network can be used in a design framework to optimize the development of the applications in an efficient manner. One of this frameworks is described in [4]. However, the abstraction layer considered in this work to model the processing nodes is too high to enable any optimization mechanism in terms of adaptation of the application to the architecture for energy-saving purposes. Another perspective is shown in [10], where the user can specify the different tasks executed by the application. Unfortunately, this approach does not consider the various network topologies and node set-ups that drive the power consumption.

**Power optimization in data centers.** In modern facilities, the infrastructure energy consumption represents nearly 30% of the total cost. The largest part of the electricity bill belongs to the cooling cost, while 10-15% is due to transformation losses.

The research on this area has been focused in two different lines. One of them targets the energy savings in computing, while the other looks for savings in the cooling mechanism. Although these two lines could be analyzed together, the state-of-the-art considers them as different realities of a single problem. Only works like [1] deal with the joint problem.

Most of the solutions proposed in the lowest abstraction levels have focused on *hot-spot* reduction. These approaches reduce the effort of the cooling techniques and alleviate the reliability issues. We can find the works related to thermal-aware floorplanning [7], and the task selection and task assignment algorithms [20].

At the operating system level, there are also some works that look for balancing the workload and delay the execution of hot tasks [16]. However, it is at the server level where the most popular techniques for low-overhead thermal optimization techniques can be found. These solutions use DVFS techniques [9], Vary-On Vary-Off and all those that take advantage of unloaded machines [22]. In terms of cooling cost, the optimization techniques also exploit the thermal heterogeneity of the data center. If heavier workloads are assigned to those spots easier to cool down, the task scheduling will be more efficient in terms of energy for cooling [14]. A similar approach can consider the load balancing between hot to cold areas in the room [19].

### 3 Holistic Optimization Approach

The proposed approach considers all the agents implied in the computing paradigm mentioned in Section 1 to build a strategy for energy/cost minimization.

Design of embedded processors for signal processing. This strategy analyzes and implements several architectures of reconfigurable instruction memories. In order to optimize the power consumption in the instruction memory for a given compiler and application, it is necessary to explore the different implementations of the memory. Starting with a distributed implementation, this phase introduces several architectural modifications considering the application mapping, the execution profile, and the compiler optimizations. The results obtained show how the proper selection of, for example, the instruction memory architecture can impact dramatically the energy consumption (Figure 1 shows the total energy consumption for different implementations of the instruction memory: SPM - scratch pad memory, SPM; central loop buffer, CELB; clustered loop buffer, CLLB).



Fig. 1. Normalized energy consumption

Also, this phase proposes the dynamic reconfiguration of complex hardware arithmetic operators. The design of functional units with a tunable architecture for both complex and simple operations, allow the fine fitting of the energy/performance curve.

**Power optimization in the radio interface** This strategy proposes and implements a framework for signal analysis to develop the compressed sensing technique for several bio-signals. Also, during this phase, an analytical model of the IEEE 802.15.4 MAC layer is developed. In this way, the impact of tuning several parameters in the MAC implementation can be understood.

At this stage, a set of case studies, for which the monitoring of bio-signals and the power optimization with Quality of Service constraints have to be satisfied, is provided.

**Design automation of applications for the processing node.** This strategy provides a generic high-level model of the architecture, where the integrating components are described and the relations among them are also captured. This model facilitates the creation of a concrete architectural instance, as well as the analysis of impact of design parameters in power consumption. The analysis shows results like those presented in Table 1 where the impact in power consumption of tuning several application-level parameters have been evaluated.

Method	Parameter	$P_{onset}$	$P_{peak}$	$P_{end}$	$QRS_{onset}$	$QRS_{end}$	$T_{peak}$	$T_{end}$
2-lead 16-bit int Cubic spline		94.02 95.15 $4.1 \pm 17.1$	94.02 95.42 12.8 $\pm$ 12.1	94.05 95.49 $-2.1 \pm 14.2$	99.67 99.23 $3.2 \pm 8.7$	99.67 99.20 7 3 $\pm$ 11 9	98.00 99.23 $1.9 \pm 17.4$	97.94 98.72 $-4.1 \pm 25.3$
2-lead 16-bit int Morphological	Se (%) P <sup>+</sup> <sub>min</sub> (%) m±s (ms)	96.24 91.19 $2.5 \pm 16.4$	96.24 91.38 14.7 ± 13.4	96.27 91.55 $-1.9 \pm 15.5$	99.75 97.07 $6.9 \pm 8.0$	$99.75 97.07 9.0 \pm 9.1$	98.33 98.58 $5.3 \pm 19.3$	$98.11 98.08 -11.4 \pm 24.9$
Single-lead 16-bit int Lead 1	Se (%) P <sup>+</sup> (%) m±s (ms)	$98.84 \\92.66 \\13.4 \pm 14.8$	$98.84 \\ 92.93 \\ 15.6 \pm 12.1$	$98.87 \\ 93.24 \\ 1.6 \pm 13.2$	$99.61 \\ 99.56 \\ 5.4 \pm 8.4$	$99.61 \\ 99.56 \\ 1.5 \pm 10.1$	$99.35 \\ 99.49 \\ 5.3 \pm 18.4$	$99.32 \\ 99.24 \\ -5.3 \pm 22.7$
Single-lead 16-bit int Lead 2	Se (%) P <sup>+</sup> <sub>min</sub> (%) m±s (ms)	$97.31 \\ 91.76 \\ 10.4 \pm 19.4$	$97.09 \\ 92.21 \\ 6.9 \pm 16.9$	$96.81 \\ 91.43 \\ -8.5 \pm 17.9$	$99.67 \\98.61 \\8.6 \pm 12.6$	$99.72 \\98.72 \\8.7 \pm 13.1$	$99.35 \\ 98.90 \\ 3.5 \pm 23.5$	$99.18 \\98.38 \\-4.6 \pm 27.2$
Tolerances (2s	CSE)	10.2	-	12.7	6.5	11.6	-	30.6

Table 1. Power analysis of algorithmic parameters (ECG delineation application)

Also, this phase develops a framework for the automatic design and optimization of applications. This framework is composed of a set of tools that start with a high-level description and help on executing tasks like the application mapping to the network nodes, or the optimization for different target architectures.

**Energy optimization in data centers.** This strategy analyzes and implements several resource managing techniques at different abstraction levels, exploiting the heterogeneity of applications and computing resources for energy minimization. For this purpose, a characterization of such computing facilities and the workload is needed in order to facilitate the efficient assignment.

As an example of the potential savings, consider the energy consumed by individual tests in Figure 2. The generated workload is composed of 12 different tasks from the SPEC CPUINT 2006 benchmark. This benchmark has been executed in three different processors. At a first glance, we can see that there is margin for improvement. For example, even though the Intel Xeon server should be better than the others, there are some tasks in which the Sparc server outperforms the Intel. On the other hand, the Sparc server behaves very bad with some specific tasks. This experiment demonstrates that a proper usage of heterogeneity and an efficient optimization algorithm could lead to significant energy savings in current data centers.



Fig. 2. Variation on energy consumption of SPEC CPUINT 2006 depending on the target processor

At this stage, the analysis of the cooling mechanism is performed, and the development of control techniques for the dynamic tuning of the room temperature. This task allows to anticipate the effect of the tasks in execution. A set-up of distributed sensors is deployed to capture the metrics: temperature, humidity, air flux and direction.

#### References

- 1. Banerjee, A., Mukherjee, T., Varsamopoulos, G., Gupta, S.: Integrating cooling awareness with thermal aware workload placement for hpc data centers. Sustainable Computing: Informatics and Systems 1(2), 134–150 (2011)
- Barrio, A.A.D., et al.: Applying speculation techniques to implement functional units. In: ICCD. pp. 74–80 (2008)
- Bonivento, A., Carloni, L.P., Sangiovanni-Vincentelli, A.: Platform based design for wireless sensor networks. Mob. Netw. Appl. 11(4), 469–485 (Aug 2006), http://dx.doi.org/10.1007/s11036-006-7194-1

- Faza, A.Z., Sedigh-Ali, S.: A general purpose framework for wireless sensor network applications. In: Proc. Annual Intl. Computer Software and Applications Conference. pp. 356–358. COMPSAC (2006)
- Feng, J., Koushanfar, F., Potkonjak, M.: System-architectures for sensor networks issues, alternatives, and directions. In: In Proceedings of the 20th International Conference on Computer Design (2002)
- Gordon-Ross, A., Vahid, F., Dutt, N.D.: Fast configurable-cache tuning with a unified second-level cache. IEEE Trans. VLSI Syst. 17(1), 80–91 (2009)
- Han, Y., Koren, I.: Simulated annealing based temperature aware floorplanning. J. Low Power Electronics 3(2), 141–155 (2007)
- Hauck, S., Fry, T.W., Hosler, M.M., Kao, J.P.: The chimaera reconfigurable functional unit. In: FCCM. pp. 87–93 (1997)
- Hsu, C.h., Feng, W.c.: A power-aware run-time system for high-performance computing. In: Proc. ACM/IEEE Conference on Supercomputing. pp. 1–. SC (2005)
- Iyengar, S., et al.: A framework for creating healthcare monitoring applications using wireless body sensor networks. In: Proc. ICST Intl. Conf. on Body Area Networks. pp. 8:1–8:2. BodyNets (2008)
- Kandemir, M.T., Kolcu, I., Kadayif, I.: Influence of loop optimizations on energy consumption of multi-bank memory systems. In: Proc. of the 11th Intl. Conference on Compiler Construction. pp. 276–292. CC '02, Springer-Verlag (2002)
- Kim, T.H., Choi, S.: Priority-based delay mitigation for event-monitoring ieee 802.15.4 lr-wpans. IEEE Communications Letters pp. 213–215 (2006)
- 13. Motorola: Morphable functional units. http://ip.com/IPCOM/000004783
- Mukherjee, T., Banerjee, A., Varsamopoulos, G., Gupta, S.K.S., Rungta, S.: Spatio-temporal thermal-aware job scheduling to minimize energy consumption in virtualized heterogeneous data centers. Comput. Netw. 53(17), 2888–2904
- Neugebauer, M., Plonnigs, J., Kabitzsch, K.: A new beacon order adaptation algorithm for ieee 802.15. 4 networks. In: Proc. European Work. on Wirel. Sens. Netw. (EWSN 2005). pp. 302–311 (2005), adaptation
- 16. Pakbaznia, E., Ghasemazar, M., Pedram, M.: Temperature-aware dynamic resource provisioning in a power-optimized datacenter. In: DATE. pp. 124–129 (2010)
- Ramachandran, I., Das, A.K., Roy, S.: Analysis of the contention access period of ieee 802.15.4 mac. TOSN 3(1), 4 (2007)
- Solihin, Y., et al.: Mutable functional units and their applications on microprocessors. ICCD 0, 0234 (2001)
- Wu, G., Xu, Z., Xia, Q., Ren, J., Xia, F.: Task allocation and migration algorithm for temperature-constrained real-time multi-core systems. In: Proc. IEEE/ACM Int'l Conference on Green Computing and Communications. pp. 189–196. GREEN-COM (2010)
- Xie, Y., Hung, W.L.: Temperature-aware task allocation and scheduling for embedded multiprocessor systems-on-chip (mpsoc) design. J. VLSI Signal Process. Syst. 45(3), 177–189 (Dec 2006)
- Yassin, Y., et al.: Ultra low power application specific instruction-set processor design for a cardiac beat detector algorithm. In: NORCHIP, 2009. pp. 1 –4 (nov 2009)
- Zheng, X., Cai, Y.: Markov model based power management in server clusters. In: Proc. IEEE/ACM Int'l Conference on Green Computing and Communications. pp. 96–102. GREENCOM (2010)