Recent Advances in Power Aware Design

Tiantian Xie, Bogdan Wilamowski
Department of Electrical and Computer Engineering, Auburn University, Auburn, USA
Email: {tzx0004, wilambm}@auburn.edu

***Abstract*—The paper presents the state-of-the-art study on the recently published literatures subjected to power aware design in various industrial applications. The basic conception of power-aware design is placed at first and then the recent progress is investigated. For each paper, a brief summary is given to introduce the related power awareness design technologies and applications.**

***Index Terms*— power awareness, embedded system, energy management**

I. Introduction

With modern electronic products becoming faster, smaller, and cheaper in the past half century, the power management emerges as a focal point in many commercial systems. In high performance systems, power-aware design techniques aim to maximize performance under power dissipation and power consumption constraints such as in portable equipment. In addition, a power-aware design can also help to reduce the energy cost.

P. Ranganathan, a distinguished technologist in the HP lab, wrote an article about the concept, current research point, and future work in Power aware computing [4] recently. It is obvious that there is a minimum amount of energy needed to perform certain tasks and a corresponding minimum amount of heat extracted to avoid thermal problems. However, designs often introduce additional inefficiencies above the actual energy. These inefficiencies are often introduced when the system design must reconcile complex trade-offs that are difficult to avoid. They are often designed for the most aggressive work load performance and worst-case risk tolerance. Such designs can lead to resource over provisioning to better handle transient peaks and offer redundancy in the case of failure. Further, sometimes the traditional design focused on the small performance improvement by the cost of power and heat extraction. How to reduce these inefficiencies is called power aware strategies. Power aware strategies can be activated either in hardware or in software. To use the more power efficient component in the design is the most direct and efficient way. Another efficient method is to create energy proportionality by scaling down energy for unused resources, which requires algorithms respond to the consequences of turning down a system. Other than having the resources adapt when not fully utilized for a given task, one can also match tasks to the resources most appropriate to the size of the task. In cluster or multi-core environments, the designs might need to explicitly introduce multiple operation modes with different power-performance trade-offs. Overlap energy events seek to combine multiple tasks into a single energy event. Decomposing system functionality into smaller subtasks can help increase the benefits from an overlapping energy event by avoiding duplication of energy consumption for similar subtasks across different larger tasks. The other less common methods include broadening the scope of the solution space, trading off other metric for energy, and spending power to save power.

In this paper, the basic conception of power aware design is placed and then the recent published literature about the power aware computing in various system architectures such as computing clusters, battery power systems, hand hold devices, distributed real time systems as well as wireless sensor network systems [1-3] will be investigated.

II. Recent Progress

In this section recent developments in power aware systems will be presented which cover energy, power and thermal management approaches in various system architectures such as computing clusters, battery power systems, hand-hold devices as well as distributed real-time systems.

A. Computing Cluster

Clusters of commodity class PCs are widely used. To reduce power, cluster power management mechanisms have been proposed in both homogenous and heterogeneous systems. There are more challenges in a heterogeneous system. For a given workload of a heterogeneous cluster, the front-end PM needs to decide how many and which back-end servers should be turned on and how much workload should be distributed to each server. Since the request rate of cluster changes from time to time, these decisions have to be reevaluated and modified regularly.

Heath *et al.* [5] argued that designing efficient servers for heterogeneous clusters requires defining an efficiency metric, modeling the different types of nodes with respect to the metric, and searching for request distributions that optimize the metric. This approach is time consuming and is not designed for real-time clusters.

In paper [6], the author presents a threshold based method for efficient power management (PM) of heterogeneous soft real time clusters. The mechanism of the algorithm is built on a sophisticated offline analysis and provides an efficient threshold-based online strategy. The offline analysis divided

the maximum request rate of workload range into small ranges using many thresholds. The power is controlled and the workload distribution decision is made for each range offline. The system just needs to measure the actual request rate of the cluster and decide which small range it falls into online, and then the corresponding PM is decided. In this approach, a PM algorithm makes three design decisions on an ordered server list, server activation thresholds and workload distribution. To decide the server list, the servers are ordered by their power consumption efficiency. The PM mechanism usually turns on a server when needed or when it leads to reduced power consumption to make sure an active server usually works under a high workload. To decide the server activation thresholds, the author combines capacity and optimal-power thresholds. At last the algorithm proposes a strategy to optimally distribute the workload among active servers. The result shows that it incurs low overhead and leads to optimal power consumption.

B. Battery Power System

Many modern consumer electronics use batteries as a power supply[7-8]. The batteries have the finite capacity which determines the time during the device can be used. Several batteries are usually provided in sequential order to extend the usability of such device, but it is not the efficient way. People found that using multiple batteries in a non-sequential way can enable the use of available energy more effectively. This is decided by two inherent properties of the battery: rate-capacity effect and recovery effect. The rate-capacity effect is that a battery delivers in total a smaller charge during its lifetime when it is discharged with a high current compared to a lower one. The recovery effect is that a battery can recover and procure some additional charge when not used or only with a low current. Recent research on power management in a battery is found in the literature.

Paper [9] proposed a stochastic model focused on the recovery effect when a pulsed discharge is applied. The recovery effect is modeled as a decreasing exponential function of the state of charge and discharged capacity. The complexity of the used model limits the battery capacity to very small batteries. The three schedulers are the round-robin, best-of-two scheduler, and a random scheduler. The results show that the best-of-two scheduler is better than the other two.

Benini *et al.* consider sequential scheduling, round-robin scheduling, and various types of best-of-two scheduling, where either the output voltage or the time that a battery has been unused determines which battery is to be scheduled [10]. The different scheduling schemes are applied to several battery configurations containing up to four batteries. The loads that have been used are simple continuous and intermitted loads and two real-life example load profiles.

Wu *et al.* described the design principles and architecture of a dual-battery power supply system [11]. They integrate two battery types with different energy capacities and current rate curves into the power supply system. The first battery has a high capacity and performs well at low discharge currents. The second one has a lower capacity but performs better at high discharge currents. This new design combines different batteries to supply power for an electronic device and utilizes different advantages of them to extend the service life of the power supply for 30%.

Paper [12] provided two different approaches to maximize system lifetime by battery scheduling. The first approach models a discretized version of the KiBaM using linearly priced-timed automata (LPTA). The scheduling decision is made at predefined points in time. With model checking techniques the best possible battery schedule is found up to an error due to the discretization. The second approach showed that the actual schedule is not important when one can change between batteries and arbitrary points in time. They proposed a greedy scheduler algorithm that only switches batteries when the one used is perceived as empty. The results show that both approaches lead to significant longer lifetimes compared to simply using the batteries sequentially. With the greedy algorithm it is possible to do computations for larger and more batteries with is not possible with the first approach. However, the first one allows for extra limitations on the time of scheduling, and the model checking techniques ensure that the best possible schedule will be found.

C. Hand-Hold Device

Modern mobile devices require high performance at system level to decode high-bitrates multimedia. The processor offloading using off-chip controller, which performs graphic processing, memory control and multimedia acceleration, is commonly exercised in this field. Due to the finer process and low-power technology, the form factor of the off-chip controller was reduced to the size of those in ultra-mobile personal computers and mobile Internet devices. Processor DVS, which is a technique for changing the processor frequency and voltage based on the workload and processor DPM which forces the processor into low-power state, is two primary methods in the processor power management. Several researches have done based on these concepts. Because SoC based processors provide DVS functionality, several policies in software such as PowerNowd and On demand were implemented as part of operating system-directed power management. These policies exploit previous load history or utilization to determine the performance of the next period.

Paper [13] considered off-chip usages while using a processor with both DVS and DPM simultaneously. They devise a system variable that identifies the off-chip dominant period and a variable to distinguish the memory intensive period. Based on the variables found, an algorithm is proposed to achieve additional power reduction by combining C-state and DVS. The first step is to find off-chip intensiveness. And then the memory intensiveness is identified by measuring the read data cache based on the rationale that applications with multimedia data induce frequent cache misses due to the nonlocality of the data streaming. The result shows that the proposed technique in a real hardware environment achieved up to 37% power reduction. Also, this technique can be easily applied to consumer products such as Smartphones or netbooks

which expand their architectures toward off-chip based architectures.

D. Distributed real time system

The dynamic voltage and frequency scaling (DVFS) technique is the most common energy management scheme for embedded systems. However, it may have a negative impact on task and system reliability. The author in [14] aimed to maximize reliability in energy-constrained settings. They took account in both frame-based and periodic task models and showed how to compute frequency assignments to maximize the overall reliability while satisfying a hard energy constraint. They presented the static optimal scheme at first and then extended with online schemes to improve reliability.

The CPU utilization troll has recently been demonstrated to be an effective way of meeting end-to-end deadlines for distributed real-time systems running in unpredictable environments. The current utilization control cannot effectively handle rate saturation and discrete task rates. Wang *et al.* [15] presented a two-layer coordinated CPU utilization architecture with theoretically guaranteed control accuracy and system stability. The primary control loop used frequency scaling to locally control the CPU utilization of each processor; while the secondary control loop adopted rate adaption to control the utilizations of all the processors at the cluster level on a finer timescale. The experimental results showed that the proposed controller achieved higher accuracy and less power consumption than a state-of-the-art utilization control algorithm. Also, the feasibility of utilization control was significantly improved using the proposed control solution.

T. Cucinotta *et al* [16]proposed a novel Quality-of-Service (QoS) management framework for soft real-time applications based on two control loops: the resource allocation controller and Global QoS Controller. The first one operates on the scheduling parameter to obtain a resource allocation that meets the temporal constraints of the application. The second one operates on the QoS level of the applications and on the power level of the resources to strike a good tradeoff between the global QoS and the energy consumption. Such design allows for the optimization of the performance of multimode real-time applications running on hardware with power-switching capabilities.

Energy-efficient real-time task synchronization protocols and the overhead of frequency switching are often considered in the form of optimization problems in real systems design. Chen et al. [17] proposed the concept of frequency locking in order to better manage the cost in frequency switching. Several algorithms were designed to minimize the energy consumption and meet the timing constrains, based on the extended synchronization protocols using the frequency locking concept. It shows good simulation and experimental results on a real case study evaluating the proposed methodology by comparing with different workloads and protocols.

E. Other Progress

Besides those researches listed above, there are also many works related to the thermal management problem. The high chip temperature not only increases package and cooling cost but also reduces the reliability and performance of the systems. The leakage plays a critical role in thermal aware techniques in the deep submicron domain. There is a strong relationship among the leakage, temperature, and supply voltage. In paper [18], the author developed a novel power model that can capture the interaction between the temperature and leakage with high accuracy. Based on this model, they developed three conditions to check the feasibility of real time tasks under the peak temperature requirement. Their results show that the leakage current model has a relative error less than 5% based on the processor using 65 nm technologies. At the same time, they highlighted that the independency of leakage, temperature, and supply voltage should be properly addressed so that the thermal aware computing becomes efficient.

Shin introduced a new dynamic thermal management (DTM) framework in Silicon devices [19]. They proposed that the die temperature can strongly affect its leakage power consumption and reliability. This method considered the thermal resistance of the heat-sink as a control variable and minimized the total power consumption both for computation and cooling. They controlled the cooling power consumption together with the microprocessor clock frequency and supply voltage and tracked the energy optimal die temperature. They spent a little bit higher cooling power than conventional DTM. However, they reduce a significant amount of the temperature-dependent leakage power consumption of the microprocessor. The total consumption of power was reduced by 8.2% compared with a baseline DTM approach. The proposed approach also enhances the Failures in Time (FIT) up to 80% in terms of the electro-migration lifetime reliability.

Rowe *et al.* [20] focused on the missed energy saving due to idle gaps between executing tasks that are shorter than the time required to enter the sleep mode. They used a family of rate-harmonized schedulers (RHS) that clusters the execution of tasks so that idle durations can be lumped together enabling transitions to the sleep mode. By adjusting the phasing between periodic tasks and removing the idle slots between periodic tasks and removing the idle slots between active execution that are too short for the processor to make a round-trip transition from idle into deep-sleep, RHS increases the percentage of processor sleep time. This approach not only improved energy-efficiency but also simplifies the design. It can also support the multicore processors. The novel RHS saves 16.8% energy compared to conventional Rate-Monotonic Scheduling.

Multiprocessor system-on-chip (SOC) platforms have been adopted for a wide range of high-performance applications. Task assignment and processing unit allocation are key steps in the design of predictable and efficient embedded systems. Schranzhofer *et al.* [21] proposed a methodology to calculate a task to processing element mapping so as to minimize the expected average power consumption in the given execution modes of applications. The proposed method allowed the

system to adapt to changing environmental conditions at runtime. Both leakage and dynamic power consumption were considered based on the underlying model. Also, a manager monitored the system at runtime and chose a proper precomputed template so as to maintain low power-consumption over the systems lifetime. Experimental results revealed the effectiveness of the proposed algorithm by comparing the derived solutions to the optimal solutions which were obtained by an integer linear program (ILP).

G. Anastasi *et al* have done some research focused on the power management in wireless sensor networks [22]. They proposed an Adaptive Staggered Sleep Protocol targeted to periodic data acquisition. This protocol dynamically adjusted the sleep schedules of nodes based on the traffic pattern and the operating conditions experienced by that node to match the network demands. The protocol is adapted to variations in the message generation rate, network topology, external conditions and so on. As the active periods are tailored to the actual needs of each single node, the proposed protocol tends to minimize both energy consumption and message latency. Also, it is conceived as an independent protocol operating above the MAC layer and can be used with different platforms.

Methods of computational intelligence, such as neural networks [23-24], neuro-fuzzy systems [25], and radial basis function networks [26], provide various strategies to design power-aware systems. A neuro-fuzzy controller proposed by Kim *et al.* [27] performed intelligent real-time estimation by mimicking the human visual system. The processor's overall pipeline states were managed using workload-aware task scheduling and applied database size control. Experimental results showed that the neuro-fuzzy controller performed workload-aware dynamic power management to reduce the proposed processor's power consumption.

III. CONCLUSION

Power-aware research is a new and vibrant area, and it is useful in the real world. Many new optimization methods and algorithms in power management have been proposed to reduce the energy depletion and improve the performance of the systems. Also, more and more novel applications are opened up and they significantly impact the continuous development of power-aware technologies.

REFERENCES

[1] J.-J. Chen, X. S. Hu, Mossé, D. , L. Thiele, "Guest Editorial Special Section on Power-Aware Computing ," in *IEEE Trans. on Industrial Informatics*, vol. 6, no. 3, Aug 2010, pp: 253-254.

[2] C.A. Boano, N. Tsiftes, T. Voigt, J. Brown, U. Roedig, "The Impact of Temperature on Outdoor Industrial Sensor net Applications ," in *IEEE Trans. on Industrial Informatics*, vol. 6, no. 3, Aug 2010, pp:451-459.

[3] A. Malinowski and H. Yu, "Comparison of Embedded System Design for Industrial Applications," in *IEEE Trans. on Industrial Informatics,* vol. 7, no. 2, pp. 244 -254 , May 2011.

[4] P. Ranganathan, "Recipe for Efficiency: Principles of Power-Aware Computing", in *Communications of The ACM*, vol. 53, no 4, 2010,, pp:60-67.

[5] T. Heath, B. Diniz, E. V. Carrera, W. Meira, Jr., and R. Bianchini, "Energy conservation in heterogeneous server clusters," in *Proc. 10th ACM SIGPLAN Symp. Principles and Practice of Parallel Programming (PPoPP 05)*, New York, 2005, pp. 186–195, ACM Press.

[6] L. Wang and Y.Lu, "An Efficient Threshold-based Power Management Mechanism for Heterogeneous Soft Real-Time Clusters". In *IEEE Trans. Industrial Informatics*, vol. 6, no. 3, 2010, pp352-364.

[7] M.B. Camara, H. Gualous, F. Gustin, A. Berthon, B. Dakyo, "DC/DC Converter Design for Supercapacitor and Battery Power Management in Hybrid Vehicle Applications—Polynomial Control Strategy ,"in *IEEE Trans. on Industrial Electronics*, vol. 57, no. 2, 2010, pp:587-597.

[8] P. Garcia, L.M. Fernandez, C.A. Garcia, F. Jurado, "Energy Management System of Fuel-Cell-Battery Hybrid Tramway ,"in *IEEE Trans. on Industrial Electronics*, vol. 57, no. 12, 2010, pp:4013-4023.

[9] C. F. Chiasserini and R. R. Rao, "Energy efficient battery management," in *IEEE J. Sel. Areas Commun.*, vol. 19, no. 7, 2001, pp. 1235–1245.

[10] L. Benini, G. Castelli, A. Macii, E. Macii, M. Poncino, and R. Scarsi, "Extending lifetime of portable systems by battery scheduling," in *Design, Automation and Test in Europe (DATE)*, Los Alamitos, CA, 2001, pp. 197–203.

[11] Q. Wu, Q. Qiu, and M. Pedram, "An interleaved dual-battery power supply for battery-operated electronics," in *Proc. Conf. Asia South Pacific Design Autom. (ASP-DAC'00)*, 2000, pp. 387–390.

[12] M. Jongerden, A. Mereacre, H. Bohnenkamp, B. Haverkort and J. Katoen, "Computing Optimal Schedules of Battery Usage in Embedded Systems", in *IEEE Trans. Industrial Informatics.* vol. 6, no. 3, 2010, pp 276-286.

[13] Jinuk Choi, Hojung Cha, "A Processor Power Management Scheme for Handheld Systems Considering Off-Chip Contributions", in *IEEE Trans. Industrial Informatics*. vol. 6, no. 3, 2010, pp255-264.

[14] Baoxian Zhao, H. Aydin, Dakai Zhu, "On Maximizing Reliability of Real-Time Embedded Applications Under Hard Energy Constraint", in *IEEE Trans. Industrial Informatics*. vol. 6, no. 3, 2010, pp316-327.

[15] Xiaorui Wang, Xing Fu, Xue Liu, Zonghua Gu, "PAUC: Power-Aware Utilization Control in Distributed Real-Time Systems", in *IEEE Trans. Industrial Informatics*. vol. 6, no. 3, 2010, pp302-315.

[16] T. Cucinotta, L. Palopoli, L. Abeni, D. Faggioli, G. Lipari, "On the Integration of Application Level and Resource Level QoS Control for Real-Time Applications ," in *IEEE Trans. on Industrial Informatics*, vol. 6, no. 4, 2010, pp479-491.

[17] Ya-Shu Chen, Chuan-Yue Yang, Tei-Wei Kuo, "Energy-Efficient Task Synchronization for Real-Time Systems", in *IEEE Trans. Industrial Informatics*. vol. 6, no. 3, 2010, pp287-301.

[18] Gang Quan, V. Chaturvedi, "Feasibility Analysis for Temperature-Constraint Hard Real-Time Periodic Tasks", in *IEEE Trans. Industrial Informatics*. vol. 6, no. 3, 2010, pp329-339.

[19] Donghwa Shin, Sung Woo Chung, Eui-Young Chung, Naehyuck Chang, "Energy-Optimal Dynamic Thermal Management: Computation and Coolong Power Co-Optimization", in *IEEE Trans. Industrial Informatics*. vol. 6, no. 3, 2010, pp340-351.

[20] A. Rowe, K. Lakshmanan, Haifeng Zhu, R. Rajkumar, "Rared-Harmonized Scheduling and Its Applicability to Energy management", in *IEEE Trans. Industrial Informatics.* vol. 6, no. 3, 2010, pp265-275.

[21] A. Schranzhofer, Jian-Jian Chen, L. Thiele, "Dynamic Power-Aware Mapping of Applications onto Heterogeneous MPSoC Platforms ," in *IEEE Trans. on Industrial Informatics,* vol. 6, no. 4, pp. 692-707, Nov 2010.

[22] G. Anastasi, M. Conti, M. Di Francesco, "Extending the Lifetime of Wireless Sensor Networks Through Adaptive Sleep," in *IEEE Trans. on Industrial Informatics*, vol. 5, no. 3, 2009, pp 351-365.

[23] B. M. Wilamowski and H. Yu, "Improved Computation for Levenberg-Marquardt Training," in *IEEE Trans. on Neural Networks*, vol. 21, no. 6, pp. 930-937, June 2010.

[24] B. M. Wilamowski and H. Yu, "Neural Network Learning Without Backpropagation," in *IEEE Trans. on Neural Networks*, vol. 21, no. 11, pp. 1793-1803, Nov. 2010.

[25] T. T. Xie, H. Yu, and B. M. Wilamowski, "Replacing fuzzy systems with neural networks," in *Proc. IEEE HSI Conf.*, HSI2010, Rzeszow, Poland, May 13-15, 2010, pp. 189-193.

[26] H. Yu, T. Xie, S. Paszczynski and B. M. Wilamowski, "Advantages of Radial Basis Function Networks for Dynamic System Design," in *IEEE Trans. on Industrial Electronics*, vol. 58, no. 12, pp. 5438-5450, Dec. 2011.

[27] J. Y. Kim, M. Kim, S. Lee, J. Oh, S. Oh and H. J. Yoo, "Real-Time Object Recognition with Neuro-Fuzzy Controller Workload Aware Task Pipelining," in *IEEE Micro*, vol. 29, no. 6, pp. 28-43, Dec. 2009.