

System-Level Synthesis of Low-Power Hard Real-Time Systems

Darko Kirovski and Miodrag Potkonjak

Computer Science Department, University of California, Los Angeles, CA 90095-1596

Abstract

We present a system-level approach for power optimization under a set of user specified costs and timing constraints of hard real-time designs. The approach optimizes all three degrees of freedom for power minimization, namely switching activity, effective capacity and voltage supply.

We first define two key associated optimization problems, processor allocation and task assignment, and establish their computational complexity. Efficient algorithms are developed for both system design problems. The statistical analysis of comprehensive experimental results and their comparison with the developed conservative and optimistic sharp lower bounds, clearly indicates the quality of the proposed optimization techniques.

1 Introduction

The application and technological factors behind the imminent convergence of computer, communications, and consumer electronic products and market impose a new set of design goals and constraints for the synthesis process of application specific systems. The short time-to-market and cost sensitivity of consumer markets imply a need for optimization intensive system-level synthesis. In addition, multiplicity of standards and a need for flexibility and customization favor programmable implementation platforms. The popularity of the portable nature of future personal digital assistants and communication devices is the principal motive for low-power design technique development [1].

The research presented in this paper answers the outlined design requirements by introducing a system of optimization intensive synthesis techniques and tools for design of hard real-time application specific systems implemented on a set of programmable units. The synthesis approach encompasses two fully modular tasks: hardware allocation and task assignment. All three degrees of freedom for power minimization, namely switching activity, effective capacity and voltage supply are considered. While the first two degrees are described by processors' power efficiency for a given task, the last one depends on both the run time of a task on a particular unit and balanced loads on all processors. In order to minimize the communication cost, tasks are assigned atomically to individual units.

Resource allocation and task assignment for power efficient hard real-time systems are difficult problems with several layers

of conceptual and computational complexity. For example, if we analyze the assignment problem, where a set of tasks is assigned to a set of units, the power can be minimized either by executing tasks on power efficient units or scaling the operational voltage. However, voltage scaling is closely related to the time efficiency of particular task-to-processor assignments and balanced minimized processor workloads. From the optimization point of view the key novelty is the use of statistical meta-algorithmic techniques for optimization of parameters that dictate the assignment algorithm. From statistical and result evaluation points of view, the key novelties are related to several new techniques for establishing the quality of heuristic algorithms.

The key trade-offs and ideas used to develop the synthesis algorithms can be introduced using the following motivational example. For the sake of brevity we restrict our attention to the assignment phase of the synthesis approach.

A set of tasks $\{T_i, i=1..6\}$ has to be assigned to a given hardware configuration so that the power consumption is minimized. Each task T_i requires a certain number of cycles $T_{i,j}$ to execute on a particular processing element $\{PE_j, j=1..3\}$. The $C_{i,j}$ data are shown on the left side of Table 1. Each unit PE_j consumes a specific amount of energy E_j per operation. Thus, the energy $PE_{i,j}$ required for one execution of each task T_i can be computed as $P_{i,j} = PE_j \cdot C_{i,j}$. The right side of the Table 1 provides the data on the task execution energies $P_{i,j}$. The tasks are independent and have to be executed exactly once per each iteration (21 μ s). A voltage supply of 3.3V is assumed. All units, when driven by 3.3V power supply, can support a clock cycle delay of 5ns, implying 200MHz operational frequency.

| | Clock Cycle | | | Energy [μ J] | | |
|-------------------------------|-------------|--------|--------|-------------------|--------|--------|
| | PE_1 | PE_2 | PE_3 | PE_1 | PE_2 | PE_3 |
| Energy per cycle of operation | 200 | 20 | 5 | | | |
| T_1 | 100 | 250 | 900 | 20 | 5 | 4.5 |
| T_2 | 120 | 250 | 800 | 24 | 5 | 4 |
| T_3 | 140 | 300 | 800 | 28 | 6 | 4 |
| T_4 | 140 | 200 | 400 | 28 | 4 | 2 |
| T_5 | 140 | 400 | 500 | 28 | 8 | 2.5 |
| T_6 | 130 | 800 | 800 | 26 | 16 | 4 |

Table 1: Task cycle and energy efficiencies on the allocated hardware.

Power can be optimized by assigning tasks to processors which consume the least amount of power for their execution and by scaling the voltage of the power supply. The voltage is scaled with respect to the maximal cycle time required by the hardware configuration and its workload. Maximally scaled voltage corresponds to the cycle time allowed by the delay vs. voltage digital CMOS dependency.

One attractive option for task assignment is to assign each task to a processor which consumes the least amount of power for its completion (PE_3). This task assignment requires 4200 cycles. The

Design Automation Conference (®)

Copyright © 1997 by the Association for Computing Machinery, Inc. Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Publications Dept, ACM Inc., fax +1 (212) 869-0481, or permissions@acm.org.

0-89791-847-9/97/0006/\$3.50

DAC 97 - 06/97 Anaheim, CA, USA

minimal frequency at which the circuit has to be clocked is $\frac{4200}{21\mu s} = 200MHz$. The power consumption totals **21mJ**.

The operational voltage of the system can be reduced further by balancing the workloads on all units. Notice that the previous optimization approach can be executed in linear time, while the problem of balancing the workloads on all processors is proven to be NP-hard. Exhaustive search has shown that the following task assignment achieves the minimal number of cycles when: T_1, T_2, T_5, T_6 on P_1 ; T_3 on P_2 ; and T_4 on P_3 . Maximally balanced workload requires 490 cycles. Now, the clock cycle can be increased $\frac{200MHz}{\frac{490}{21\mu s}} = 8.57$ times, resulting in voltage decrease from 3.3V to 1.29V. The power consumption is $(P_{1,1} + P_{2,1} + P_{3,2} + P_{4,3} + P_{5,1} + P_{6,1}) \frac{V_1^2}{V_0^2} = 106mJ \frac{1.29^2 V^2}{3.3^2 V^2} = 16mJ$.

Although maximally lowered operational voltage appears to be the most important criteria considered for power optimization, proper analysis shows that the best solution can be obtained when the advantages of the cycle-, power- and balance-greedy optimization criteria are combined. For example, consider the following assignment: T_1, T_2, T_3, T_4 on P_2 ; T_5, T_6 on P_3 ; which requires 1300 cycles. The cycle time can be increased $\frac{200MHz}{\frac{1300}{21\mu s}} = 3.23$ times, yielding 1.70V operational voltage. The total power consumption is $(P_{1,2} + P_{2,2} + P_{3,2} + P_{4,2} + P_{5,3} + P_{6,3}) \frac{V_1^2}{V_0^2} = 26.5mJ \frac{1.70^2 V^2}{3.3^2 V^2} = 7mJ$.

| Technique | Greedy | | Optimal |
|---------------|--------|---------|---------|
| | Power | Balance | Power |
| MaxFrequency | 200MHz | 23.4MHz | 62MHz |
| Voltage | 3.3V | 1.29V | 1.70V |
| Cycles | 4200 | 490 | 1300 |
| Energy (3.3V) | 21mJ | 106mJ | 26.5mJ |
| Scaled Energy | 21mJ | 16.3mJ | 7mJ |

Table 2: Comparison of greedy optimization techniques to the optimal solution.

This is an improvement of 130% and 200% in comparison to the balanced, and power-greedy solution respectively. The motivational example shows that high energy savings can be obtained even on small examples when optimization is treated as a combination rather than selection of the stated power optimization concepts.

The rest of this paper is organized in the following way. In Section 2 we present related research. Section 3 outlines the computational, task and power consumption models. Analysis of design trade-offs and overall design flow are presented in Section 4. Synthesis optimization problems are defined and their complexity is established in Section 5. The algorithms are described in detail in Section 6 and 7. Sections 8 and 9 contain a report on conducted experiments and conclusion respectively.

2 Related Work

System level synthesis, including hardware-software codesign and partitioning, is premier design and CAD research topic [12], [2]. The most relevant system research subdomain is hardware-software partitioning, where a great variety of techniques are proposed [3], [4], [5].

A number of synthesis and compilation techniques for power optimization at all levels of abstractions during the design process have been proposed [9]. Although power optimization is often most effective when applied early in the design process at the algorithmic and architectural level, majority of power minimization techniques have been proposed for logic synthesis and physical design

[9]. Potkonjak and Rabaey presented a CAD system-level synthesis approach which targets low-power [8].

Hard-real time scheduling efforts have been an important topic for research since early computing days. A survey of the key results in hard real-time scheduling and assignment is given in [10].

3 Preliminaries

We assume that all tasks are periodic and defined as one-way infinite streams of data. Each task has its own period of appearance. This computational model is the standard abstraction for many video, digital signal processing, communication, and multimedia applications.

In the period equal to the Least Common Multiple (LCM) of all periods of all tasks, all instances of a single task have to be executed. Their execution has to satisfy the deadline that is set by the LCM of all task periods. Tasks are independent which means there is no cost for intertask communication. The assumption of task independence is done with no loss of generality whatsoever, because atomic independent tasks can be treated as a single composite task.

The targeted architecture template is shown in Figure 1. The system contains N programmable processors (CPUs). Each CPU has a reserved segment of data memory. Programs that execute the tasks are stored in EPROMs. All units are subject to the common voltage supply source.




Figure 1: High-level architectural template.

In CMOS digital circuits the dynamic power consumption is given by the following formula

$$Energy \approx SwitchingActivity \cdot C_{eff} \cdot V_{dd}^2$$

Since the power consumption is proportional to the square of the voltage supply, its lowering is often the main concern in every power optimization technique. However, lowering the supply voltage causes increased delay of CMOS circuits according to the following formula [1]:

$$Delay = K \cdot \frac{V_{dd}}{(V_{dd} - V_i)^2}$$

where K and V_i are constants entirely dependent upon the implementation technology.

When assigning tasks to the allocated hardware we minimize and balance the workloads on all processors. The deadline, set by the periodic nature of the incoming datastream, imposes a boundary on the delay of a single clock cycle of the allocated resources: $Deadline > Cycles \cdot Delay$, and therefore, limits system designer's ability to lower the supply voltage.

The other two parameters that impact power consumption, the effective capacity and switching activity, are optimized implicitly in two phases of our synthesis approach. First, for a given set of tasks we select hardware that has the least effective capacity and

switching activity for the execution of the entire set of tasks. Then, during the task assignment phase, the processor of a minimal effective capacity and switching activity is given a certain preference.

4 The New Synthesis Approach

The synthesis procedure starts with the resource allocation step. The allocation algorithm searches for the processor configuration which consumes the least power for a given set of tasks under a given cost constraint of the system. The allocation algorithm is based on a multistart gradient search. For the evaluation of the proposed solution, the allocation algorithm invokes the task assignment procedure. Once the allocation is completed, the final assignment starts. The assignment heuristic algorithm is developed by a new statistical meta-algorithm methodology explained in Section 7. The assignment procedure iteratively assigns one task at a time. The decision of selecting tasks and target processors is made globally; each time a task is considered for assignment, its impact on further assignment of remaining tasks is encountered. The task is selected upon its relative size, and the most efficient power and cycle requirements for execution. Processor selection depends upon the processors' current workloads, and power and cycle efficiencies for the execution of the selected task. Once all tasks are assigned, the voltage of the power supply is scaled as dictated by the overall timing deadline and maximum cycle workload. Detailed explanations of the developed algorithms for each synthesis step are given in Sections 6. and 7.

5 Optimization Problems

We now formulate the optimization problems associated with system-level synthesis of low power hard real-time multitask designs and establish their computational complexity. Our synthesis approach has two optimization intensive phases: allocation and assignment.

Problem: Allocation for synthesis of low power hard real-time application specific systems.

Instance: A set of K processors and a set of N independent periodic hard real-time tasks are given. Each processor has an associated cost and power consumption per cycle. Each task has its period and execution time for assignment to any of the processors. Two positive numbers C and P.

Question: Select a multisubset of processors (subset where some processors can be included more than once) such that each task is assigned to exactly one processor, the sum of costs of selected processors is at most C, and the power consumed by all tasks is at most P.

Problem: System-level Assignment of Hard Real-Tasks.

Instance: A set of K processors and a set of N independent periodic hard real-time tasks are given. Each task has its period and execution time (number of cycles) for each of the available processors.

Question: Assign each task to one of the processors in such a way that all tasks can be scheduled within their timing constraints and that the power consumed by all tasks is at most P.

We proved that system-level synthesis optimization problems for low power hard real-time multitask systems are NP-complete in [6].

6 Resource Allocation

The first synthesis step is resource allocation. The goal is to select a multiset from the set of available processors, so that the power consumption of the hard real-time system is minimized under the maximum cost constraint. The inputs to the resource allocation

program are costs and power consumptions per cycle of all available processing units, denoted by C_i and PE_i respectively, and the total user specified system cost C. Also, the information about the execution time of each task on each type of processor is available. This information can be obtained either by executing tasks on a given processor or by using simulation and estimation techniques. Since the resource allocation is computationally intractable problem, we opted to use heuristic search which provides a favorable run-time vs. quality trade-off. The synthesis algorithm starts with a preprocessing step, which eliminates all types of processors dominated by other types of processors. Processor PE_x is inferior to processor type PE_y , if PE_x is more expensive, performs all tasks in more clock cycles and requires more energy for the execution of each task. The minimum set of non-inferior processors is selected using the standard 3-D dominance algorithm [11].

```

Remove all inferior processors from the hardware menu
While (number of attempts j K)
Repeat
Generate initial random feasible solution
For all types of processors add or delete a processor from
the current configuration  $HC_i$ , and calculate the OPC as
well as the objective function  $OF_i = OPT_i + \alpha \cdot MC_i$ ,
where  $MC_i$  is the marginal cost of  $HC_i$ .
If  $C_i > C$  then  $MC_i = 0$  else  $MC_i = C_i - C$ .
Depending on the values of the objective function for all
types of processors, delete or add a processor of the
winning type and action.
Increase  $\alpha$ .
Until there is improvement and the cost of the configuration
satisfies the system requirements.
  
```

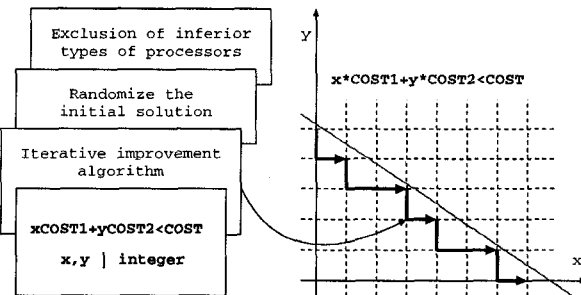


Figure 2: Resource allocation algorithm (example for two types of processors).

Once all inferior processing types are excluded from the hardware menu, the heuristic that searches for the most power-efficient configuration starts. The heuristic defines a random hardware configuration HC_0 which is used as a starting point for an iterative improvement heuristic. This configuration is constructed in such a way that it is not possible to add a single processor of any type to HC_0 so that it stays within the range of the system cost C. The initial configuration is established by randomly adding processing units, as long as the cost of the system is lower than C.

The iterative heuristic at each step tries all the options for addition or removal of any of the considered types of resources. The heuristic is gradient-driven, i.e. the configuration is upgraded by following the most beneficial action. The objective function is a weighted sum of the optimized power consumed (OPC) by the current configuration and the marginal cost of the system. The power is calculated by invoking a simplified task assignment procedure. If the difference is positive, the marginal cost of the system is defined as the excess of the cost of the current configuration over the user specified system cost C. If the configuration cost satisfies the

cost requirement, the marginal cost is equal to zero. The weighted sum of the two components is time dependent. As the optimization process progresses, the importance of the marginal cost over power increases, since it is more important to produce feasible solutions. Thus, at the beginning of the heuristic search, we impose relatively non-strict cost enforcement, while at the end, we strongly favor solutions that satisfy the cost constraint.

Finally, when no further improvements cannot be obtained by adding or deleting a processor of any type, we are likely to be stuck in one of the local minima of the space of all possible solutions. Therefore, we terminate our search and memorize the power-optimal hardware configuration. The core of the allocation procedure is repeated as long as no improvement is retrieved from the last k iterations of the allocation procedure.

The resource allocation procedure can be summarized using the pseudo-code on Figure 2.

7 Task Assignment

Whenever the allocation algorithm generates a hardware configuration it invokes a simplified task assignment heuristic in order to estimate the power consumption of the hardware configuration. The task assignment algorithm iteratively selects a task for assignment and then assigns the selected task to the allocated hardware. The post-processing step aims at additional load balancing. The three key steps of the algorithm are described in the remainder of this section.

7.1 Task Selection

The first step during each iteration of the assignment procedure is to select a task which will be assigned to one of the processors. The selection is conducted using the following objective function:

$$Select_{task_i} = \alpha_1(T_{2,i} - T_{1,i}) + \alpha_2(P_{2,i} - P_{1,i}) + \alpha_3 \frac{T_{2,i} + T_{1,i}}{\sum_{j=1}^{tasks} (T_{2,j} + T_{1,j})} + \alpha_4 + \alpha_5 + \alpha_6$$

where $T_{1,i}$ denotes the minimal number of cycles required to execute the task T_i , $T_{2,i}$ denotes the second minimal number of cycles required to execute the task T_i , $P_{1,i}$ is the most power-efficient processor for executing T_i , and $P_{2,i}$ is the second most power-efficient processor for executing T_i .

The selection function components represent the key trade-offs among the considered mechanisms for power minimization while the weight parameters dictate their relative importance. The first component of the objective function favors assigning a task to a processor where the task has minimal execution time. Since short execution times enable voltage scaling, the difference $T_{2,i} - T_{1,i}$ encountered in the objective function includes the lower bound for not following the most run-time efficient option. The second component favors assigning a task to a processor which consumes the minimal amount of energy at the nominal voltage. The difference $P_{2,i} - P_{1,i}$ captures the minimal penalty for not following the most power-efficient choice. Long tasks often cause workload balance problems if not scheduled early. A task T_i is termed long compared with the length of all other tasks if $\alpha_3 \cdot (T_{2,i} + T_{1,i}) / \sum_{j=1}^{tasks} (T_{2,j} + T_{1,j})$ is much greater than the same expression calculated for all other tasks. Parameters α_4 , α_5 , α_6 are values added to the selection function if the most power-efficient and the most cycle-efficient processor are the same (α_4), if the most power-efficient processor is the same as the second most cycle-efficient processor or if the most cycle-efficient processor is the same as the second most power-efficient processor (α_5), or if the second most power- and cycle-efficient processors are the same (α_6). Those three parameters capture the relative difficulty of making an optimal assignment

choice for the task. The task that has the highest value of its selection function is the one selected for assignment.

7.2 Processor Selection

Processor selection is based on the execution characteristics for the selected task on each processor. The selection function used in our approach is

$$Select_{processor_j, task_i} = \beta_1 T_{1, processor_j, task_i} + \beta_2 P_{1, processor_j, task_i} + \beta_3 workload_{offset}(processor_j)$$

where $workload_{offset}(processor_j)$ is equal to zero if the workload of $processor_j$ does not exceed the current maximum workload in the system. Otherwise it is equal to the difference of the current maximum workload and the workload after the task is assigned to this processor.

As in the task selection function, here we also deal with several trade-offs. The selected processor is the one with the smallest value of the selection objective function. The first and the second component favor the processor that performs the selected task in the most time- and power-efficient way respectively. The third component facilitates voltage scaling by load balancing. The quality of the final solution is a function of parameter sets α and β . In the next Section we describe how parameters for this function are obtained by using meta-algorithmics and statistical techniques.

Application of the described synthesis algorithms to the motivational example is presented in [6].

7.3 Meta-Algorithms

In the previous subsection, we discussed the key trade-offs and optimization mechanisms which impact the decision process and the quality of the final design produced by the iterative task selection and assignment process. Common engineering strategy is to start with ad-hoc values for the parameters using insights obtained on small manually done examples. In the second phase, a majority of traditional approaches usually use the feedback from executed benchmarks to tune the values of the parameters and therefore to determine their relative importance. However, the traditional approaches most often do not provide any formal, statistical, or even intuitive guarantees of the optimality of the achieved result.

In this subsection we outline a new meta-algorithmic approach for statistical parameter determination and validation and demonstrate its application on our assignment procedure for synthesis of hard real-time systems. The key steps of the new methodology for CAD algorithm development are described using the following pseudo-code:

```
Meta-Algorithmic Procedure() {
    generate_learning_examples();
    generate_optimization_mechanism();
    parameter_evaluation();
}
```

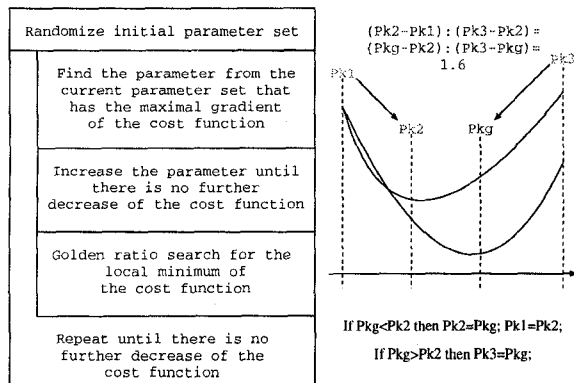
As a first step to this strategy, we generated sets of statistically identical benchmark examples using the set of available real-life examples [6]. We prepared a set of K statistically random sets of M tasks and N different hardware options. For the results presented in this and next section, we used the following set of values: $K = 150$, $M = 20 - 320$, and $N = 5$.

The second step, the generation of optimization mechanisms is explained in Section 7. The core of the meta-algorithmics approach is the parameter evaluation procedure given in Figure 3.

The core procedure performs multistart gradient search. One iteration of the program for parameter evaluation first sets the parameters to random values. Then it checks the total power consumptions (TPC) when each parameter is positively and negatively

offset by $k\%$ of its value. In our experimentations we used $k = 2$. Other parameters in range $\{0.01-10\}$ produced very similar results. The parameter that has the smallest TPC is then offset to the value that caused this improvement, following the deepest descent paradigm. Then the procedure continues exponentially increasing or decreasing the selected parameter using the golden ratio search approach. The one iteration of the golden ratio search is terminated once the improvement is smaller than a user specified value ϵ (we used value 0.1%).

Once the golden ratio search is completed, the list of the parameters is updated to this new value and the procedure is iteratively repeated as long as local minimum is not reached using our stopping criteria. Then we start a search for the minimum value of TPC by adjusting the parameter P_k . The one iteration of the search is terminated once when no improvement along directions indicated by any of the parameters is not detected.



```

Set all parameters  $P_i$  to random values.
For each parameter  $P_i$ 
  Calculate  $TPC_i^+$  when  $P_i = P_i \cdot (1 + w)$ 
  and the  $TPC_i^-$  when  $P_i = P_i \cdot (1 - w)$ .
 $TPC_k^@ = \min(TPC_i)$ , where  $@ = \text{sign}(\text{variation}(P_k))$ .
Set  $P_k$  to  $P_k \cdot (1 @ w)$ .
Repeat
  Set  $P_k$  to  $P_k \cdot (1 @ w)$ .
until TPC of the current set of parameters is greater
than the TPC of the previous set of parameters.
Memorize last  $P_k, P_{k1}, P_{k2}, P_{k3}$ .
Repeat
   $P_k = P_{kg} = \text{GoldenRatio}(P_{k1}, P_{k3})$ .
  If the  $TPC_{kg}^-$  is greater than the  $TPC_{k2}$  then  $P_{k3} = P_{kg}$ 
  else  $P_{k1} = P_{k2}$  and  $P_{k2} = P_{kg}$ .
until  $TPC_{kg}^- - TPC_{k2} > \epsilon$ .

```

Figure 3: Meta-algorithmic evaluation of parameters.

The intuition behind the evaluated parameters is the following: the most important is to early select a task for which there is a cleanly preferable assignment as indicated by parameters α_1, \dots, β_3 . During the assignment it is relatively important to follow balancing, unless there are high timing and power advantages. The evaluated parameters were used in the next section to conduct experiments that proved our assumptions about the efficiency of our synthesis paradigm.

8 Experimental Results

In this section we describe our benchmark examples, experimental results evaluation strategy, and present and analyze the experimen-

tal results. Our goal is to evaluate the effectiveness of proposed algorithms on task mixes that approximate the expected task mixes of application specific space program computers, wireless portable personal communicators, and module application specific servers. Initially we assembled 40 typical DSP, video, and communication tasks [6]. These sets of tasks provide a solid approximation of expected loads for emerging DSP, video, and broadband servers.

| Task(cpu) | Bounds | | Results | | Ratio |
|-----------|--------|------|---------|-----------|--------|
| | Lower | Meta | Random | Heuristic | % |
| 20(5) | 217 | 247 | 285 | 255 | 105.85 |
| 40(5) | 480 | 533 | 635 | 565 | 112.39 |
| 40(10) | 390 | 616 | 842 | 632 | 133.23 |
| 80(5) | 901 | 997 | 1186 | 1030 | 115.15 |
| 80(10) | 767 | 1130 | 1591 | 1155 | 137.75 |
| 80(20) | 698 | 925 | 1779 | 943 | 188.65 |
| 160(5) | 1816 | 2004 | 2384 | 2067 | 115.34 |
| 160(10) | 1543 | 2255 | 3096 | 2296 | 134.84 |
| 160(20) | 1449 | 1827 | 3270 | 1839 | 177.81 |
| 160(40) | 1399 | 1684 | 4160 | 1736 | 239.63 |
| 320(5) | 3629 | 3984 | 4833 | 4110 | 117.59 |
| 320(10) | 3038 | 4389 | 6009 | 4414 | 136.14 |
| 320(20) | 2998 | 3774 | 6297 | 3807 | 165.41 |
| 320(40) | 2828 | 3311 | 7259 | 3388 | 214.26 |
| 320(80) | 3156 | 3746 | 12838 | 3783 | 339.36 |

Table 3: Power consumptions for 15 classes of sets of tasks (assignment; resources allocated).

The execution times of tasks on a variety of processors are either obtained through direct measurements on the available hardware platforms (e.g. SUN Sparcstation) or by using estimation formulas developed in [7] for a number of DSP processors. The power consumption data is obtained by combining data from manufacturers product descriptions and the well known fact that power consumption of an arbitrary task on majority of modern processors is linearly proportional to its execution time [7]. In order to test the performances of the proposed algorithms on large instances of input data, we used a statistical technique to generate synthetic mixes of tasks with identical statistical properties as the original test set.

First, histograms of run-times for each processor and all available tasks are constructed. The run times of additional tasks are generated using a random number generator that follows the distribution dictated by the histogram. All statistically generated examples are divided into two groups: learning and testing benchmarks. The learning examples are used to evaluate parameters of the assignment heuristics, the testing examples along with the original set are used for the evaluation purposes.

In order to properly evaluate our synthesis approach we developed the following two-phase approach. In the first phase the assignment algorithm is evaluated. First, we defined a lower bound solution. Then, we recorded the best results achieved for any proposed set of parameters during the application of the meta-algorithmic technique. We also generated a large set of random solutions and recorded the best random results. Finally, we applied the new assignment algorithm on the test examples.

The provably conservative lower bound used for the evaluation of our synthesis algorithms is developed in the following way. For each task the minimal amount of energy required to execute the task on any of the available processors is selected and added up to the overall power consumption of the entire set of tasks. The resulting power consumption is further reduced by scaling the voltage of the power supply according to timing requirements of the fastest perfectly balanced solution. In order to obtain a perfectly balanced

solution, for each task the minimal number of cycles required to execute the task is selected, added up and divided by the total number of processors in the configuration.

Note that meta-algorithmic approach by itself can be considered as an assignment algorithm. While the search for optimal assignment parameters progresses, the meta-algorithmics generates numerous assignments. We run the meta-algorithmics approach on all test examples and record the best solutions (resulting in minimal power consumption) regardless of the used parameters. The main disadvantage of the solution developed by the meta-algorithmics itself is the execution time of the program.

| Task(cpu) | Bound | Results | | Ratio |
|-----------------------|-------|---------|-----------|---------|
| | Lower | Random | Heuristic | % |
| 20(\$ ₁) | 148 | 442 | 255 | 173.307 |
| 40(\$ ₁) | 327 | 1263 | 565 | 223.557 |
| 40(\$ ₂) | 270 | 3602 | 632 | 569.143 |
| 80(\$ ₁) | 589 | 1799 | 1030 | 174.529 |
| 80(\$ ₂) | 510 | 6573 | 1155 | 568.698 |
| 80(\$ ₃) | 476 | 12488 | 943 | 1322.89 |
| 160(\$ ₁) | 1268 | 4452 | 2067 | 215.394 |
| 160(\$ ₂) | 1012 | 12939 | 2296 | 563.557 |
| 160(\$ ₃) | 970 | 18720 | 1839 | 1017.73 |
| 160(\$ ₄) | 942 | 37940 | 1736 | 2184.69 |
| 320(\$ ₁) | 2574 | 9524 | 4110 | 231.734 |
| 320(\$ ₂) | 2102 | 25331 | 4414 | 573.777 |
| 320(\$ ₃) | 2016 | 26581 | 3807 | 698.174 |
| 320(\$ ₄) | 1986 | 46620 | 3388 | 1375.95 |
| 320(\$ ₅) | 2249 | 188588 | 3783 | 4984.81 |

Table 4: Power consumptions for 15 sets of tasks (allocation and assignment).

In the second phase the allocation algorithm is evaluated. To obtain the lower bound solution we relax several constraints imposed by the proper integral solution requirement: we assume that fractional number of processors can be allocated, that best power and timing characteristics of processors can be combined, and that perfect balancing is achievable. These approximations transform the problem which has several layers of NP-hard subproblems into one which can be solved rapidly in polynomial time. This is achieved in the following way.

First, for each task the most time and power efficient processors are identified. For each task the nominal power consumption at the initial operational voltage is accepted as the most power efficient option. Next, the cost of cycle per dollar for each task on its most time efficient processor is calculated. This cost is multiplied by the ratio of the number of cycles vs. the iteration period for each task, and added up for all tasks. By dividing the total specified cost by this sum we obtain the factor by which the cycle time can be prolonged as well as the bound on how much we can reduce the supply voltage.

Similarly to the case of assignment evaluation, we generated the best random solution by using lower bound assignments in 300 random allocation runs. The relative performance of the new allocation algorithm versus the described three comparison options are given in Table 4.

The experimental results for the task assignment heuristic are shown in Table 3. The table shows power consumption for 15 different classes of sets of tasks and processors. Each class is represented by 10 different sets of constant number of tasks and units. The first column provides information about the number of tasks and processors. The first number in the first column corresponds to the number of tasks followed by the corresponding number of processors. The next four columns provide power consumption data

per iteration for lower and meta bound, best random result in 1500 runs, and the proposed new approach respectively.

Table 4 provides information about the effectiveness of the allocation algorithm. The first column contains information about the number of tasks and the number of processors used by our allocation algorithm. The next three columns provide power consumption produced by the lower bound algorithm, best random solution, and our solution.

9 Conclusion

We developed a system-level approach for power minimization of cost-constrained hard real-time designs. The approach simultaneously optimizes all three degrees of freedom for power minimization, namely switching activity, effective capacity and supply voltage. Novel meta-algorithmics approach strategy is used for computer-aided statistical optimization of the assignment algorithm. The comprehensive experimental results demonstrate the quality of the proposed optimization techniques.

Acknowledgments

The authors wish to acknowledge the support given by Okawa Foundation for this work.

REFERENCES

- [1] A.P. Chandrakasan and et. al. Low-power cmos digital design. *IEEE Journal of Solid-State Circuits*, 27(4):473-484, April 1992.
- [2] M. Chiodo and et. al. A case study in computer-aided co-design of embedded controllers. *Design Automation for Embedded Systems*, 1(1-2):51-67, January 1996.
- [3] D. D. Gajski, F. Vahid, and S. Narayan. A system-design methodology: Executable specification refinement. In *Euro-DAC*, pages 458-463, 1994.
- [4] R. K. Gupta and G. De Micheli. Hardware-software cosynthesis for digital systems. *IEEE Design and Test of Computers*, 10(7):29-41, September 1993.
- [5] T. B. Ismail, K. O'Brien, and A. Jerraya. Interactive system-level partitioning with partif. In *Euro-DAC*, pages 464-468, 1994.
- [6] D. Kirovski and M. Potkonjak. System level synthesis of low power hard real time systems. Technical Report 960051, CSD, UCLA, September 1996.
- [7] M.T.-C. Lee and et. al. Power analysis and low-power scheduling techniques for embedded dsp software. *Fujitsu Scientific and Technical Journal*, 31(2):215-229, 1995.
- [8] M. Potkonjak and W.H. Wolf. Cost optimization in asic implementation of periodic hard-real time systems using behavioral synthesis techniques. In *ICCAD*, pages 446-451, 1995.
- [9] D. Singh and et al. Power conscious cad tools and methodologies. *Proceedings of IEEE*, 83(4):570-594, 1995.
- [10] J.A. Stankovic, M. Spuri, M. Di Natale, and G.C. Buttazzo. Implications of classical scheduling results for real-time systems. *IEEE Computer*, 28(6):16-25, June 1995.
- [11] R.L. Rivest T.H. Cormen, C.E. Leiserson. *Introduction to algorithms*. MIT Press, Cambridge, MA, 1990.
- [12] W.H. Wolf. Hardware-software co-design of embedded systems. *Proceedings of IEEE*, 82(7):967-989, July 1994.