Dynamic and Leakage Power-Composition Profile Driven Co-Synthesis for Energy and Cost Reduction

D. Wu* and B. M. Al-Hashimi*
University of Southampton, UK

P. Eles
Linköpings University, Sweden

Abstract

Recent research has shown that combining dynamic voltage scaling (DVS) and adaptive body bias (ABB) techniques achieve the highest reduction in embedded systems energy dissipation [1]. In this paper we show that it is possible to produce comparable energy saving to that obtained using combined DVS and ABB techniques but with reduced hardware cost achieved by employing processing elements (PEs) with separate DVS or ABB capability. A co-synthesis methodology which is aware of tasks' power-composition profile (the ratio of the dynamic power to the leakage power) is presented. The methodology selects voltage scaling capabilities (DVS, ABB, or combined DVS and ABB) for the PEs, maps, schedules, and voltage scales applications given as task graphs with timing constraints, aiming to dynamic and leakage energy reduction at low hardware cost. We conduct detailed experiments, including a real-life example, to demonstrate the effectiveness of our methodology. We demonstrate that it is possible to produce designs that contain PEs with only DVS or ABB technique but have energy dissipation that are only 4.4% higher when compared with the same designs that employ PEs with combined DVS and ABB capabilities.

Keywords: low power system-level co-synthesis

1. Introduction

Energy reduction is an essential consideration in the design of embedded systems. Based on the fact that there are periods when applications do not require the maximum performance provided by the processing elements (PEs), a number of dynamic voltage scaling (DVS) [2-4] and adaptive body biasing (ABB) [5-8] techniques have been reported which enable a trade-off between energy and performance. DVS scales down the circuit supply voltage and operational frequency to reduce dynamic power, whilst ABB increases the circuit threshold voltage through body biasing to reduce leakage power. DVS is effective in reducing dynamic power, however, in deep sub-micron (DSM) designs, the leakage power contributes significantly to the total power consumption [9]. As an example, in [5] it was demonstrated that the overall leakage power of six benchmarks is only 9.8% of the total power for 0.35µm technology, but it is 56.2% for 0.07µm technology. In [10], it was shown that ABB would become comparable to DVS in achieving energy saving for future technologies. The energy reduction techniques reported in [2-8] employ either DVS or ABB to reduce energy dissipation.

* Supported in part by EPSRC under grant GR/S95770

However, for foreseeable future systems where dynamic power and leakage power are comparable to each other, neither DVS nor ABB in isolation can achieve the best energy efficiency. In [1], it was shown that a combined DVS and ABB will provide the highest energy saving. In such a combined scheme, DVS and ABB compete for the redundant system performance (i.e., slack time) to scale down the system's operational frequency. Therefore, the difficulty in employing simultaneous DVS and ABB is to determine the trade-off between them [11]. In [9], analytical models were developed to produce optimal supply voltage and threshold voltage values for energy minimisation. In [12, 13] authors addressed the supply voltage and body bias voltage scaling problem for task graphs executed on multi-processor systems, unlike the work in [9] which targets single processor systems.

Although the combination of DVS and ABB [9, 11-13] (referred to as DVS+ABB from now on) provides higher energy saving than DVS or ABB in isolation, it increases hardware cost and complexity. This is because the implementation of DVS requires an additional voltage converter and impacts processor verification [14]. The implementation of ABB, on the other hand, requires a bias voltage generator and a substrate-bias distribution, i.e., additional wiring [6, 7]. Clearly, combined DVS+ABB implies larger hardware cost than separate DVS or ABB. Cost and energy are two main design considerations in embedded systems; therefore, it is important to carefully balance these two aspects as early as possible in the design phase. In this paper we will show that, depending on the ratio of the dynamic power to the leakage power (i.e., power-composition profile), it is possible to employ PEs with separate DVS or ABB capability to achieve comparable energy saving while avoiding the extra cost of PEs with DVS+ABB capability. This is achieved by using a power-composition profile aware co-synthesis methodology developed for both dynamic and leakage energy minimisation, which selects voltage scaling capabilities (DVS, ABB, or DVS+ABB) for the PEs, maps, schedules, and voltage scales applications given as task graphs with timing constraints. The co-synthesis methodology enables designers to identify suitable voltage scaling capabilities according to the tasks power-composition profiles, and is most suitable for designs where the designer has the flexibility to decide which PE should be equipped with either DVS, ABB, or DVS+ABB capability. Consideration of the ratio of the dynamic power to the leakage power in DSM designs is justified due to the dynamic power variation and increased leakage power contribution to the total power consumption. The

rest of this paper is organized as follows: Preliminaries are given in Section 2. This is followed by motivational examples in Section 3. The problem formulation and the proposed co-synthesis methodology are presented in Section 4. Experimental results and conclusion are given in Section 5 and 6 respectively.

2. Preliminaries

Power consumption of digital CMOS circuits has two major sources: dynamic power and leakage power [9]:

$$P = C_{eff} \cdot f \cdot V_{dd}^2 + L_g \cdot V_{dd} \cdot K_3 \cdot e^{K_4 V_{dd}} \cdot e^{K_5 V_{bs}} \quad (1)$$

where the first term is dynamic power and the second term is leakage power. C_{eff}, f, V_{dd} and V_{bs} denote the average switched capacitance per cycle, the circuit operational frequency, the circuit supply voltage and the body bias voltage respectively. K_3, K_4 and K_5 are technology dependent constants and L_g denotes the number of gates. In this paper, we define the ratio of the dynamic power to the leakage power at nominal V_{dd} and V_{bs} as *power-composition profile* (PCP). Equation (1) shows that, for a given processing element running at nominal V_{dd}, V_{bs} and f, leakage power is the same for all tasks, while dynamic power varies with the value of C_{eff} [3]. Thus, the PCP varies from task to task. On the other hand, for a task that can be potentially mapped to several PEs, its dynamic power, leakage power, and therefore PCP varies with the task's mapping. So PCP is a feature of a task and its particular mapping to a PE. The PCPs of the tasks can be estimated using, for example, the technique and component library outlined in [15]. DVS and ABB techniques scale V_{dd} and V_{bs}, respectively, in order to reduces the power consumption, but it is at the expense of prolonged circuit delay, which results in lower circuit operational frequency:

$$f = \frac{1}{d} = \frac{\left[(1+K_1)\cdot V_{dd} + K_2 \cdot V_{bs} - V_{th1}\right]^{\alpha}}{L_d \cdot K_6 \cdot V_{dd}} \quad (2)$$

where K_1, K_2, K_6 and V_{th1} are technology dependent constants, α is a measure of velocity saturation, L_d is the logic depth of the datapath. The energy dissipation of a task is (N_c is the cycle number needed to execute the task):

$$E_{task} = \left(\frac{N_c}{f}\right) \cdot P \quad (3)$$

In this paper we consider that an application is specified as a *task graph G(V, E)* (Fig. 1(a)). Each node, $n_i \in V$, represents a task, an atomic unit to be executed without being preempted. An edge $e_{ij} \in E$ from n_i to n_j indicates that the output of n_i is the input of n_j. A node can only be activated after all inputs have arrived. There are two nodes, *source* and *sink*, which are the first and last node respectively, so that all other nodes in the graph are successors of the source and predecessors of the sink. Release times of some tasks as well as multiple deadlines

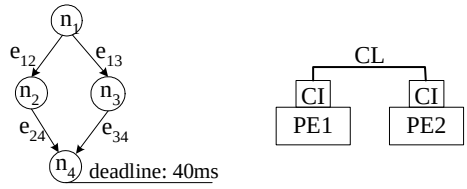

(a) task graph (b) architecture

Figure 1. Task graph and architecture example

can be easily modeled by inserting dummy nodes between certain tasks and the source or the sink respectively. An implementation is only feasible when the sink finishes earlier than the imposed deadline. The architecture considered in this paper consists of multiple PEs. The PEs are equipped with either DVS or ABB or DVS+ABB capability. An infrastructure of communication links (CLs) connects the PEs through communication interfaces (CIs). Fig. 1(b) shows an example architecture. Each task in a task graph can be potentially mapped to several PEs able to execute this task. Tasks take a finite number of clock cycles N_c to be executed and consume certain amount of dynamic power P_{dyn} and leakage power P_{leak}, depending on the PE to which they are mapped. If two tasks, n_i and n_j with an edge e_{ij} connecting them, are mapped to different PEs, a communication takes place over a CL, involving a certain amount of communication time and power.

3. Motivational examples

This section illustrates the main idea behind the proposed co-synthesis methodology by means of two motivational examples. The first example demonstrates the impact of power-composition profile (PCP) in the selection of various voltage scaling techniques (DVS, ABB, or DVS+ABB) for the PEs. The second example shows the influence of mapping on the PCP and why it is important to consider PCP during the mapping process.

3.1 Example 1

To demonstrate the impact of PCP in the selection of voltage scaling techniques for the PEs, consider a task with an execution time of 10ms at nominal supply voltage and body bias voltage and a deadline of 13ms (i.e., 3ms slack). Fig. 2 shows the corresponding energy dissipations of the task when various voltage scaling techniques are employed for two different PCPs. Consider PCP 1, where both dynamic power and leakage power are 4.67mW (i.e., dynamic power : leakage power = 1). Such power values are derived using Equation (1) with the 0.07µm Crusoe processor constants provided in [9]. As it can be seen, the energy dissipation without any voltage scaling is 93.5µJ (V_{dd} = 1V, V_{bs} = 0V), employing DVS or ABB reduce the energy dissipation to 73.8µJ (V_{dd} = 0.85V, V_{bs} = 0V) and 50.6µJ (V_{dd} = 1V, V_{bs} = –0.66V) respectively. ABB produces more energy reduction than DVS because

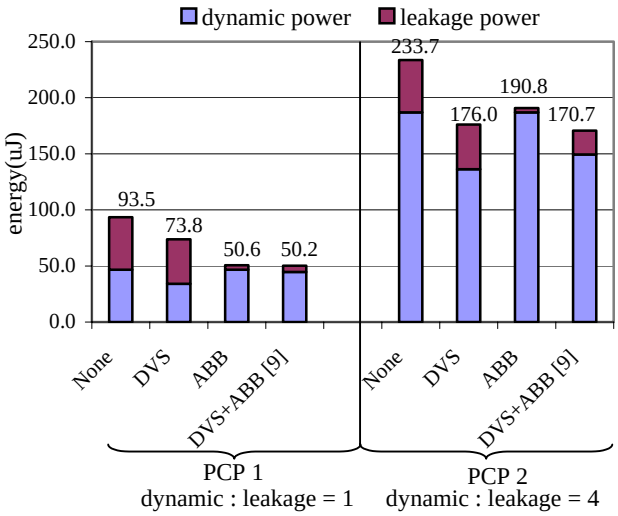

Figure 2. PCP and energy dissipation of example 1

according to Equation (1), dynamic power is quadratically dependent on supply voltage, while leakage power is exponentially dependent on body bias voltage (leakage power is also exponential to V_{dd}, but the constant K_4 for V_{dd} is smaller than K_5 for V_{bs}). Employing DVS+ABB [9] reduces the energy dissipation to 50.2µJ (V_{dd} = 0.98V, V_{bs} = –0.55V). This shows that ABB produces similar energy reduction as DVS+ABB. In this case, a suitable voltage scaling technique is ABB because of its similar energy reduction to that of DVS+ABB but less hardware cost.

Now consider PCP 2, where dynamic power and leakage power are 18.68mW and 4.67mW respectively (i.e., dynamic power : leakage power = 4). The energy dissipation without any voltage scaling is 233.7µJ (V_{dd} = 1V, V_{bs} = 0V). Employing DVS, ABB, DVS+ABB respectively reduces the energy dissipation to 176.0µJ (V_{dd} = 0.85V, V_{bs} = 0V), 190.8µJ (V_{dd} = 1V, V_{bs} = –0.66V), and 170.7µJ (V_{dd} = 0.89V, V_{bs} = –0.18V). As expected, the DVS+ABB again achieves the highest energy reduction. DVS performs slightly worse than DVS+ABB, but much better than ABB, because the dynamic power is dominating in the total power. We have a trade-off in the selection of the voltage scaling technique. One can choose DVS+ABB for lowest energy consumption at the expense of extra hardware cost, or choose DVS with slight degradation in energy reduction and reduced hardware cost. Fig. 3 shows the energy dissipation of the same task using various voltage scaling techniques for different power-composition profiles (PCP) in the range of 1 to 4, taking into account current and future CMOS technologies [5, 15]. First, consider the energy reduction of ABB and DVS+ABB. It can be seen that both techniques have comparable energy reduction in the PCP range of 1 to 2.25 (Zone I), but DVS+ABB has much higher energy reduction than ABB in Zone II and III.

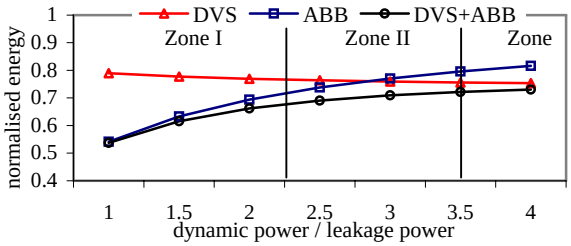

Figure 3. Energy reduction of different voltage scaling techniques for a range of PCPs

Now consider the energy reduction of DVS and DVS+ABB. As it can be seen, the two techniques have comparable energy reduction in Zone III, but DVS+ABB has much higher energy reduction than DVS in Zone I and II. Based on these comparisons, it can be observed that, for a single task to be executed on a PE, if the PCP is in Zone I, then ABB can be selected with little degradation in energy reduction but less hardware cost when compared with DVS+ABB. Similarly, if the PCP is in Zone III, then DVS can be selected. Finally, if the PCP is in Zone II, DVS+ABB should be selected because it performs much better than either DVS or ABB. Note that the zone ranges of Fig. 3 are for the case of 0.07µm Crusoe processor [9] as an example. This motivational example shows the importance of taking into account the PCP when identifying the PEs' voltage scaling capabilities in order to produce energy and cost effective designs.

3.2 Example 2

As mentioned in Section 2, a task's PCP depends on its mapping. To examine this influence, consider the same task, but this time it can be possibly mapped to two different PEs which have already been associated with certain type of voltage scaling techniques: PE_1 with DVS capability and PE_2 with ABB capability. We assume that the execution properties of the task on the two PEs are as shown in Table 1. It can be seen that, the PCP is 3 when the task is mapped to PE_1, and the PCP is 2 when the task is mapped to PE_2. Hence the two mappings generate different PCPs. Next we decide to which PE the task should be mapped with the aim of high energy saving. Based on the knowledge that the total power consumption of the task is lower on PE_1 (13.44mW) than on PE_2 (14.01mW), one may possibly map the task to PE_1. In this case, the energy dissipation without any voltage scaling is 134.4µJ (V_{dd} = 1V, V_{bs} = 0V). Applying DVS reduces the energy dissipation to 102.2µJ (V_{dd} = 0.85V, V_{bs} = 0V). However, if one maps the task to PE_2, although the energy dissipation without any voltage scaling (VS) is worse (140.1µJ), applying ABB reduces the energy dissipation to 97.3µJ (V_{dd} = 1V, V_{bs} = –0.66V), which is better than being mapped to PE_1. This is because PE_2 has ABB

Table 1. Execution properties for different mappings

	execution time (ms)	dynamic pwr (mW)	leakage pwr (mW)	total pwr (mW)	PCP	energy before VS (µJ)	energy after VS (µJ)
PE_1 (DVS)	10	10.08	3.36	13.44	3	134.4	102.2
PE_2 (ABB)	10	9.34	4.67	14.01	2	140.1	97.3

capability, which is able to efficiently reduce the energy of the task with a PCP of 2, as shown in Fig. 3. On the other hand, the DVS capability available on PE_1 is less efficient in reducing energy for the task with a PCP of 3. This motivational example shows the influence of the mapping on the PCP and the importance to consider the PCP during the mapping to achieve good energy saving.

Motivational examples have demonstrated that it is necessary to select voltage scaling techniques for PEs appropriately depending on the PCP. This selection is straightforward if there is only one task in the application, or all tasks in the application have the same PCP value. However, PCP depends on not only the technology but also the properties of the tasks and PEs. An application usually consists of multiple tasks, and the PCPs of the tasks are distributed over a wide range (e.g. from 1 to 4 as shown in Fig. 3). Considering this PCP diversity and the influence of mapping on the PCP, this paper addresses the selection of voltage scaling capability during co-synthesis to achieve energy and cost efficiency at the same time.

4. Proposed co-synthesis methodology

The problem formulation and the proposed co-synthesis methodology are presented next.

4.1 Problem formulation

The input given by the designer to the proposed co-synthesis methodology includes a task graph, an architecture (Section 2), and a technology library. Each task in the task graph can be potentially mapped to several PEs able to execute it. For each possible mapping, the number of clock cycle N_c required to execute the task, the dynamic power P_{dyn} and the leakage power P_{leak} at the nominal supply voltage and body bias voltage are given in the technology library. These values are either based on previous design experience or on estimation techniques [16]. In the architecture, the number and type of PEs have been decided, but which type of voltage scaling capability (DVS, ABB, or DVS+ABB) they should have is not yet fixed. Depending on the available voltage scaling capability, the PE can vary its supply voltage V_{dd} and/or body bias voltage V_{bs} within certain continuous ranges. There is an associated cost for each type of voltage scaling capability, which is also given in the technology library. The goal of the co-synthesis is to decide which type of voltage scaling capability should be associated with each PE, and find mapping, scheduling, supply voltage and body bias voltage assignment for the tasks, such that the imposed deadline is met, and at the same time, the total energy dissipation and hardware cost are reduced. An assumption is made that the tasks are of sufficiently coarse granularity and that the PEs can continue operation during the voltage scaling, which allows neglecting the scaling overhead in terms of power and time [17]. If these overheads cannot be neglected, techniques such as those presented in [13] can be used to take the overhead issues into consideration during the optimisation process.

4.2 Power-composition profile driven co-synthesis

A co-synthesis methodology addressing both dynamic power and leakage power for embedded systems is presented. Taking into account the power-composition profile (PCP), the methodology selects suitable voltage scaling capabilities (DVS, ABB, or DVS+ABB) for the PEs and decides the mapping, scheduling, supply voltage and body bias voltage assignments for the tasks. There are two optimisation objectives in this problem: energy dissipation and hardware cost. Our methodology provides the designer the trade-off between these two objectives.

The methodology is based on two nested optimisation loops, as shown in Fig. 4. The inner loop is a power-composition profile aware mapping (Section 4.3). It is responsible for the optimisation of the design (mapping, scheduling and voltage assignments) for a given architecture and a given selection of voltage scaling capabilities, which is identified in the outer loop. The outer loop iteratively identifies a selection of voltage scaling capabilities for the PEs. It starts from a selection with the highest hardware cost (i.e., all PEs are associated with DVS+ABB capability) and incrementally moves towards selections with lower hardware costs. In each iteration, the outer loop (1) identifies a selection of voltage scaling capabilities; (2) passes the identified selection to the inner loop to generate the corresponding design; (3) checks whether there exists another selection with equal or lower hardware cost; (4) if such a selection exists, then an actual selection is identified, e.g., associate a PE with separate DVS or ABB instead of DVS+ABB, and the newly identified selection is passed to the inner loop. This procedure repeats until there is no selection with equal or lower hardware cost. Exhaustive search is employed in this outer loop to identify all possible selections. Finally, all the solutions (including the selections of voltage scaling capabilities and the corresponding designs) are evaluated in terms of two characteristics: energy dissipation and hardware cost. We

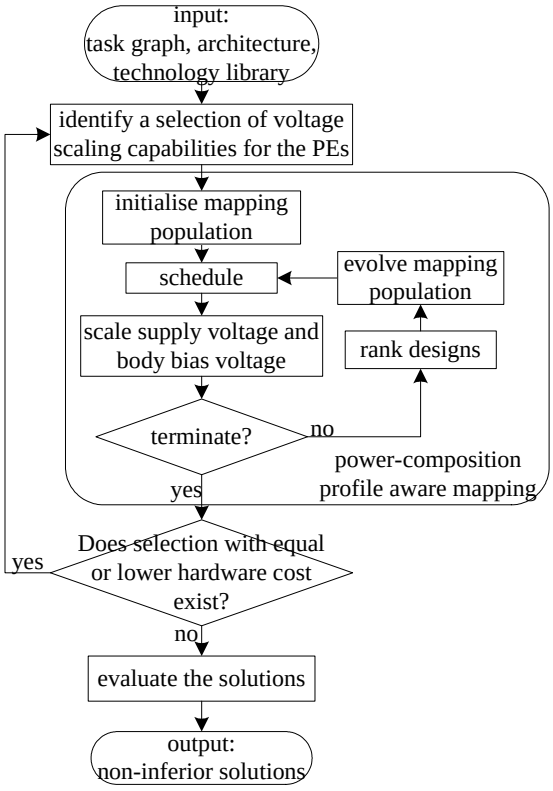

Figure 4. Proposed PCP aware co-synthesis

define a solution A is *inferior* if there exists another solution B which meets one of the following conditions: (1) both characteristics of B are better than A, or (2) one characteristic of B is the same as for A, while the other characteristic of B is better than A (there are two possibilities meeting this condition: equal energy dissipation but better hardware cost, or equal hardware cost but better energy dissipation). Inferior solutions are discarded and the remaining non-inferior solutions, i.e., the Pareto set, are presented to the designer.

4.3 Power-composition profile aware mapping

The key characteristic of the proposed co-synthesis methodology is the power-composition profile (PCP) aware mapping (see the inner loop of Fig. 4). Taking advantage of the available PCP information, it optimises the design (mapping, scheduling, supply voltage and body bias voltage assignment) in terms of energy dissipation and timing feasibility for a given selection of voltage scaling capabilities, which is identified in the outer loop.

As shown in the motivational example, it is essential to consider PCP during the mapping process to achieve good energy saving at low hardware cost. The basic idea of the PCP aware mapping is to, for each task, look at the PCPs resulted from all possible mappings, then decide which mapping is preferable. The complexity of this process is *O*(number of tasks * number of PEs). Denote a specific task as τ_i, a specific PE as PE_j, the PCP of τ_i when it is mapped to PE_j as PCP_{ij}. According to Fig. 3, there are three cases that τ_i should be preferably mapped to PE_j: (1) PCP_{ij} is in Zone I (Fig. 3) and PE_j has ABB capability or DVS+ABB capability; (2) PCP_{ij} is in Zone II and PE_j has DVS+ABB capability; (3) PCP_{ij} is in Zone III and PE_j has DVS capability or DVS+ABB capability. In practice, there will be a set of dependent tasks in an application. Ideally all tasks need to be mapped according to the PCP. However this is not always possible due to the imposed time constraints, area constraints, or lack of PEs with suitable voltage scaling capabilities. As will be shown later in this section, to meet timing feasibility, some optimised task mappings may not be in accordance with their PCP, although this increases energy dissipation.

The PCP aware mapping is based on a genetic algorithm (GA). For further details concerning the application of genetic algorithm in task mapping the reader is referred to [18]. As shown in Fig. 4, the first step of the PCP aware mapping is mapping initialisation. In this step an initial population (i.e., a set of mappings) is generated taking into consideration the PCP. A detailed description of the mapping initialisation is given in Fig. 5 and explained as follows. According to the previous discussion of this section, there are three cases that τ_i should be preferably mapped to PE_j depending on PCP_{ij} (the PCP of τ_i when it is mapped to PE_j). Correspondingly, steps 01 – 10 identify the PEs preferable for τ_i to be mapped to, and include them into $PE\text{-}SET_i$. Therefore, $PE\text{-}SET_i$ includes the PEs that have suitable voltage scaling capabilities to efficiently reduce the energy dissipation of τ_i. If there is no PE preferable for τ_i, then τ_i should be mapped to any PEs able to execute τ_i (Steps 11 – 12). In steps 14 – 15, $PE\text{-}SET_i$ is used to generate a proportion (50% is found to work practically well) of the initial mapping population. In step 16, the remaining proportion of initial mapping population is generated randomly in order to increase the diversity of the initial population. Although some of the initialised mappings might be infeasible in terms of timing behaviour, this initialisation has been found to improve the optimisation procedure significantly by introducing individuals that are likely to evolve into high quality solutions.

As shown in Fig. 4, after mapping initialisation, the initial mapping population is passed to the scheduling, the supply voltage and body bias voltage scaling to produce the designs. We adopt a critical path list scheduling [19] and a supply voltage and body bias voltage scaling similar to the technique in [3]. The technique in [3] addresses DVS only. To enable it to address both DVS and ABB, a modification is made, which basically involves the distribution of slack time that simultaneously allows the

```
01 for (each task τ_i) {
02   for (each PE_j able to execute τ_i) {
03     PCP_ij = P_dyn(i,j) / P_leak(i,j)
04     if ((PCP_ij ∈ Zone I) and (PE_j has ABB or DVS+ABB))
05       PE-SET_i = PE-SET_i + PE_j
06     if ((PCP_ij ∈ Zone II) and (PE_j has DVS+ABB))
07       PE-SET_i = PE-SET_i + PE_j
08     if ((PCP_ij ∈ Zone III) and (PE_j has DVS or
                 DVS+ABB))
09       PE-SET_i = PE-SET_i + PE_j
10   }
11   if (PE-SET_i is empty )
12     PE-SET_i = all PEs able to execute τ_i
13 }
14 generate an initial mapping: for each τ_i, choose a PE ∈ PE-
     SET_i randomly
15 repeat step 14 to generate a proportion of the initial mappings
16 generate the remaining proportion of the initial mappings
     randomly
```

Figure 5. PCP aware mapping initialisation

adjustment of V_{dd} and V_{bs}. After voltage scaling, the designs are evaluated and ranked, the mapping individuals with high fitness (i.e., leading to low energy dissipation and timing feasibility) are selected to evolve a new population by mating and mutating them. This evaluation – ranking – evolvement procedure continues to optimise the mappings until the termination criteria (i.e., no improved individual has been produced for a certain number of generations) is reached. The evaluation of mappings is guided by the following fitness function:

$$Fitness = \left(\sum_i E(n_i)\right) \cdot \left(\frac{\max(T_d, T_e)}{T_d}\right)^2 \qquad (4)$$

where $E(n_i)$ is the energy dissipation of task n_i, T_d is the deadline of the task graph, and T_e is the actual delay of the task graph. The first part of the fitness function is the total energy dissipation of all tasks. The second part introduces a penalty factor due to deadline violations. Thus the evolvement is driven towards solutions with reduced energy dissipation, while, at the same time, the real-time constraints are satisfied. In the final optimised mappings, it is possible that some tasks' mappings are not in accordance with the PCP. This is because the optimisation algorithm changes the initial mappings during the evolvement in order to meet timing feasibility, at the expense of increased energy dissipation.

5. Experimental results

The proposed power-composition profile (PCP) aware co-synthesis methodology has been implemented on a Pentium III 866/256MB PC running Linux. In order to evaluate its capability of reducing energy dissipation at reduced hardware cost, a set of experiments has been carried out on 5 automatically generated examples (tg1 – tg5) containing 25-69 nodes and 29-84 edges. The target architecture contains 2-3 PEs, which have technology constants of the 0.07μm Crusoe processor given in [9]. The PCPs of the tasks resulted from possible mappings range from 1 to 4. The deadline of all the examples is 125 – 130% of the minimum delay. The supply voltage of the DVS and body bias voltage of the ABB can be scaled continuously within 0.5 – 1V and –1.0 – 0V respectively. To estimate the hardware cost, we assume that the cost of a PE with no voltage scaling capability is 1, the cost of a PE with DVS or ABB capability is 1.1, and a PE with DVS+ABB capability is 1.2. Such costs have been chosen mainly to give an indication of the cost difference between DVS, ABB and DVS+ABB.

Using the proposed co-synthesis methodology, the results for example tg3 (48 nodes, 60 edges, 3 PEs) are given in Fig. 6. Each numbered point stands for a solution featured by two characteristics: energy dissipation (relative to the energy dissipation without voltage scaling) and hardware cost, where '×' denote inferior solutions and 'o' denote non-inferior solutions. As mentioned in Section 4.2, a solution A is *inferior* if there exists another solution B which meets one of the following conditions: (1) both characteristics of B are better than A, or (2) one characteristic of B is same as A, while the other characteristic of B is better than A. For example, consider the solutions with hardware cost of 3.3 (i.e., all PEs have separate DVS or ABB capability). Solutions 1 – 3 are inferior and solution 4 is non-inferior. Similarly, solutions 5, 6 and 8 are inferior, while solutions 7, 9 and 10 are non-inferior. The proposed co-synthesis discards the inferior solutions and provides the non-inferior solutions for the designer. In this example, solutions 4, 7, 9 and 10 are provided for the designer to give a trade-off between energy dissipation and hardware cost.

Similarly, the co-synthesis methodology is applied to

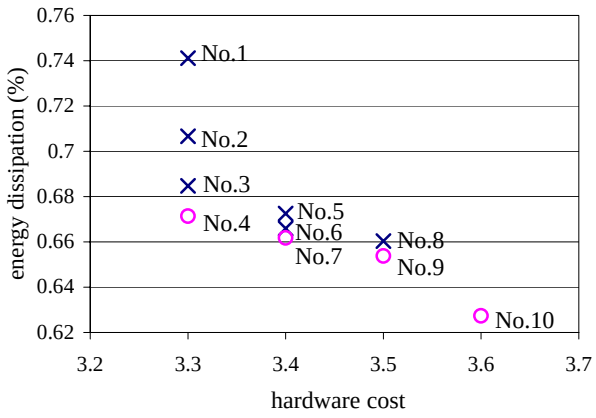

Figure 6. Co-synthesis results for tg3 example ('×' inferior solution, 'o' non-inferior solution)

Table 2. Co-synthesis results of various selection of voltage scaling capabilities

exa-mple	number of PE			hardware cost	energy (%)
	DVS	ABB	DVS+ ABB		
tg1	0	0	2	2.4	61.5
	1	0	1	2.3	64.3
	1	1	0	2.2	69.1
tg2	0	0	3	3.6	63.3
	0	1	2	3.5	63.5
	2	0	1	3.4	65.3
	2	1	0	3.3	68.2
tg3	0	0	3	3.6	62.7
	0	1	2	3.5	65.4
	1	1	1	3.4	66.2
	2	1	0	3.3	67.1
tg4	0	0	3	3.6	65.0
	1	0	2	3.5	66.1
	1	1	1	3.4	67.8
	2	1	0	3.3	69.6
tg5	0	0	3	3.6	62.1
	1	0	2	3.5	65.4
	1	1	1	3.4	66.1
	2	1	0	3.3	69.4

tg1, tg2, tg4 and tg5. The non-inferior solutions for tg1 – tg5 are given in Table 2, where columns 2 – 4 give the numbers of PEs with DVS capability, PEs with ABB capability, and PEs with DVS+ABB capability respectively. Employing the PEs indicated in columns 2 – 4, columns 5 and 6 give the corresponding hardware cost and the achieved energy dissipation (relative to the energy dissipation without voltage scaling) of the examples. As expected, for all the examples, using PEs with DVS+ABB capability achieves the best energy saving. However, by carefully selecting separate DVS or ABB capability for the PEs, comparable energy saving can be achieved with reduced hardware cost. Consider tg3, in the case when all the three PEs have DVS+ABB capability (first row), the energy dissipation is 62.7%. The energy dissipation is increased to 66.2%, when the three PEs have different voltage scaling capabilities (third row). This 3.5% of increase in energy dissipation has been achieved through the reduction of hardware cost from 3.6 to 3.4. Furthermore, as it can be seen, the hardware cost can be reduced further to 3.3 (fourth row) at the expense of a slight increase in energy dissipation (0.9%). Similarly, the non-inferior results for tg1, tg2, tg4 and tg5 also provide trade-offs between energy dissipation and hardware cost. Based on these results, a designer can choose a suitable selection of voltage scaling capabilities for the PEs according to the design constraints of the application. The execution time of the examples ranges from 2 hours (tg1) to up to 21 hours (tg5) due to the exhaustive search nature as outlined in Section 4.2.

The results of Table 2 have been produced assuming the supply voltage of the DVS and body bias voltage of the ABB can be scaled continuously within 0.5 – 1V and –1.0 – 0V respectively. We have applied our co-synthesis methodology when supply voltage and body bias voltage can only have a known number of fixed values. For each voltage scaling technique, we assume five different voltage settings, so that there is an equal increment of operational frequency between two neighbouring voltage settings. For example, for combined DVS+ABB, the five voltage settings are (1, 0), (1, –0.23), (1, –0.47), (0.96, –0.55) and (0.91, –0.56). From the results of continuous voltage scaling, Table 2, we can derive the results with discrete voltage scaling, using the technique described in [2, 13]. Table 3 gives the results for tg5. For example, using two PEs with DVS capability and one PE with ABB capability (last row) reduces the energy dissipation of tg5 to 73.1% in the case of discrete voltage scaling (Table 3), whilst the reduction is 69.4% in the case of continuous voltage scaling (Table 2).

To further validate the proposed co-synthesis methodology, we have applied it to a real life GSM voice CODEC [17]. This example has a task graph of 87 nodes and 139 edges, and an architecture of 3 PEs. Table 4 shows the non-inferior solutions and some inferior solutions generated using the co-synthesis methodology. As can be seen, there is a trade-off between energy dissipation and hardware cost. Using three PEs with DVS+ABB capability (first row of non-inferior solutions), the hardware cost is 3.6 and the energy dissipation is 64.3% (relative to the energy dissipation without voltage scaling), while using three PEs with DVS capability (last row of non-inferior solutions), the hardware cost and energy dissipation is 3.3 and 71.3% respectively. This

Table 3. Co-synthesis results of discrete voltage scaling

exa-mple	number of PE			hardware cost	energy (%)
	DVS	ABB	DVS+ ABB		
tg5	0	0	3	3.6	68.6
	1	0	2	3.5	71.3
	1	1	1	3.4	71.6
	2	1	0	3.3	73.1

Table 4. Co-synthesis results for CODEC example

	number of PE			HW cost	energy (%)
	DVS	ABB	DVS+A BB		
non-inferior solution	0	0	3	3.6	64.3
	1	0	2	3.5	67.1
	1	1	1	3.4	68.2
	3	0	0	3.3	71.3
inferior solution	2	1	0	3.3	73.9
	1	2	0	3.3	75.9
	0	3	0	3.3	78.3

shows that, in some situations, comparable energy saving can be achieved without using PEs with DVS+ABB capability. However, a careful identification of voltage scaling capability should be proceeded to achieve the desired energy saving. For example, comparing the non-inferior and inferior solutions with the hardware cost of 3.3, the energy dissipation of using three PEs with ABB capability (last row of inferior solutions) is 78.3%, much higher than that of using three PEs with DVS capability (71.3%). For this example, DVS capability should be selected for the three PEs. This is because, according to the PCPs, most tasks in this example should preferably be mapped to PEs with DVS capability.

6. Conclusions

We have proposed a new co-synthesis methodology for reducing dynamic and leakage energy in multi-processor embedded systems. Unlike previous approaches, the presented co-synthesis methodology takes into account the power-composition profile of the tasks, optimises designs not only towards energy reduction but, additionally hardware cost reduction. The key contribution includes the voltage scaling technique selection and task mapping strategy that consider the influence of power-composition profile, and the co-synthesis methodology that provides trade-offs between energy dissipation and hardware cost. We have validated our methodology using several experiments including a real-life GSM voice CODEC example. These experiments have demonstrated that, depending on the power-composition profiles, it is possible to achieve significant energy reduction without the employment of PEs with combined DVS+ABB capability, i.e., reduced hardware cost. The presented co-synthesis methodology is most suitable for designs where the designer has the flexibility to decide which processing element should be equipped with either DVS, ABB, or DVS+ABB capability.

7. References

[1] D. Duarte, N. Vijaykrishnan, M. J. Irwin, H.-S. Kim and G. McFarland, "Impact of scaling on the effectiveness of dynamic power reduction schemes," in *Proc. ICCD, 2002*, pp.382-7, Germany.

[2] T. Ishihara and H. Yasuura, "Voltage scheduling problem for dynamically variable voltage processors," in *Proc. ISLPED, 1998*, pp.197-202, USA.

[3] M. T. Schmitz and B. M. Al-Hashimi, "Considering power variations of DVS processing elements for energy minimisation in distributed systems," in *Proc. ISSS, 2001*, pp.250-5, Canada.

[4] Y. Zhang, X. Hu and D. Z. Chen, "Task scheduling and voltage selection for energy minimization," in *Proc. DAC, 2002*, pp.183-8, USA.

[5] K. S. Khouri and N. K. Jha, "Leakage power analysis and reduction during behavioral synthesis," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol.10, no.6, pp.876-85, Dec. 2002.

[6] T. Kuroda, T. Fujita, S. Mita, T. Nagamatsu, S. Yoshioka, K. Suzuki, F. Sano, M. Norishima, M. Murota, M. Kako, M. Kinugawa, M. Kakumu and T. Sakurai, "A 0.9-V, 150-MHz, 10-mW, 4 mm^2, 2-D discrete cosine transform core processor with variable threshold-voltage (VT) scheme," *IEEE Journal of Solid-State Circuits*, vol.31, no.11, pp.1770-9, Nov. 1996.

[7] M. Miyazaki, G. Ono and K. Ishibashi, "A 1.2-GIPS/W microprocessor using speed-adaptive threshold-voltage CMOS with forward bias," *IEEE Journal of Solid-State Circuits*, vol.37, no.2, pp.210-17, Feb. 2002.

[8] K. Nose, M. Hirabayashi, H. Kawaguchi, S. Lee and T. Sakurai, "V_{TH}-hopping scheme to reduce subthreshold leakage for low-power processors," *IEEE Journal of Solid-State Circuits*, vol.37, no.3, pp.413-9, Mar. 2002.

[9] S. M. Martin, K. Flautner, T. Mudge and D. Blaauw, "Combined dynamic voltage scaling and adaptive body biasing for lower power microprocessors under dynamic workloads," in *Proc. ICCAD, 2002*, pp.721-5, USA.

[10] C. H. Kim and K. Roy, "Dynamic V_{TH} scaling scheme for active leakage power reduction," in *Proc. DATE, 2002*, pp.163-7, France.

[11] J. T. Kao, M. Miyazaki and A. P. Chandrakasan, "A 175-MV multiply-accumulate unit using an adaptive supply voltage and body bias architecture," *IEEE Journal of Solid-State Circuits*, vol.37, no.11, pp.1545-54, Nov. 2002.

[12] L. Yan, J. Luo and N. K. Jha, "Combined Dynamic Voltage Scaling and Adaptive Body Biasing for Heterogeneous Distributed Real-time Embedded Systems," in *Proc. ICCAD, 2003*, pp.30-37, United States.

[13] A. Andrei, M. T. Schmitz, P. Eles, Z. Peng and B. M. Al-Hashimi, "Overhead-conscious voltage selection for dynamic and leakage energy rduction of time-constrained systems," in *Proc. DATE, 2004*, pp.518-523, France.

[14] T. D. Burd and R. W. Brodersen, *Energy efficient microprocessor design*, ISBN 0-7923-7586-6, Kluwer academic publishers, 2002.

[15] D. Duarte, Y.-F. Tsai, N. Vijaykrishnan and M. J. Irwin, "Evaluating run-time techniques for leakage power reduction," in *Proc. 15th International Conference on VLSI Design, 2002*, pp.31-8, India.

[16] V. Tiwari, S. Malik and A. Wolfe, "Power analysis of embedded software: a first step towards software power minimization," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol.2, no.4, pp.437-45, Dec. 1994.

[17] M. T. Schmitz, B. M. Al Hashimi and P. Eles, *System-level design techniques for energy-efficient embedded systems*, ISBN 1-4020-7750-5, Kluwer Academic Publishers, 2004.

[18] D. Wu, B. M. Al-Hashimi and P. Eles, "Scheduling and mapping of conditional task graphs for the synthesis of low power embedded systems," in *Proc. DATE, 2003*, pp.90-5, Germany.

[19] P. Eles, A. Doboli, P. Pop and Z. Peng, "Scheduling with bus access optimization for distributed embedded systems," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol.8, no.5, pp.472-91, Oct. 2000.