Dynamic Voltage Scaling And Power Management For Portable

Dynamic voltage scaling (DVS) is a promising method to reduce the power consumption of CMOS-based embedded processors. However, pure DVS techniques do not perform well for dynamic systems where the execution times of different jobs vary significantly. A novel DVS scheme with feedback control mechanisms for hard real-time systems is proposed in this work. It produces energy-efficient schedules for both static and dynamic workloads. Task-splitting, slack-passing and preemption-handling schemes are proposed to aggressively reduce the speed of each task. Different feedback control structures are integrated into the DVS algorithm to make it adaptable to workload variations. This scheme relies strictly on operating system support. It is evaluated in simulation as well as on an embedded platform. For given task sets, simulation experiments demonstrate the benefits of this scheme with savings of up to 29\% in energy over previous work. This scheme exhibits up to 24\% additional energy savings over other DVS algorithms on the embedded platform. The feedback-based DVS scheme is further extended to be leakage aware, which considers not only dynamic but also static power consumption caused by leakage current in circuits. A combined DVS, delay and sleeping scheme is proposed for architectures where static power exceeds dynamic power in some cases. DVS is used when dynamic power dominates the total power consumption, while a sleep mode is entered when static power becomes dominant. The extended algorithm, DVSleak, shows 30\% additional energy savings on average over a pure DVS algorithm in the simulation experiment.

Energy consumption has become a primary concern in the last decade. One highly effective way to reduce CPU energy while still executing applications is dynamic voltage scaling (DVS). While DVS makes runtime transitions between power levels possible, thus far the scheduling of DVS has only been implemented at the system levels. The primary reason for this is that a transition has significant time and energy costs and therefore must be restricted. On the other hand, if developers are given control over DVS and the flexibility to apply it as necessary, then DVS scheduling decisions can include application-specific knowledge. We have developed a runtime support module for developer-driven dynamic voltage scaling (D3VS). The module allows applications to be densely populated with DVS scale requests, yet restricts the DVS overhead to 4% under reasonable assumptions. To do this, the module does not make power level transitions at every request. Instead, using the past history as hints, it picks a single power level that is representative of the application's behavior. In this thesis we present the analytical models and simulations used in the design of the D3VS runtime support module.

Multiprocessor platforms play important roles in modern computing systems, and appear in various applications, ranging from energy-limited hand-held devices to large data centers. As the performance requirements increase, energy-consumption in these systems also increases significantly. Dynamic Voltage and Frequency Scaling (DVFS), which allows processors to dynamically adjust the supply voltage and the clock frequency to operate on different power/energy levels, is considered an effective way to achieve the goal of energy-saving. This book surveys existing works that have been on energy-aware task scheduling on DVFS multiprocessor platforms. Energy-aware scheduling problems are intrinsically optimization problems, the formulations of which greatly depend on the platform and task models under consideration. Thus, Energy-aware Scheduling on Multiprocessor Platforms covers current research on this topic and classifies existing works according to two key standards, namely, homogeneity/heterogeneity of multiprocessor platforms and the task types considered. Under this classification, other sub-issues are also included, such as, slack reclamation, fixed/dynamic priority scheduling, partition-based/global scheduling, and application-specific power consumption, etc.

The Dynamic Voltage Scaling (DVS) technique has proven to be ideal in regard to balancing performance and energy consumption of a processor since it allows for almost cubic reduction in dynamic power consumption with only a nearly linear reduction in performance. Due to its virtue, the DVS technique has been used for the two main purposes: energy-saving and temperature reduction. However, recently, a Dynamic Voltage Scaled (DVS) processor has lost its appeal as process technology advances due to the increasing Process, Voltage and Temperature (PVT) variations. In order to make a processor tolerant to the increasing uncertainties caused by such variations, processor designers have used more timing margins. Therefore, in a modern-day DVS processor, reducing voltage requires comparatively more performance degradation when compared to its predecessors. For this reason, this technique has a lot of room for improvement for the following facts. (a) From an energy-saving viewpoint, excessive margins to account for the worst-case operating conditions in a DVS processor can be exploited because they are rarely used during run-time. (b) From a temperature reduction point of view, accurate prediction of the optimal performance point in a DVS processor can increase its performance. In this dissertation, we propose four performance improvement ideas from two different uses of the DVS technique. In regard to the DVS technique for energy-saving, in this dissertation, we introduce three different types of margin reduction (or margin decision) techniques. First, we introduce a new indirect Critical Path Monitor (CPM) to make a conventional DVS processor adaptive to its given environment. Our CPM is composed of several Slope Generators, each of which generates similar voltage scaling slopes to those of potential critical paths under a process corner. Each CPR in the Slope Generator tracks the delays of potential critical paths with minimum difference at any condition in a certain voltage range. The CPRs in the same Slope Generator are connected to a multiplexer and one of them is selected according to a current voltage level. Calibration steps are done by using conventional speed-binning process with clock duty-cycle modulation. Second, we propose a new direct CPM that is based on a non-speculative pre-sampling technique. A processor that is based on this technique predicts timing errors in the actual critical paths and undertakes preventive steps in order to avoid the timing errors in the event that the timing margins fall below a critical level. Unlike direct CPM that uses circuit-level speculative operation, although the shadow latch can have timing error, the main Flip-Flop (FF) of our direct CPM never fails, guaranteeing always-correct operation of the processor. Our non-speculative CPM is more suitable for high-performance processor designs than the speculative CPM in that it does not require original design modification and has lower power overhead. Third, we introduce a novel method that determines the most accurate margin that is based on the conventional binning process. By reusing the hold-scan FFs in a processor, we reduce design complexity, minimize hardware overhead and increase error detecting accuracy. Running workloads on the processor with Stop-Go clock gating allows us to find which paths have timing errors during the speed binning steps at various, fixed temperature levels. From this timing error information, we can determine the different maximum frequencies for diverse operating conditions. This method has high degree of accuracy without having a large overhead. In regard to the DVS technique for temperature reduction, we introduce a run-time temperature monitoring scheme that predicts the optimal performance point in a DVS processor with high accuracy. In order to increase the accuracy of the optimal performance point prediction, this technique monitors the thermal stress of a processor during run-time and uses several Look-Up Tables (LUTs) for different process corners. The monitoring is performed while applying Stop-Go clock gating, and the average EN value is calculated at the end of the monitoring time. Prediction of the optimal performance point is made using the average EN value and one of the LUTs that corresponds to the process corner under which the processor was manufactured. The simulation results show that we can achieve maximum processor performance while keeping the processor temperature within threshold temperature.

This lecture provides an introduction to the problem of managing the energy demand of mobile devices. Reducing energy consumption, primarily with the goal of extending the lifetime of battery-powered devices, has emerged as a fundamental challenge in mobile computing and wireless communication. The focus of this lecture is on a systems approach where software techniques exploit state-of-the-art architectural features rather than relying only upon advances in lower-power circuitry or the slow improvements in battery technology to solve the problem. Fortunately, there are many opportunities to innovate on managing energy demand at the higher levels of a mobile system.

Increasingly, device components offer low power modes that enable software to directly affect the energy consumption of the system. The challenge is to design resource management policies to effectively use these capabilities. The lecture begins by providing the necessary foundations, including basic energy terminology and widely accepted metrics, system models of how power is consumed by a device, and measurement methods and tools available for experimental evaluation. For components that offer low power modes, management policies are considered that address the questions of when to power down to a lower power state and when to power back up to a higher power state. These policies rely on detecting periods when the device is idle as well as techniques for modifying the access patterns of a workload to increase opportunities for power state transitions. For processors with frequency and voltage scaling capabilities, dynamic scheduling policies are developed that determine points during execution when those settings can be changed without harming quality of service constraints. The interactions and tradeoffs among the power management policies of multiple devices are discussed. We explore how the effective power management on one component of a system may have either a positive or negative impact on overall energy consumption or on the design of policies for another component. The important role that application-level involvement may play in energy management is described, with several examples of cross-layer cooperation. Application program interfaces (APIs) that provide information flow across the application-OS boundary are valuable tools in encouraging development of energy-aware applications. Finally, we summarize the key lessons of this lecture and discuss future directions in managing energy demand.

The monolithic implementation of a Zero-Voltage-Switching Quasi-Square Wave (ZVS-QSW) buck converter capable of meeting the future challenges of low-voltage point-of-load converters is presented. The design uses novel high-speed Dead-Time-Locked-Loops with precision dead-time control to achieve ZVS under dynamic loads and a variable output voltage. By eliminating switching losses, high efficiency is achieved at switching frequencies beyond 1MHz. The ZVS-QSW converter is implemented in a 0.18mum CMOS process and has a measured efficiency of 82% at 5MHz with a 1.4V output. In order to meet the demand for scalable performance, the Dynamic Voltage scaling concept is applied to the ZVS-QSW converter using a dual-mode configuration. In DVS, the supply voltage is scaled down to reduce power while satisfying a variable target clock frequency. An experimental DVS system is demonstrated using a state-of-the-art CPLD. The fully digital DVS controller can achieve the maximum Vdd transition in 30mus.

Please note that the content of this book primarily consists of articles available from Wikipedia or other free sources online. Pages: 28. Chapters: Active State Power Management, Case modding, Cool'n'Quiet, CPU locking, Dynamic frequency scaling, Dynamic voltage scaling, HD Tune, Jumper (computing), LongHaul, Memory divider, Modchip, Overclocking, Pentium OverDrive, Performance acceleration technology, Performance tuning, PowerNow!, Quiesce, RivaTuner, RMClock, SpeedFan, SpeedStep, Tweaking, Underclocking.

This book offers the first comprehensive coverage of digital design techniques to expand the power-performance tradeoff well beyond that allowed by conventional wide voltage scaling. Compared to conventional fixed designs, the approach described in this book makes digital circuits more versatile and adaptive, allowing simultaneous optimization at both ends of the power-performance spectrum. Drop-in solutions for fully automated and low-effort design based on commercial CAD tools are discussed extensively for processors, accelerators and on-chip memories, and are applicable to prominent applications (e.g., IoT, AI, wearables, biomedical). Through the higher power-performance versatility techniques described in this book, readers are enabled to reduce the design effort through reuse of the same digital design instance, across a wide range of applications. All concepts the authors discuss are demonstrated by dedicated testchip designs and experimental results. To make the results immediately usable by the reader, all the scripts necessary to create automated design flows based on commercial tools are provided and explained.

State-of-the-Art Approaches to Advance the Large-Scale Green Computing Movement Edited by one of the founders and lead investigator of the Green500 list, The Green Computing Book: Tackling Energy Efficiency at Large Scale explores seminal research in large-scale green computing. It begins with low-level, hardware-based approaches and then traverses up the software stack with increasingly higher-level, software-based approaches. In the first chapter, the IBM Blue Gene team illustrates how to improve the energy efficiency of a supercomputer by an order of magnitude without any system performance loss in parallelizable applications. The next few chapters explain how to enhance the energy efficiency of a large-scale computing system via compiler-directed energy optimizations, an adaptive run-time system, and a general prediction performance framework. The book then explores the interactions between energy management and reliability and describes storage system organization that maximizes energy efficiency and reliability. It also addresses the need for coordinated power control across different layers and covers demand response policies in computing centers. The final chapter assesses the impact of servers on data center costs.

The Multi-Level Computing Architecture (MLCA) is a novel architecture for parallel systems-on-a-chip. We propose and evaluate a profile-driven compiler technique for power optimizations of MLCA applications using dynamic voltage scaling (DVS). Our technique combines dependence analysis of loops with profiling in order to identify the slack in parallel execution of coarse-grain tasks. DVS is applied to slow down processors executing tasks outside the critical path, saving power with little or no impact on execution time. Evaluation of our technique using an MLCA simulator and three realistic MLCA multimedia applications shows that up to 10% savings in processor power consumption can be achieved with no more than 1.5% increase in execution time. The achieved power savings are significantly greater than those that could be achieved by uniformly slowing down all computations with only a similar increase in overall execution time.

This thesis investigates Dynamic Voltage Scaling (DVS) techniques to lower power consumption in video decoding. A DVS scheme called the Frame-data Computation Aware (FDCA) method has been presented. This method is adaptable not only to stored video applications but also to real-time video scenarios. Unlike DVS schemes for video decoding proposed earlier, the FDCA scheme does not require any preprocessing mechanisms. Results from simulations performed using the scheme are presented and compared with prior existing DVS schemes. The results indicate that the FDCA method provides power saving of up to an average of about 68%.

In this work, we propose a new dynamic migration (DM) heuristic method integrating dynamic voltage scaling (DVS), dynamic power management (DPM) and task migration in multi-core real-time systems which can feasibly balance the task load and reduce energy consumption during execution to achieve energy efficiency. Meanwhile, voltage scaling based dynamic core scaling (VSDCS) is presented for reducing leakage power consumption under low task load

conditions. The framework used for the proposed methods is composed of a partitioner, a local earliest deadline first (EDF) scheduler, a power-aware manager, a dynamic migration module, and a dynamic core scaling module. The primary unit is the power-aware manager which controls the frequency for the power consumption and the voltage scaling based on the feedback of the dynamic migration module and the dynamic core scaling module. Simulation results show that the DM heuristic can produce further energy savings of about 3 percent compared with the closest previous work. That is (1-(1 - 8\%)x(1 - 3%))=11% energy saved with the new DM techniques. This work also greatly reduces the cost of task migration among the multi-core processors. The results show that VSDCS can achieve up to 33 percent of energy savings under low load conditions as compared with previous methods.

Circuit designers perform optimization procedures targeting speed and power during the design of a circuit. Gate sizing can be applied to optimize for speed, while Dual-VT and Dynamic Voltage Scaling (DVS) can be applied to optimize for leakage and dynamic power, respectively. Both gate sizing and Dual-VT are design-time techniques, which are applied to the circuit at a fixed voltage. On the other hand, DVS is a run-time technique and implies that the circuit will be operating at a different voltage than that used during the optimization phase at design-time. After some analysis, the risk of non-critical paths becoming critical paths at run-time is detected under these circumstances. The following questions arise: 1) should we take DVS into account during the optimization phase? 2) Does DVS impose any restrictions while performing design-time circuit optimizations? This thesis is a case study of applying DVS to a circuit that has been optimized for speed and power, and aims at answering the previous two questions. We used a 45-nm CMOS design kit and flow. Synthesis, placement and routing, and timing analysis were applied to the benchmark circuit ISCAS'85 c432. Logical Effort and Dual-VT algorithms were implemented and applied to the circuit to optimize for speed and leakage power, respectively. Optimizations were run for the circuit operating at different voltages. Finally, the impact of DVS on circuit optimization was studied based on HSPICE simulations sweeping the supply voltage for each optimization. The results showed that DVS had no impact on gate sizing optimizations, but it did on Dual-VT optimizations. It is shown that we should not optimize at an arbitrary voltage. Moreover, simulations showed that Dual-VT optimizations should be performed at the lowest voltage that DVS is intended to operate, otherwise non-critical paths will become critical paths at run-time.

This book constitutes the thoroughly refereed post-proceedings of the Second International Workshop on Power-Aware Computer Systems, PACS 2002, held in Cambridge, MA, USA, in February 2002. The 13 revised full papers presented were carefully selected for inclusion in the book during two rounds of reviewing and revision. The papers are organized in topical sections on power-aware architecture and microarchitecture, power-aware real-time systems, power modeling and monitoring, and power-aware operating systems and compilers.

In recent years, many innovative researches have been conducted on dynamic voltage scaling (DVS), such as Razor [1]. This thesis presents an error-tolerant DVS design that can enhance the reliability and reduce the power consumption of a pipeline circuit simultaneously. Based on delay distributions of all pipeline stages, an efficient voltage island partitioning method is developed to cluster all pipeline stages into several voltage islands. By assigning the best voltages to stages, the DVS design can enable the pipeline stages to work at an optimal energy consumption with least performance penalty. Experimental results obtained by HSPICE and Matlab simulations demonstrate the feasibility of this method.

Power consumption becomes the most important design goal in a wide range of electronic systems. There are two driving forces towards this trend: continuing device scaling and ever increasing demand of higher computing power. First, device scaling continues to satisfy Moore’s law via a conventional way of scaling (More Moore) and a new way of exploiting the vertical integration (More than Moore). Second, mobile and IT convergence requires more computing power on the silicon chip than ever. Cell phones are now evolving towards mobile PC. PCs and data centers are becoming commodities in house and a must in industry. Both supply enabled by device scaling and demand triggered by the convergence trend realize more computation on chip (via multi-core, integration of diverse functionalities on mobile SoCs, etc.) and finally more power consumption incurring power-related issues and constraints. Energy-Aware System Design: Algorithms and Architectures provides state-of-the-art ideas for low power design methods from circuit, architecture to software level and offers design case studies in three fast growing areas of mobile storage, biomedical and security. Important topics and features: - Describes very recent advanced issues and methods for energy-aware design at each design level from circuit and architecture to algorithm level, and also covering important blocks including low power main memory subsystem and on-chip network at architecture level - Explains efficient power conversion and delivery which is becoming important as heterogeneous power sources are adopted for digital and non-digital parts - Investigates 3D die stacking emphasizing temperature awareness for better perspective on energy efficiency - Presents three practical energy-aware design case studies; novel storage device (e.g., solid state disk), biomedical electronics (e.g., cochlear and retina implants), and wireless surveillance camera systems. Researchers and engineers in the field of hardware and software design will find this book an excellent starting point to catch up with the state-of-the-art ideas of low power design.

[Truncated abstract] Recent advances in semiconductor fabrication technology have enabled the concept of a single camera-on-a-chip, which integrates all camera functions onto a single piece of silicon. To enable the concept of low cost and high performance miniature camera-on-a-chip, the proposed research aims at developing advanced real-time digital image processing (DIP) cores or modules that could be integrated together with the photo-sensing pixel array. The developed DIP cores address the stringent requirements of compactness, low-power, and real-time operation to enable their integration into a variety of low cost portable consumer imaging products. This work first presents a multi-precision reconfigurable multiplier, which incorporates variable precision, parallel processing, razor-based dynamic voltage scaling (DVS), and dedicated multi-precision operands scheduling to realize full energy and performance flexibility and efficiency.

According to user s arbitrary requirements (eg. throughput), the dynamic voltage and frequency scaling management unit first configures the multiplier operating at the proper precision and frequency. Adapting to the run-time workload of the targeted application, razor flip-flops and the dithering voltage unit assist the voltage and frequency scaling management units to autonomously configure the multiplier to work at the minimum possible power operating point to achieve the minimized power consumption. This multi-precision multiplier is coupled with an operands scheduler, which can analyze and rearrange the input data to achieve the optimal voltage and frequency combinations to further reduce the overall power consumption. Our work successfully demonstrates through a fabricated prototype that multi-precision architecture can reap the benefits from dynamic voltage scaling techniques more effectively, which enables the efficient use for DIP applications. A new digital image sensor (DPS) architecture based on a multi-reset integration scheme was proposed to reduce memory needs, sensor size and consumption of the pixel level. The operation of the DPS exploits the chronological change of the code. In the proposed implementation, a 4-bit in-pixel memory is used to reduce the pixel size, and an 8-bit resolution is achieved with the multi-reset scheme. In addition, full complementary metal-oxide-semiconductor (CMOS) 2T DRAM and selective refresh scheme are adopted to implement the memory elements and maximize the area savings. The proposed architecture was validated by a prototype chip fabricated using AMS 0.35U+006fm CMOS technology. Our scheme achieves a 20% fill factor for a 22U+006fm {604} 22U+006fm digital pixel sensor, with a power consumption per pixel reduced by around 37.5% compared to previous implementations. Finally, we reduced image processing precision to 1 single bit. The resulting binary images and corresponding Boolean operations reduce the computational complexity of DIP cores/modules significantly compared to traditional full- precision arithmetic processing. The proposed approach exploits the typically long integration times of each image plane in order to carry out processing at low operating frequency and supply voltage...

Although users of high-performance computing are most interested in raw performance, both energy and power consumption have become critical concerns. As a result improving energy efficiency of nodes on HPC machines has become important and the importance of power-scalable clusters, where the frequency and voltage can be dynamically modified, has increased. This thesis investigates the energy consumption and execution time of applications on a power-scalable cluster. It studies intra-node and inter-node effects of memory and communication bottlenecks. Results show that a power-scalable cluster has the potential to save energy by scaling the processor down to lower energy levels. This thesis presents a model that predicts the energy-time trade-off for larger clusters. On power-scalable clusters, one opportunity for saving energy with little or no loss of performance exists when the computational load is not perfectly balanced. This situation occurs frequently, as keeping the load balanced between nodes is one of the long standing fundamental problems in parallel and distributed computing. However, despite the large body of research aimed at balancing load both statically and dynamically, this problem is quite difficult to solve. This thesis presents a system called Jitter that reduces the frequency on nodes that are assigned less computation and therefore have idle time or slack time. This saves energy on these nodes, and the goal of Jitter is to attempt to ensure that they arrive 'just in time' so that they avoid increasing overall execution time. Specifically, we dynamically determine which nodes have enough slack time so that they can be slowed down, which will greatly reduce the consumed energy on that node. Thus a superior energy-time trade-off can be achieved. This thesis studies a suite of MPI benchmarks, which are profiled, gathering information about the computation and communication occuring in the application. This information is used to analyse various ene.

With the advent of portable and autonomous computing systems, power con sumption has emerged as a focal point in many research projects, commercial systems and DoD platforms. One current research initiative, which drew much attention to this area, is the Power Aware Computing and Communications (PAC/C) program sponsored by DARPA. Many of the chapters in this book include results from work that have been supported by the PACIC program. The performance of computer systems has been tremendously improving while the size and weight of such systems has been constantly shrinking. The capacities of batteries relative to their sizes and weights has been also improv ing but at a rate which is much slower than the rate of improvement in computer performance and the rate of shrinking in computer sizes. The relation between the power consumption of a computer system and it performance and size is a complex one which is very much dependent on the specific system and the technology used to build that system. We do not need a complex argument, however, to be convinced that energy and power, which is the rate of energy consumption, are becoming critical components in computer systems in gen eral, and portable and autonomous systems, in particular. Most of the early research on power consumption in computer systems ad dressed the issue of minimizing power in a given platform, which usually translates into minimizing energy consumption, and thus, longer battery life.

Power Aware Design Methodologies was conceived as an effort to bring all aspects of power-aware design methodologies together in a single document. It covers several layers of the design hierarchy from technology, circuit logic, and architectural levels up to the system layer. It includes discussion of techniques and methodologies for improving the power efficiency of CMOS circuits (digital and analog), systems on chip, microelectronic systems, wirelessly networked systems of computational nodes and so on. In addition to providing an in-depth analysis of the sources of power dissipation in VLSI circuits and systems and the technology and design trends, this book provides a myriad of state-of-the-art approaches to power optimization and control. The different chapters of Power Aware Design Methodologies have been written by leading researchers and experts in their respective areas. Contributions are from both academia and industry. The contributors have reported the various technologies, methodologies, and techniques in such a way that they are understandable and useful.

"Previously, research and design of Network-on-Chip (NoC) paradigms where mainly focused on improving the performance of the interconnection networks. With emerging wide range of low-power applications and energy constrained high-performance applications, it is highly desirable to have NoCs that are highly energy efficient without incurring performance penalty. In the design of high-performance massive multi-core chips, power and heat have become dominant constrains. Increased power consumption can raise chip temperature, which in turn can decrease chip reliability and performance and increase cooling costs. It was proven that Small-world Wireless Network-on-Chip (SWNoC) architecture which replaces multi-hop wire-line path in a NoC by high-bandwidth single hop long range wireless links, reduces the overall energy dissipation when compared to wire-line mesh-based

NoC architecture. However, the overall energy dissipation of the wireless NoC is still dominated by wire-line links and switches (buffers). Dynamic Voltage Scaling is an efficient technique for significant power savings in microprocessors. It has been proposed and deployed in modern microprocessors by exploiting the variance in processor utilization. On a Network-on-Chip paradigm, it is more likely that the wire-line links and buffers are not always fully utilized even for different applications. Hence, by exploiting these characteristics of the links and buffers over different traffic, DVFS technique can be incorporated on these switches and wire-line links for huge power savings. In this thesis, a history based DVFS mechanism is proposed. This mechanism uses the past utilization of the wire-line links & buffers to predict the future traffic and accordingly tune the voltage and frequency for the links and buffers dynamically for each time window. This mechanism dynamically minimizes the power consumption while substantially maintaining a high performance over the system. Performance analysis on these DVFS enabled Wireless NoC shows that, the overall energy dissipation is improved by around 40% when compared Small-world Wireless NoCs."--Abstract.

Abstract: A multiprecision multiplier blends variable precision, parallel processing, dynamic voltage scaling and operand arrangement techniques to give improved performance over a wide range of operating conditions. The sub-blocks of the multiplier can function either as low precision multipliers when separated or as high precision multipliers when integrated. A dynamic voltage and frequency management unit influences the operation of the multiplier by adjusting the voltage and frequency so that the multiplier works with improved accuracy. The razor flip-flops along with the voltage dithering unit adjust themselves to the real world run-time of the load and increase the performance of the multiplier to obtain the least power usage. This combination decreases the error margins for voltage safety. Furthermore, the operand scheduler organizes the inputs, which helps the entire system in achieving improved voltage and frequency values, decreasing power consumption. This process iterates until power consumption can not be further reduced.

Increasing performance demands in integrated circuits, together with limited energy budgets, force IC designers to find new ways of saving power. One innovative way is the presented adaptive voltage scaling scheme, which tunes the supply voltage according to the present process, voltage and temperature variations as well as aging. The voltage is adapted “on the fly” by means of in-situ delay monitors to exploit unused timing margin, produced by state-of-the-art worst-case designs. This book discusses the design of the enhanced in-situ delay monitors and the implementation of the complete control-loop comprising the monitors, a control-logic and an on-chip voltage regulator. An analytical Markov-based model of the control-loop is derived to analyze its robustness and stability. Variation-Aware Adaptive Voltage Scaling for Digital CMOS Circuits provides an in-depth assessment of the proposed voltage scaling scheme when applied to an arithmetic and an image processing circuit. This book is written for engineers interested in adaptive techniques for low-power CMOS circuits.

Circuit designers perform optimization procedures targeting speed and power. Gate sizing can be applied to optimize for speed, while Dual-VT and Dynamic Voltage Scaling (DVS) can be applied to optimize for leakage and dynamic power, respectively. Both gate sizing and Dual-VT are design-time techniques applied to the circuit at a fixed voltage. DVS is a run-time technique and implies that the circuit will be operating at a different voltage than that used during optimization at design-time. After some analysis, the risk of non-critical paths becoming critical paths at run-time is detected under these circumstances. The following questions arise: 1) should we take DVS into account during optimization? 2) Does DVS impose any restrictions to design-time circuit optimizations? This is a case study of applying DVS to a circuit that has been optimized for speed and power, and aims at answering the previous two questions. We used a 45-nm CMOS design kit and flow for ISCAS'85 c432. Results showed that we should not optimize using Dual-VT at an arbitrary voltage but at the lowest in the DVS range, otherwise non-critical paths might become critical paths at run-time.

System-Level Design Techniques for Energy-Efficient Embedded Systems addresses the development and validation of co-synthesis techniques that allow an effective design of embedded systems with low energy dissipation. The book provides an overview of a system-level co-design flow, illustrating through examples how system performance is influenced at various steps of the flow including allocation, mapping, and scheduling. The book places special emphasis upon system-level co-synthesis techniques for architectures that contain voltage scalable processors, which can dynamically trade off between computational performance and power consumption. Throughout the book, the introduced co-synthesis techniques, which target both single-mode systems and emerging multi-mode applications, are applied to numerous benchmarks and real-life examples including a realistic smart phone.

An important issue in the development of mobile embedded systems is the optimization of energy consumption. The increasing complexity of these systems leads to an increased power consumption. Unfortunately, the battery capacity grows at a smaller rate than the power requirements from the applications, leaving a gap that must be filled by various system-level optimizations. Dynamic voltage scaling and adaptive body-biasing are well-known techniques to reduce dynamic and leakage energy effectively. These voltage settings can be calculated either offline, before the actual execution of tasks or online, during runtime. The advantage of online optimizations is the utilization of dynamic slack that results from variations in the execution time. However, these calculations are expensive and time-consuming. This book introduces a quasi-static algorithm for combined supply voltage scaling and body biasing. Voltage settings are computed offline and readjusted during runtime with very low online time complexity. This work considers time-constrained and discrete multiprocessor systems as well as energy and time overheads for the recalculation and readjustement of the voltages.

1. 1 Power-dissipation trends in CMOS circuits Shrinking device geometry, growing chip area and increased data-processing speed performance are technological trends in the integrated circuit industry to enlarge chip functionality. Already in 1965 Gordon Moore predicted that the total number of devices on a chip would double every year until the 1970s and every 24 months in the 1980s. This prediction is widely known as "Moore's Law" and eventually culminated in the Semiconductor Industry Association (SIA) technology road map [1]. The SIA road map has been a guide for the in dustry leading them to continued wafer and die size growth, increased transistor density and operating frequencies, and defect density reduction. To mention a few numbers; the die size increased 7% per year, the smallest feature sizes decreased 30% and the operating frequencies doubled every two years. As a consequence of these trends both the

number of transistors and the power dissi pation per unit area increase. In the near future the maximum power dissipation per unit area will be reached. Down-scaling of the supply voltage is not only the most effective way to reduce power dissipation in general it also is a necessary precondition to ensure device reliability by reducing electrical fields and device temperature, to prevent device degradation. A draw-back of this solution is an increased signal propa gation delay, which results in a lower data-processing speed performance.

To the hard-pressed systems designer this book will come as a godsend. It is a hands-on guide to the many ways in which processor-based systems are designed to allow low power devices. Covering a huge range of topics, and co-authored by some of the field’s top practitioners, the book provides a good starting point for engineers in the area, and to research students embarking upon work on embedded systems and architectures.

This book constitutes the refereed proceedings of the 15th International Workshop on Power and Timing Optimization and Simulation, PATMOS 2005, held in Leuven, Belgium in September 2005. The 74 revised full papers presented were carefully reviewed and selected from numerous submissions. The papers are organized in topical sections on low-power processors, code optimization for low-power, high-level design, telecommunications and signal processing, low-power circuits, system-on-chip design, busses and interconnections, modeling, design automation, low-power techniques, memory and register files, applications, digital circuits, and analog and physical design.

200-MHz Digital Low-dropout Regulator with Dynamic Voltage Scaling for Power ManagementPower Switch Characterization for Fine-grained Dynamic Voltage ScalingDynamic Voltage Scaling Techniques for Power-efficient MPEG DecodingImpact of Dynamic Voltage Scaling on Nano-Scale Circuit Optimization45-nm CMOS TechnologyLAP Lambert Academic Publishing