

Integer linear programming and heuristic techniques for system-level low power scheduling on multiprocessor architectures under throughput constraints

Krishnan Srinivasan, Karam S. Chatha*

Department of CSE, PO BOX 875406, Arizona State University, Tempe, AZ 85287-5406, USA

Received 3 April 2005; received in revised form 21 December 2005; accepted 11 January 2006

Abstract

The increased complexity and performance requirements of embedded systems has led to the advent of programmable multiprocessor architectures. The paper presents system-level design techniques for minimizing the power consumption of throughput constrained periodic applications (such as multimedia and network processing) that are mapped to multiprocessor architectures. The paper discusses several design techniques that integrate dynamic voltage scaling (DVS) along with loop transformations (pipelining and unrolling), and apply dynamic power management (DPM) as the final design step. The paper presents an optimal mixed-integer linear programming (MILP) formulation along with three modifications that trade-off solution quality for reduced run times. The paper also presents a heuristic technique along with deterministic ($LPPWU_{det}$) and simulated annealing based ($LPPWU_{sa}$) optimization strategies for solving the system-level low power design problem.

The proposed techniques are evaluated by extensive experimentation with multimedia applications (MPEG-1 decoder, JPEG decoder, MP3 encoder), and synthetic taskgraphs (with 10–40 nodes). The proposed techniques are compared with two existing strategies that apply (i) loop pipelining (FP) with DPM but no DVS, and (ii) DVS and DPM but no loop transformations (LPS [Luo, Jha, Power conscious joint scheduling of periodic task graphs and aperiodic tasks in distributed real-time embedded systems, Proceedings of the International Conference on Computer Aided Design, November 2000]), respectively. The optimal $MILP$ formulation, $LPPWU_{det}$ and $LPPWU_{sa}$ give an average power reduction of 50.68%, 48.57% and 49.23%, respectively, for multimedia applications when compared against FP . While all our techniques are able to satisfy the performance constraints for JPEG and MPEG-1 decoding applications, the LPS technique fails in many cases. Further, the results produced by our deterministic and simulated annealing based techniques for multimedia benchmarks are on an average within 8.04% and 2.75%, respectively, of the optimum solution produced by the MILP based approach. The experimentation with large synthetic taskgraphs demonstrate that the run times of the heuristic techniques scale very well.

© 2006 Elsevier B.V. All rights reserved.

Keywords: Lowpower design; Dynamic voltage scaling; Dynamic power management

1. Introduction

Embedded system applications in multimedia and network processing domains have witnessed an increase in complexity and performance requirements. These two factors coupled with the need for shorter design turn around times and ease of future upgrades have led to the advent of programmable multiprocessor architectures for

design of such applications. Examples of commercial multiprocessor architectures aimed at these embedded applications include System-on-Chip (SoC) designs such as Intel IXP series (IXP1200, IXP2400, IXP2800) processors [2], TI TMS320C8x [3], Motorola C-port C-5 [4], and board level implementations such as Sun SX2500 board [5], Alacorn FastImage 1500 [6], Synergy microsystem's MantaQX [7]. These architectures are deployed in portable devices (such as DVD players, digital cameras), set top boxes (HDTV), edge and backbone routers.

All these implementations have low power consumption as a key design requirement. The portable devices are

*Corresponding author. Tel.: +1 480 727 7850; fax: +1 480 965 2751.

E-mail addresses: ksrinivasan@asu.edu (K. Srinivasan),
kchatha@asu.edu (K.S. Chatha).

constrained by battery lifetime. The set top boxes, edge and back bone routers have thermal budgets which in turn translate into power consumption constraints. Consequently, innovative system level low power design techniques are required for the implementation of these embedded applications on multiprocessor architectures.

System-level low power optimization is enabled by utilization of the dynamic power management (DPM) [8] and dynamic voltage scaling (DVS) [9] (also known as dynamic voltage frequency scaling) capabilities of a processor. Both, DPM and DVS were developed to address the challenge of increased power consumption in CMOS devices. DPM exploits the idle times in application behavior and turns off the power supply to several sub-systems of the processing element. Therefore, DPM reduces the stand-by power or leakage power consumption of the application. For example, the Intel SA1100 StrongARM processor [10] has two low power DPM modes namely idle and sleep in addition to the normal run mode. In the idle mode the CPU clock is switched off and in the sleep mode the power supply to majority of the chip is turned off.

DVS trades-off the active power consumption with performance while executing the non-timing critical portions of the application. In CMOS technologies the active power consumption of a chip can be specified as $P = C_{\text{switch}} \times V^2 \times f$ where C_{switch} is the average capacitance switched per clock cycle, V is the operating voltage and f is the frequency of operation. The propagation delay through a CMOS inverter can be approximated as $t_p = C_L/2V(1/k_p + 1/k_n)$ [11] where C_L is the load capacitance, V is the operating voltage, and k_p and k_n are the process gain factors of the p- and n-type devices, respectively. The frequency of the device can then be calculated as $f = 1/t_p = 2V/C_L(k_p k_n/(k_p + k_n)) = KV$, where $K = 2/C_L(k_p k_n/(k_p + k_n))$. Thus, the frequency of a processor in CMOS technology is linearly dependent on the operating voltage. Therefore, reduction in the supply voltage results in a cubic reduction in power consumption at the expense of a linear slow down in the processor speed. DVS exploits this relationship to provide variable operating voltages and corresponding frequencies for the processor. For example, the Intel SA1100 StrongARM processor [10] supports supply voltages that range from 0.8 to 1.5V with corresponding operating frequencies ranging from 59 to 206 MHz.

DPM and DVS can be currently applied to multiprocessor board level architectures. In SoC based architectures, the emerging globally synchronous locally asynchronous design methodology [12] with multiple voltage and clock islands would also include DPM and DVS techniques.

Embedded system applications in the multimedia and network processing domains demonstrate periodic behavior. Real life implementations of these applications allow a greater time to process each data block than the period (period = 1/throughput). Time to process each data block is typically denoted by latency or deadline. For example,

the MPEG-1 video decoder is required to support a stream of upto 1.5 Mbps. If we consider that the decoder in each execution generates a macroblock of 8×8 (64 values) of 8 bits each, then the throughput requirement can be specified as 2929 macroblocks/s (approximately) or a period of 345 μ s/block. If the stream is intended for human viewing, an initial latency of 100 μ s will not result in detectable degradation. In other words, an MPEG-1 decoder that generates a steady stream of 2929 macroblocks/s (period = 345 μ s/block) with a latency of 445 (100 + 345) μ s/block will satisfy the performance requirements. Performance constraints with deadline greater than period coupled with multiprocessor target architecture enable the application of system level loop transformations such as pipelined scheduling and unrolling.

Pipelined scheduling and loop unrolling are two powerful loop transformations for throughput maximization of applications. Pipelined scheduling constructs a steady state that overlaps instances of tasks belonging to different iterations of the original specification. In the steady state there is only one instance of every task in the application. Unrolling as the name suggests transforms the original specification by replicating successive iterations. Thus, the transformed specification includes more than one instance of each task in the original loop. The two transformations can also be applied in an integrated manner where the specification is first unrolled and then scheduled into a pipeline.

This paper presents system level design techniques that integrate loop transformations with DPM and DVS to minimize the power consumption of applications mapped to multiprocessor architectures. In the following two sections we present a motivating example for the techniques discussed in this paper.

1.1. Embedded multiprocessor architecture

The techniques that are presented in this paper are applicable to a large class of embedded multiprocessor SoC architectures. In a typical SoC architecture each processing engine has its own code and data memory. The SoC architectures (such as Intel IXP series processors) provide sufficient on-chip memory to map the code and data segments of typical tasks. Embedded system application domains such as multimedia and network processing demonstrate clear demarcation of producer and consumer tasks with well defined inter-task communication behavior. Due to these characteristics and the application specific nature of the design, the designer can statically allocate data and code memory for each task, and data memory for inter-processor communication. Support for mutual exclusion during memory accesses is provided by the memory controller in hardware.

Existing SoC architectures implement a multitude of memory configurations to support inter-processor communication. The memory configurations for inter-processor

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات