

# Cosynthesis of Energy-Efficient Multimode Embedded Systems With Consideration of Mode-Execution Probabilities

Marcus T. Schmitz, Bashir M. Al-Hashimi, *Senior Member, IEEE*, and Petru Eles, *Member, IEEE*

**Abstract**—In this paper, we present a novel co-design methodology for the synthesis of energy-efficient embedded systems. In particular, we concentrate on distributed embedded systems that accommodate several different applications within a single device, i.e., multimode embedded systems. Based on the key observation that operational modes are executed with different probabilities, that is, the system spends uneven amounts of time in the different modes, we develop a new co-design technique that exploits this property to significantly reduce energy dissipation. Energy and cost savings are achieved through a suitable synthesis process that yields better hardware-resource-sharing opportunities. We conduct several experiments, including a realistic smart phone example, that demonstrate the effectiveness of our approach. Reductions in power consumption of up to 64% are reported.

**Index Terms**—Embedded systems, energy efficiency, multimode systems, power minimization, system-level cosynthesis.

## I. INTRODUCTION

OVER the last several years, the popularity of portable applications has explosively increased. Millions of people use battery-powered mobile phones, digital cameras, MP3 players, and personal digital assistants (PDAs). To perform major parts of the system's functionality, these mass products rely, to a great extent, on sophisticated embedded computing systems with *high performance* and *low-power dissipation*. One key characteristic of many current and emerging embedded systems is their need to work across a set of different interacting applications and operational modes. For instance, modern mobile phones often contain not solely the functionality required for communication purpose (e.g., voice coding and protocol handling), but additionally integrate applications like digital cameras, games, and complex multimedia functions (MP3 players and video decoders) into the same single device. Throughout this article, such embedded systems are referred to as *multimode embedded systems*. This paper introduces a novel cosynthesis methodology particular suitable for the design of energy-efficient multimode embedded systems. Starting

Manuscript received May 8, 2003; revised October 14, 2003. This work was supported in part by the UK Engineering and Physical Sciences Research Council (EPSRC) under Grant GR/S95770/01. This paper was recommended by Associate Editor M. F. Jacome.

M. T. Schmitz and P. Eles are with the Department of Computer and Information Science, Linköping University, S-58333 Linköping, Sweden (e-mail: g-marsc@ida.liu.se; petel@ida.liu.se).

B. M. Al-Hashimi is with the School of Electronics and Computer Science, University of Southampton, Southampton SO17 1BJ, U.K. (e-mail: bmah@ecs.soton.ac.uk).

Digital Object Identifier 10.1109/TCAD.2004.837729




Fig. 1. Typical activation profile of a cellular phone.

from a specification model that captures both mode interaction and functionality, the developed cosynthesis technique maps the application under consideration of *mode-execution probabilities* to a heterogeneous, distributed architecture with the aim to reduce the energy consumption through an appropriate resource sharing between tasks. Mode-execution probabilities refer to the activation time of operational modes that are user typical. Consider, for instance, the typical activation profile of a mobile phone, which is shown in Fig. 1. According to this profile, the phone stays most of the time in a radio link control (RLC) mode, in order to maintain network connectivity. While the network search and calling modes are only active for small periods of the overall time. The main principle by which the proposed cosynthesis process achieves energy-efficiency is an implementation tradeoff between the different operational modes. In general, modes with high execution probability should be implemented more energy efficient (e.g., by moving more tasks to hardware) than modes with a low execution probability. Nonetheless, the implementation of modes is heavily interrelated, due to the fact that different modes share the same resources (architecture). For example, mapping an energy-critical task of a highly active mode into energy-efficient hardware might prohibit to implement a timing-critical task into fast hardware due to the restricted area (see motivational example in Section IV). Clearly, a well-balanced implementation of the operational modes is vital for a good system design.

In addition, the cosynthesis approach further reduces the energy dissipation by adapting the system performance to the particular needs of the active mode, using dynamic voltage scaling (DVS) as well as component shutdown. That is, instead of wasting energy through overperformance, the computational power is adapted according to the individual performance requirements of each mode and each task. Furthermore, we introduce a transformation-based method to extend existing

DVS approaches in order to allow the scaling of hardware-processing elements that are capable of executing tasks in parallel, however, which rely on a single scalable supply voltage source.

The remainder of this paper is organized as follows. Section II introduces relevant previous work. Preliminaries, outlining a new multimode specification and an architectural model, are given in Section III. Motivational examples exemplify the need for a suitable multimode synthesis approach in Section IV. The problem at hand is formulated in Section V. Section VI describes our multimode cosynthesis approach, and Section VII presents experimental results. Finally, in Section VIII we draw some conclusions.

## II. PREVIOUS WORK

In the last decade, numerous methodologies for the design of low power-consuming embedded systems have been proposed, including approaches that leverage power management techniques, such as dynamic power management and DVS. Nevertheless, a crucial feature of many modern embedded systems is their capability to execute several different applications (multimodes), which are integrated into a single device.

Approaches for the schedulability analysis of systems with several modes of operations can be found in the real-time research community [24], [29]. However, these approaches solely concentrate on scheduling aspects (i.e., they investigate if the mode change events fulfill the imposed timing constraints) and do not address implementation aspects. Three recent approaches have addressed various problems involved in the design of multimode embedded systems [19], [23], [30]. Shin *et al.* [30] proposed a schedulability-driven performance analysis technique for real-time multimode systems. They show that it is possible, through a sophisticated performance estimation, to identify timing-critical tasks, which are active in different operational modes. This identification allows an improvement of the execution times of the most crucial tasks, in order to achieve system schedulability. In their work, the optimization of the identified tasks is up to the designer. For example, reductions in the execution times can be made by handcrafted code tuning and outsourcing of core routines into hardware. Kalavade and Subrahmanyam [19] have introduced a hardware/software-partitioning approach for systems that perform multiple functions. Their technique classifies tasks, found within similar applications, into groups of task types. The implementation of frequently appearing task types is biased toward hardware. This can be intuitively justified by the fact that costly hardware implementations are shared across a set of applications, hence, exploiting the allocated hardware more cost effective. Oh and Ha [23] address the problem in a slightly different way. Their cosynthesis framework for multimode systems is based on a combined scheduling and mapping technique for heterogeneous multiprocessor systems (HMP [22]). Taking a processor-utilization criterion into account, an allocation-controller selects the required processing elements (PEs) such that the schedulability constraint is satisfied and the system cost is minimized. The main principle behind all three approaches is to consider the possibility of resource sharing, i.e., computational tasks of the same type, which can

be found in different modes, utilize the same implementations. Thereby, multiple hardware implementations of the same task type are avoided, which, in turn, reduces the hardware cost. Other noticeable approaches are the works by Chung *et al.* [12] and Yang *et al.* [32]. In [12], energy efficiency is achieved by leveraging information regarding the execution-time variations, which is supplied to the mobile terminal by the contents provider. That is, the performance of the mobile terminal can be influenced directly by the contents provider, in accordance to the processing requirements of the sent content. The approach presented in [32] uses a two-phase scheduling method. In the first stage, which is performed offline (during design time), a Pareto-optimal set of schedules is generated. These schedule provide different execution time/energy tradeoffs. During runtime, a runtime scheduler selects points along the Pareto set, in order to account for the dynamic behavior of the application. As opposed to these approaches, the work presented in this paper addresses the design of low energy consuming multimode systems that exhibit variations in the mode activation profile; hence, it differs in several aspects from the previous works. To the authors' knowledge, there has been no prior work investigating the co-design problem of energy minimization taking into account mode-execution probabilities. This paper makes the following contributions.

- 1) The consideration of mode-execution probabilities and their effect on the energy-efficiency of multimode embedded systems is analyzed and demonstrated.
- 2) A co-design methodology for the design of energy-efficient multimode systems is presented. The proposed cosynthesis maps and schedules a system specification that captures both mode interaction and mode functionality onto a distributed heterogeneous architecture. Four mutation strategies are introduced that aid the GA-based optimization process in finding solutions of high quality by pushing the search into promising design-space regions.
- 3) DVS is investigated in the context of multimode embedded systems. A transformation-based approach is used to tackle the problem of DVS on PEs that execute different tasks in parallel, but offer only a single scalable supply voltage source.

## III. PRELIMINARIES

This section introduces the functional specification model (Section III-A) and the architectural model (Section III-B), which are fundamental to the proposed cosynthesis framework.

### A. Functional Specification of Multimode Systems

The abstract specification model used for multimode embedded systems consists of two parts. In précis, it is based on a combination of finite state machines and task graphs, capturing both the interaction between different operational modes as well as the functionality of each individual mode. Structurally, each node in the finite state machine represents an operational mode and further contains the task graphs which are active during this mode. The following two sections introduce this model, which is henceforth referred to as operational mode

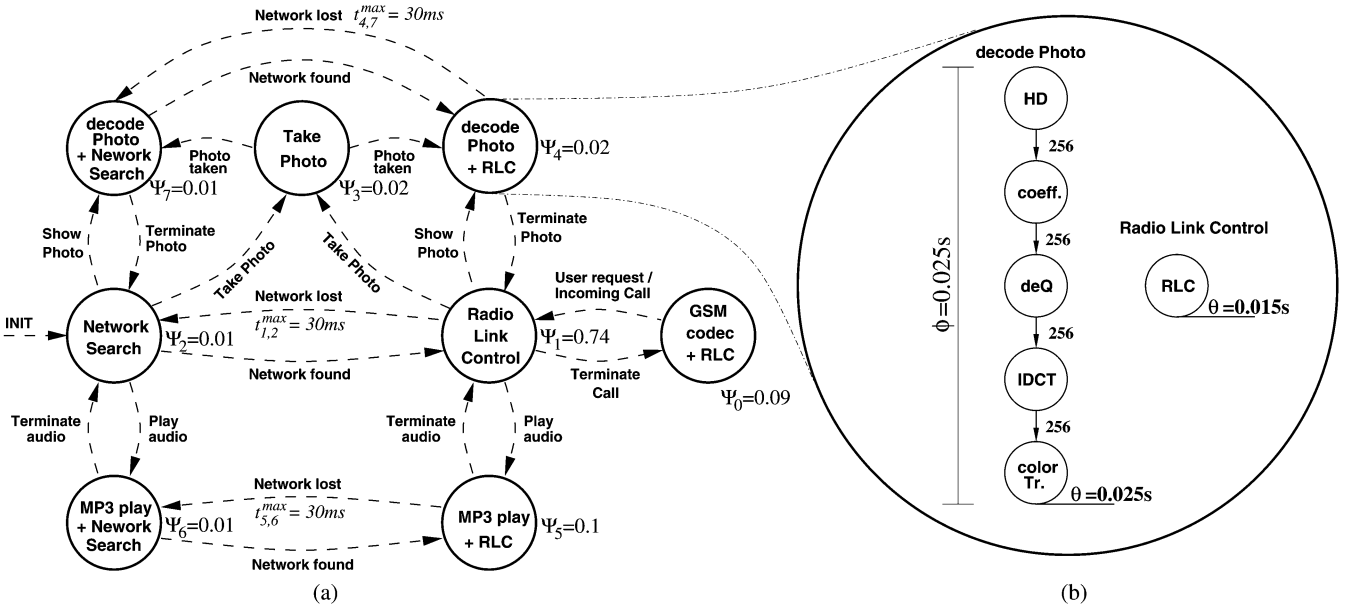


Fig. 2. Relation between OMSM and individual task graph specifications. (a) Operational models. (b) Task graph of a single mode.

state machine (OMSM). A similar abstract model was mentioned in [16]. However, here, this model is extended toward system-level design and includes transition time limits as well as mode-execution probabilities. The following introduces this model, using the smart phone example shown in Fig. 2.

1) *Top-Level Finite State Machine*: In this work, it is considered that an application is given as a directed cyclic graph  $\Upsilon(\Omega, \Theta)$ , which represents a finite state machine. Within this top-level model, each node  $O \in \Omega$  refers to an operational mode and each edge  $T \in \Theta$  specifies a possible transition between two different modes. If the system undergoes a change from mode  $O_x$  to mode  $O_y$ , where  $x \neq y$ , the transition time  $t_T^{max}$  associated with the transition edge  $T = (O_x, O_y)$  has to be met. For instance, as indicated in Fig. 2(a), upon losing the network connection the system needs to activate the network search mode within 30 ms. Such transition overheads can originate from the reconfiguration of field programmable gate arrays (FPGAs) as well as from loading the application software of the particular mode into the local PE's memory. At any given time, there is only one active mode, i.e., the modes execute mutually exclusive. To exemplify the proposed model, consider Fig. 2(a). This figure shows the OMSM for a smart phone example with eight different modes. A possible activation scenario could look like this. When switched on, the phone initializes into network search mode. The system stays in this mode until a suitable network has been found. Upon finding a network the phone undergoes a mode change to RLC. In this mode, it maintains the connection to the network by handling cell handovers, radio link failure responses, and adaptive RF power control. An incoming phone call necessitates to switch the system into GSM codec + RLC mode. This mode is responsible for speech encoding and decoding, while simultaneously maintaining network connectivity. Similarly, the remaining modes have different functionalities and are activated upon mode change events. Such events originate upon user requests (e.g., MP3-player activation) or are initiated by the system itself (e.g., loss of network connection

necessitates to switching the system into network search mode). Furthermore, based on the key observation that many multimode systems spend their operational time *unevenly* in each of the modes, an execution probability  $\Psi_O$  is associated with each operational mode  $O$ , i.e., it is known what percentage of the operational time the device spends in each mode. For instance, in accordance with Fig. 2(a), the smart-phone stays 74% of this operational time in RLC mode, 9% in GSM codec + RLC mode, and 1% in network search mode. The remaining 16% of the operation time are associated with the remaining modes. In practice the mode probabilities vary from user to user, depending on the personal usage behavior. Nevertheless, it is possible to derive an average activation profile based on statistical information collected from several different users. Taking this information into account will prove to be important when designing systems with a prolonged battery lifetime. It is interesting to note that different operational modes do not necessarily correspond to different functionalities of the system. It is possible to use alternative modes to model the same functionality under different working conditions (such as different workloads). For instance, in order to account for variations in the wireless channel quality, we could exchange the GSM voice transcoder mode  $O_0$  in Fig. 2(a) with three transcoder schemes, each responsible for the coding at a specific signal-to-interference ratio (SIR) on the channel. During runtime, the appropriate transcoder scheme would be selectively activated, depending on the actual channel quality.

2) *Functional Specification of Individual Modes*: The functional specification of each operational mode  $O \in \Omega$  in the top-level finite state machine is expressed by a task graph  $G_S^O(\mathcal{T}, \mathcal{C})$ . This relation is shown in Fig. 2. Each node  $\tau \in \mathcal{T}_O$  in a task graph represents a task, i.e., a fragment of functionality that needs to be executed without preemption. The level of granularity is coarse, i.e., tasks refer to functions such as Huffman decoder, dequantizer, FFT, IDCT, etc. Therefore, each task is further associated with a task type

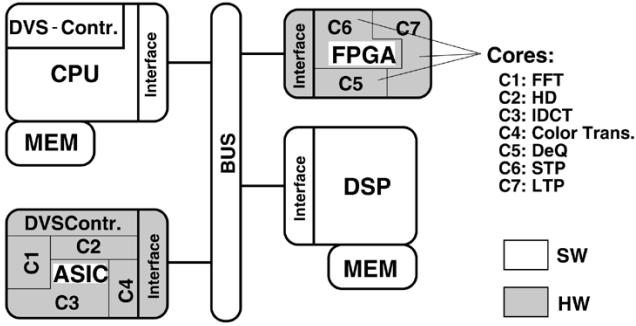


Fig. 3. Distributed architecture model.

$\eta \in \Gamma = \{\text{HD, deQ, FFT, IDCT, } \dots\}$ . A distinctive feature of multimode systems is that task type sets  $\Gamma^O \subseteq \Gamma$  of different modes  $O \in \Omega$  can intersect, i.e., tasks of the same type are executed in different modes. Such modes can share the same hardware resource (intermode sharing). Resource sharing is also possible for multiple tasks of identical type that are found in a single mode (intramode sharing); however, due to task communalities among different modes, the chances to share resources are increased. Further, tasks might be annotated with deadlines  $\theta_\tau$  (with  $\tau \in \mathcal{T}_O$ ) by which the execution has to be finished, in order to guarantee correct functioning. Similarly, the whole task graph has to be successively repeated according to a period  $\phi_O$ . Edges  $\gamma \in \mathcal{C}$  in the task graph refer to precedence constraints and data dependencies between the computational tasks, i.e., if two tasks,  $\tau_i$  and  $\tau_j$ , are connected via an edge, then task  $\tau_i$  must be finished and transfer data to task  $\tau_j$  before  $\tau_j$  can be executed. A feasible implementation of a certain mode  $O$  needs to respect all task deadlines  $\theta$ , task graph period  $\phi$ , and precedence relations.

### B. Architectural Model and System Implementation

The proposed system-level synthesis approach targets distributed architectures that possibly consist of several heterogeneous PEs, such as general-purpose processors (GPPs), application-specific instruction set processors (ASIPs), application-specific integrated circuits (ASICs), and FPGAs. These components are connected through an infrastructure of communication links (CLs). A directed graph  $G_A((\mathcal{P} \cup \mathcal{L}), \mathcal{E})$  captures such an architecture, where nodes  $\pi \in (\mathcal{P} \cup \mathcal{L})$  denote PEs and CLs, while edges  $\lambda \in \mathcal{E}$  impose the connections between those components. Fig. 3 shows an architecture example. Since each task might have multiple implementation alternatives, it can be potentially mapped onto several different PEs that are capable of performing this type of task. Tasks mapped to software-programmable components (i.e., GPP or ASIP) are placed into local memory. However, if a task is mapped to a hardware component (i.e., ASIC or FPGA), a core for this task type needs to be allocated. A feasible solution needs to obey the imposed area constraints, i.e., only a restricted number of cores can be implemented on hardware components. The subdivision of hardware components (ASICs and FPGAs) into hardware cores is shown in Fig. 3. Each core is capable of performing a single task of type  $\eta \in \Gamma$  at a time. Tasks assigned to GPPs or ASIPs (software tasks) need to be sequenced, whilst the tasks mapped onto FPGAs and ASICs (hardware tasks) can be performed in par-

allel if the necessary resources (cores) are not already engaged. However, contention between two or more tasks assigned to the same hardware core requires a sequential execution order, similar to software tasks. Cores implemented on FPGAs can be dynamically reconfigured during a mode change, involving a time overhead, which needs to respect the imposed maximal mode transition times  $t_T^{\max}$ . Further, PE's might feature DVS to enable a tradeoff between power consumption and performance that can be exploited during runtime. The relation between the dynamic power dissipation  $P_{\text{dyn}}$  and the circuit delay  $d$  (inverse proportional to performance) can be expressed using the following two equations [8], [11]:

$$P_{\text{dyn}} = C_{\text{eff}} \cdot V_{\text{dd}}^2 \cdot f$$

$$d \propto 1/f = k \cdot \frac{V_{\text{dd}}}{(V_{\text{dd}} - V_t)^\alpha}$$

where  $C_{\text{eff}}$  is the effectively switched capacitance,  $V_{\text{dd}}$  denotes the circuit supply voltage,  $f$  represents the clock frequency,  $k$  and  $\alpha$  are a circuit dependent constants, and  $V_t$  denotes the threshold voltage. As we can see from these equations, by varying the circuit supply voltage  $V_{\text{dd}}$ , it is possible to trade off between power consumption and performance. In reality, DVS processors are often restricted to run at discrete voltage levels [5], [6]. Therefore, a set  $\mathcal{V}_\pi$  specifies the available discrete voltages of DVS-PE  $\pi$ . For such PEs, a voltage schedule needs to be derived, in addition to a timing schedule. To implement a multimode application captured as OMSM, the tasks and communications of all operational modes need to be mapped onto the architecture, and a valid schedule for these activities  $\epsilon \in \mathcal{A}$ , where  $\mathcal{A} = \mathcal{T} \cup \mathcal{C}$ , needs to be constructed. As mentioned above, for tasks mapped to DVS-enabled components an energy reducing voltage schedule has to be determined. According to these aspects, an implementation candidate can be expressed through four functions, which need to be derived for each operational mode  $O \in \Omega$ :

**Task mapping:**  $M_\tau^O : \mathcal{T} \rightarrow \mathcal{P}$

**Communication mapping:**  $M_\gamma^O : \mathcal{C} \rightarrow \mathcal{L}$

**Timing schedule:**  $S_\epsilon^O : \mathcal{A} \rightarrow \mathbb{R}^+$

**Voltage schedule:**  $V_\tau^O : \mathcal{T}_{\text{DVS}} \rightarrow \mathcal{V}_\pi$

where  $M_\tau^O$  and  $M_\gamma^O$  denote task and communication mapping, respectively, assigning tasks to PEs and communications to CLs. Activity start times are specified by the scheduling function  $S_\epsilon^O$ , while  $V_\tau^O$  defines the voltage schedule for all tasks  $\tau \in \mathcal{T}_{\text{DVS}}$  mapped to DVS-PEs, where  $\mathcal{V}_\pi$  is the set of the possible discrete supply voltages of PE  $\pi$ . Clearly, the mappings as well as the corresponding schedules are defined for every mode separately, i.e., during the change from mode  $O_x$  to mode  $O_y$ , the execution of activities found in mode  $O_x$  are finished and the activities of mode  $O_y$  are activated.

## IV. MOTIVATIONAL EXAMPLES

The aim of this section is to motivate the key ideas behind the new multimode cosynthesis, that is, the consideration of mode-execution probabilities and multiple task-type implementations. First, the influence of mapping in the context of multimode embedded systems with different mode-execution probabilities is demonstrated. Second, it is illustrated that multiple

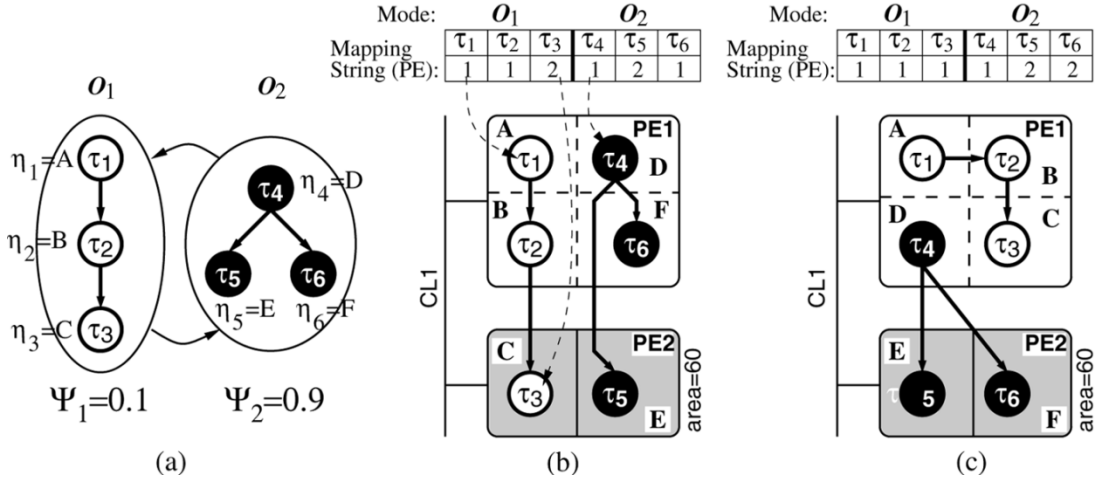


Fig. 4. Example of mode-execution probabilities. (a) Application specified by two interacting modes. (b) Optimized without mode consideration. (c) Optimized with mode consideration.

TABLE I  
TASK EXECUTION AND IMPLEMENTATION PROPERTIES

task type	PE1 (SW)		PE2 (HW)		area (mm <sup>2</sup> )
	execution time (ms)	dynamic energy (mJ)	execution time (ms)	dynamic energy (mJ)	
A	20	10	2	0.010	24.0
B	28	14	2.2	0.012	30.0
C	32	16	1.6	0.023	27.5
D	26	13	3.1	0.047	24.5
E	30	15	1.8	0.015	21.0
F	24	14	2.2	0.032	28.0

task implementations can further reduce the energy dissipation of multimode embedded systems.

1) *Example: Mode-Execution Probabilities:* For simplicity, timing and communication issues are neglected in the following example. Consider the application shown in Fig. 4(a), which consists of two operational modes  $O_1$  and  $O_2$  each specified by a task graph with three tasks. The system spends 10% of its operational time in mode  $O_1$  and the remaining 90% in mode  $O_2$ , i.e., the execution probabilities are given by  $\Psi_1 = 0.1$  and  $\Psi_2 = 0.9$ . The specification needs to be mapped onto a target architecture built of one general-purpose processor (PE1) and one ASIC (PE2), linked by a bus (CL1). Depending on the task mapping to either of the components, the execution properties of each task are shown in Table I. In general, hardware implementations of tasks achieve a higher performance and are more energy efficient [9]. It can be observed that all tasks are of different types, therefore, if a task is mapped to HW, a suitable core needs to be allocated explicitly for that task. Hence, in this particular example, no hardware sharing is considered. Each allocated core uses area on the hardware component that offers 60 mm<sup>2</sup>, i.e., at most two cores can be allocated at the same time without violating the area constraint (see Table I, Column 6). Note that although the two modes execute mutually exclusive, the task types implemented in hardware (HW cores) cannot be changed during runtime, since their implementation is static (nonreconfigurable ASIC); as opposed to software-programmable components. Consider the mapping shown in Fig. 4(b) in which

the highest energy-consuming tasks ( $\tau_3$  and  $\tau_5$ , when implemented in software) are executed using a more energy-efficient hardware implementation. According to the task characteristics given in Table I, the energy dissipation during modes  $O_1$  and  $O_2$  are  $E_1 = 10 \text{ mJ} + 14 \text{ mJ} + 0.023 \text{ mJ} = 24.023 \text{ mJ}$  and  $E_2 = 13 \text{ mJ} + 0.015 \text{ mJ} + 14 \text{ mJ} = 27.015 \text{ mJ}$ . Neglecting the mode-execution probabilities by assuming that both modes are active for even amounts of time (50% mode  $O_1$  and 50% mode  $O_2$ ), the energy consumption can be calculated as  $E_e = 0.5 \cdot 24.023 \text{ mJ} + 0.5 \cdot 27.015 \text{ mJ} = 25.519 \text{ mJ}$ . Nevertheless, taking the real behavior into account, mode  $O_1$  is active for 10% of the operational time, i.e., its energy dissipation can then be calculated as  $E_{r1} = 0.1 \cdot 24.023 \text{ mJ} = 2.4023 \text{ mJ}$ . Similarly, mode  $O_2$  is active 90% of the operational time, hence, its energy is given by  $E_{r2} = 0.9 \cdot 27.015 \text{ mJ} = 24.3135 \text{ mJ}$ . Thus, the real energy dissipation results in  $E_r = E_{r1} + E_{r2} = 26.7158 \text{ mJ}$ . Now, consider an alternative mapping, shown in Fig. 4(c), for the same task graphs. In this configuration tasks  $\tau_5$  and  $\tau_6$ , i.e., the most energy dissipating tasks of the highly active mode  $O_2$ , use energy-efficient hardware implementations on PE2, while task  $\tau_3$  of the less active model  $O_1$  is shifted into the software-programmable processor (PE1). According to this solution, the energy consumptions of modes  $O_1$  and  $O_2$  are given by  $E_1 = 10 \text{ mJ} + 14 \text{ mJ} + 16 \text{ mJ} = 40 \text{ mJ}$  and  $E_2 = 13 \text{ mJ} + 0.015 \text{ mJ} + 0.032 \text{ mJ} = 13.047 \text{ mJ}$ . Considering the even execution of each mode (neglecting the execution probabilities), the energy consumption can be calculated as  $0.5 \cdot 40 \text{ mJ} + 0.5 \cdot 13.047 \text{ mJ} = 26.524 \text{ mJ}$ . Note that this value is higher than the corresponding energy of the first mapping ( $E_e = 25.519 \text{ mJ}$ ). Thus, a cosynthesis approach that neglects the mode-execution probabilities would optimize the system toward the first mapping. However, in real-life the modes are active for different amount of time and hence the real energy dissipation is given by  $E_r = 0.1 \cdot 40 \text{ mJ} + 0.9 \cdot 13.047 \text{ mJ} = 15.7423 \text{ mJ}$ . This is 41% lower compared to the first mapping ( $E_r = 26.7158 \text{ mJ}$ ) shown in Fig. 4(b), which is not optimized for an uneven task execution probability. Furthermore, the second task mapping allows to switch off PE2 and CL1 during mode  $O_1$ , since all tasks of this mode are assigned to PE1. This results in a

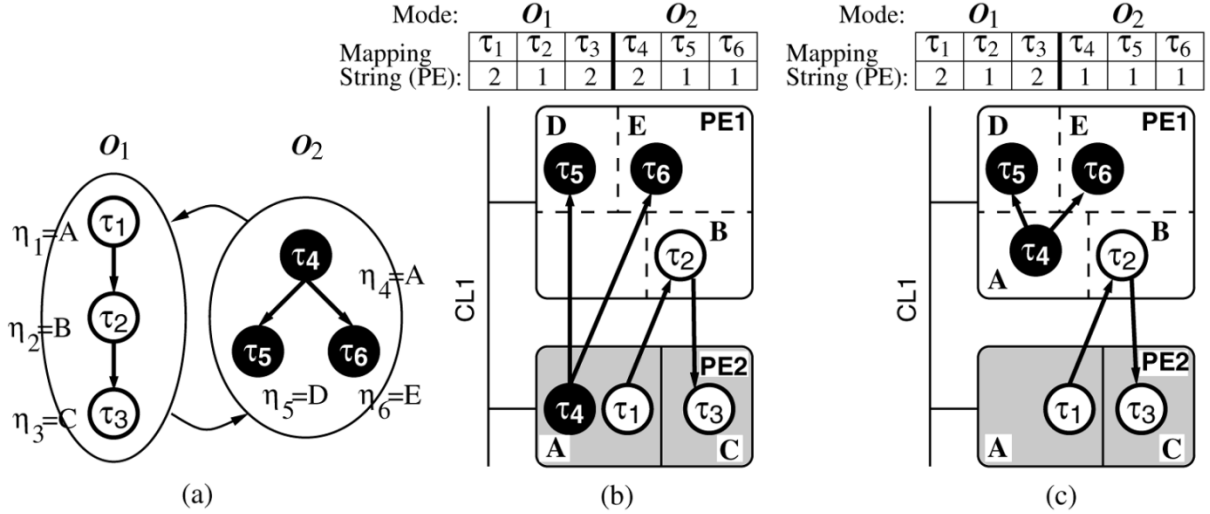


Fig. 5. Multiple task type implementations. (a) Application with resource sharing possibility. (b) Resource sharing, but no shutdown possible. (c) No resource sharing, but component shutdown.

significant reduction of the static power, additionally increasing the energy savings.

2) *Example: Multiple Task-Type Implementations:* An important characteristic of multimode systems is that tasks of the same type might be found in different modes, i.e., resources can be shared among the different modes in a time-multiplexed fashion. To increase the possibility of component shutdown, it might be necessary to implement the same task type multiple times, however, on different components. The following example, shown in Fig. 5, clarifies this aspect. Here tasks  $\tau_1$  and  $\tau_4$  are of type A [see Fig. 5(a)], allowing resource sharing between these tasks. The sharing is possible without contention due to the mutual exclusive execution of these tasks (only one mode is active at a given time). In the first mapping, given in Fig. 5(b), both tasks utilize the same HW core. However, implementing task  $\tau_4$  in software (additional task type A on PE1), as shown in Fig. 5(c), allows to shut down PE2 and CL1 during the execution of mode  $O_2$ . Hence, multiple implementations of task types can help to reduce power dissipation.

These two examples have demonstrated that it is essential to guide the synthesis process by: 1) an energy model that takes into account the mode-execution probability as well as 2) allowing multiple task implementations.

## V. PROBLEM FORMULATION

The goal of our cosynthesis approach is an energy-efficient implementation of application  $\Upsilon$ , which is modeled as OMSM, such that timing and area constraints are satisfied. This involves the derivation of the mapping and schedule functions,  $M_\tau^O$ ,  $M_\gamma^O$ ,  $S_\epsilon^O$ , and  $V_\tau^O$  (outlined in Section III-B), under the consideration of static and dynamic power as well as mode-execution probabilities. Although static power consumption is often neglected in system-level design approaches, since until recently dynamic power has been the dominating power dissipation, emerging submicron technologies with reduced threshold voltage levels show increased leakage currents that are becoming comparable to the dynamic currents [10]. In multimode systems, this static power consumption can have a

significant impact on the overall energy efficiency. The reasons for this are the different performance requirements of the various operational modes. For instance, the minimal performance requirements of the hardware architecture are imposed by the most computational intensive mode, i.e., the minimal allocated architecture has to provide enough computational power to execute this performance critical mode. However, the allocated architecture might be far more powerful than actually needed for the execution of modes with low performance requirements. Furthermore, low-performance modes, such as the standby-mode of mobile phones (i.e.,  $RC$ ), often account for the greatest portion of the system time. During such circumstances, the static energy dissipation of unnecessarily switched-on PEs and CLs can outweigh the dynamic energy consumption caused by tasks of a “lightweight” mode. Thus, switching-off the unneeded components becomes an important aspect particularly in multimode embedded systems. In accordance, an accurate estimation of the average power consumption of an implementation alternative should consider both static and dynamic power, and further the mode-execution probabilities. The average power consumption  $\bar{p}$  can be expressed using the following equation:

$$\bar{p} = \sum_{O \in \Omega} (P_O^{\text{stat}} + P_O^{\text{dyn}}) \cdot \Psi_O \quad (1)$$

where  $P_O^{\text{stat}}$ ,  $P_O^{\text{dyn}}$ , and  $\Psi_O$  refer to the static power dissipation, the dynamic power dissipation, and the execution probability of mode  $O$ , respectively. The static and dynamic power consumptions are given as

$$P_O^{\text{stat}} = \sum_{\xi \in \mathcal{K}_O} P^{\text{stat}}(\xi) \quad (2)$$

and

$$P_O^{\text{dyn}} = \left( \sum_{\epsilon \in \mathcal{A}_O} E^{\text{dyn}}(\epsilon) \right) \cdot \frac{1}{\text{hp}_O} \quad (3)$$

where  $P^{\text{stat}}(\xi)$  refers to the static power consumption of a component  $\xi$ , which is found in the set of all active components  $\mathcal{K}_O \subseteq (\mathcal{P} \cup \mathcal{L})$  of mode  $O$ . Please note that this static power consumption also includes the additional power required for the dc/dc converter of voltage-scalable processors. Further,  $\mathcal{A}_O$  and  $\text{hp}_O$  denote all activities and the hyper-period of mode  $O$ , respectively. With respect to the type of activities, the dynamic energy consumption  $E^{\text{dyn}}(\epsilon)$  can be calculated as

$$E^{\text{dyn}}(\epsilon) = \begin{cases} P_{\max}(\epsilon) \cdot t_{\min}(\epsilon) \cdot \frac{V_{\text{dd}}^2(\epsilon)}{V_{\max}^2(\epsilon)} & \text{if } \epsilon \in \mathcal{T}_{\text{DVS}} \\ P_{\max}(\epsilon) \cdot t_{\min}(\epsilon) & \text{if } \epsilon \in \mathcal{T} \setminus \mathcal{T}_{\text{DVS}} \\ P_C(\epsilon) \cdot t_C(\epsilon) & \text{if } \epsilon \in \mathcal{C} \end{cases} \quad (4)$$

where  $P_{\max}$  is the dynamic power consumption and  $t_{\min}$  is the execution time of tasks when executed at nominal supply voltage  $V_{\max}$ . Tasks  $\tau \in \mathcal{T}_{\text{DVS}}$  mapped to DVS-PEs can execute at a scaled supply voltage  $V_{\text{dd}}$ , resulting in reduced energy consumption. Further, communications consume power  $P_C$  over a time  $t_C$ . If the DVS-enabled processors are restricted to a limited set of discrete voltages, the continuous selected supply voltage  $V_{\text{dd}}$  is split into its two neighboring discrete voltages  $V_{\text{low}}$  and  $V_{\text{high}}$ . The corresponding execution times in each voltage are calculated as given in [25]. The mode-execution probabilities used in (1) are either based on approximations or statistical information collected from several real users. In the case that statistical information is available from a set of different users  $U$ , the average execution probabilities  $\Psi_O$  of a single operational mode  $O \in \Omega$  can be calculated.

The cosynthesis goal is to find a task mapping  $M_\tau^O$ , a communication mapping  $M_\gamma^O$ , a starting time schedule  $S_\epsilon^O$ , as well as a voltage schedule  $V_\tau^O$  for each operational mode  $O$ , such that the total average power  $\bar{p}$ , given in (1), is minimized and the deadlines are satisfied. Furthermore, a feasible implementation candidate needs to fulfill the following requirements.

- 1) The mapping of tasks  $M_\tau^O$  does not violate area constraints in terms of memory and hardware area, i.e.,  $(\sum_{\eta \in \Gamma_\pi} a_\eta) \leq a_\pi^{\max}, \forall \pi \in \mathcal{P}$ , where  $\Gamma_\pi$  is the set of all task types implemented on PE  $\pi$ , and  $a_\eta$  and  $a_\pi^{\max}$  refer to the area used by task type  $\eta$  and the available area on PE  $\pi$ , respectively. Please note that for DVS-enabled HW,  $a_\pi^{\max}$  represents the available area including the area overhead required for the dc/dc converter.
- 2) The timing schedule  $S_\epsilon^O$  and the voltage schedule  $V_\tau^O$ , based on task and communication mapping, do not exceed any task deadlines  $\theta_\tau$  or task graph repetition periods  $\phi_O$ , therefore,  $t_S(\tau) + t_{\text{exe}}(\tau) \leq \min(\theta_\tau, \phi), \forall \tau \in \mathcal{T}$ , where  $t_S(\tau)$  and  $t_{\text{exe}}$  refer to task start time and task execution time (potentially based on voltage scaling).
- 3) The system reconfiguration time  $t_T$  between mode changes does not exceed the imposed maximal mode transition times  $t_T^{\max}$ . Hence,  $t_T \leq t_T^{\max}, \forall T \in \Theta$  needs to be respected for all mode transitions.

## VI. COSYNTHESIS OF ENERGY-EFFICIENT MULTIMODE SYSTEMS

Fig. 6 shows an overview of our cosynthesis flow for multimode embedded systems. This design flow is primarily based

on two nested optimization loops. The outer loop optimizes task mapping and core allocation, while the inner loop is responsible for the combined optimization of communication mapping and scheduling. Although we concentrate on task mapping and core allocation in this paper, we will briefly outline this overall design flow. As we can observe from Fig. 6, an initial system specification has to be translated into the final hardware and software implementations. In our approach, the specification includes information regarding the execution probabilities. Along the design flow we can identify five major synthesis steps.

- 1) An adequate target architecture needs to be allocated, i.e., it is necessary to determine the quantity and the types of the different interconnected components (PEs and communication links). Available components are specified in the technology library.
- 2) The tasks of the system specification and the required communications have to be uniquely mapped among the allocated components. Based on the component to which a task has been mapped, its execution properties are determined. This is done in accordance to previously established execution estimations and profile information. Furthermore, hardware cores are allocated based on the available hardware area and the application parallelism.
- 3) Depending on the task mapping, the communications are mapped onto the allocated communication links, and the activities are scheduled with the aim to meet the imposed task deadlines, while, at the same time, achieve a good slack distribution in order to reduce energy via DVS.
- 4) DVS and component shutdown possibilities are exploited to reduced the system energy consumption.
- 5) The system implementation candidates, specified by the synthesis Steps 1–4, are evaluated in terms of power consumption, performance (deadline satisfaction), and cost (architecture and hardware area). This step is used to provide feedback to the previous synthesis steps in order to refine the design.

For more information we refer the interested reader to [25]–[27]. A major goal of this design flow is to support the system designer with a methodology that aids to find suitable target architectures for a given system specification.

In this paper, we concentrate on the task mapping and core allocation step, since the scheduling and communication mapping can be carried with standard single-mode techniques (e.g., [17] and [26]). This is due to the fact that the modes are executing mutually exclusive. Thus, here we present new techniques and algorithms for task mapping, hardware core allocation, and DVS that suit the particular problems of multimode embedded systems. However, since these approaches are targeted toward the exploitation of mode-execution probabilities, we first discuss how such probabilities can be obtained in practice.

### A. Estimation of Mode-Execution Probabilities

As we have demonstrated in the motivational example of Section IV, the consideration of mode-execution probabilities during design time can help to significantly reduce the energy consumption of the embedded system. Certainly, to achieve a good design, it is necessary that the execution probabilities

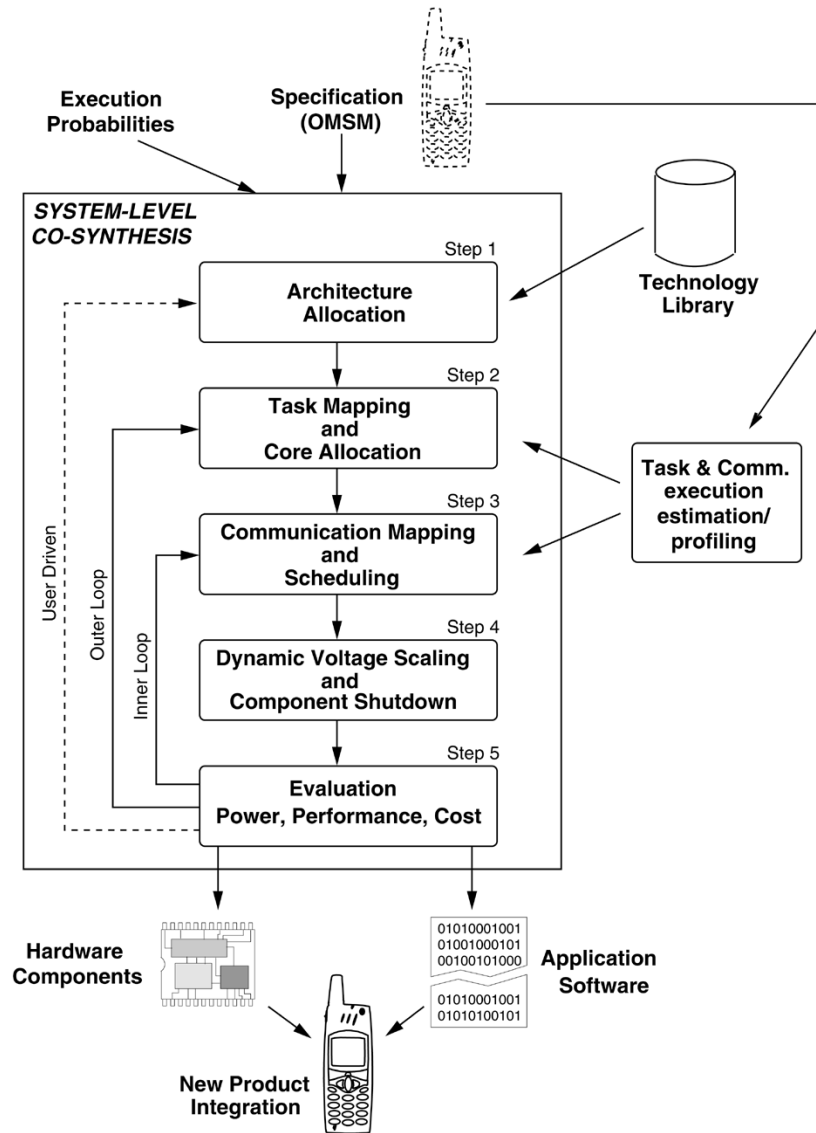


Fig. 6. Multimode embedded systems design flow.

(estimations) used during design time reflect the real usage probabilities (in-field) accurately. In the following, we outline how to obtain adequate execution probabilities using two different design scenarios.

- 1) *The new design is an upgrade of an existing product which is connected to a service provider* (e.g., a new version of a mobile phone). For such product types, it is possible to use information regarding the activation profile that has been collected on the provider side during the operation of the previous product generation. For instance, the cellular network-base stations can record the activation profile of the mobile terminals (e.g., phones with RLC and calling mode) directly in-field. This information could then be evaluated and used during the design of the new product.
- 2) *The product is a completely new design.* In this situation, it is common practice to evaluate the market acceptance before the final product is introduced, using a limited number of prototypes that are distributed among a set of

evaluation users. During this evaluation phase, the prototypes can gather information regarding the activation profile. This information could then be used during the final design of the product to optimize the energy consumption. Of course, it is also possible to use application-specific insight of the designer to estimate the execution probabilities. As we will show in the experiments given in Section VII-A, even if the estimated execution probabilities do not reflect the user activation with absolute accuracy, but are sufficiently close to the real values, energy savings can be still achieved.

### B. Multimode Cosynthesis Algorithm

The task-mapping approach, which simultaneously determines  $M_\tau$  for all modes of application  $\Upsilon$ , is driven by a genetic algorithm (GA). GAs optimize a population of individuals over several generations by imitating and applying the principles of natural selection. That is, the GA iteratively



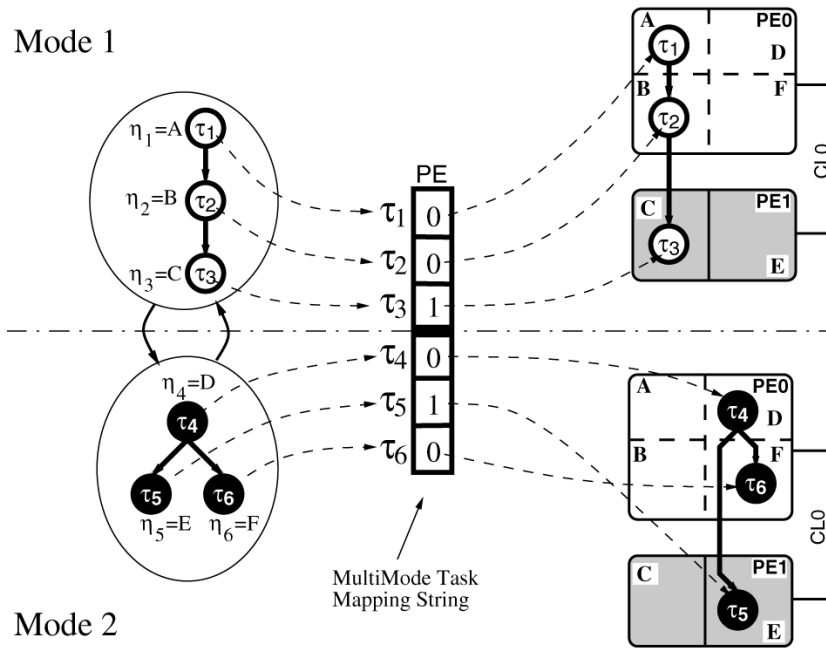


Fig. 7. Task mapping string for multimode systems.

evolves new populations by mating (crossover) the fittest individuals (highest quality) of the current population pool until a certain convergence criterion is met. In addition to mating, mutation, i.e., the random change of genes in the genome (string), provides the opportunity to push the optimization into unexplored search space regions. GA-driven task-mapping approaches have already been shown to provide a powerful tool in deriving mappings for single mode systems [15], [25]. Here, we enhance such approaches toward multimode aspects. These enhancements include the consideration of resource sharing, component shutdown, and mode transition issues. As opposed to the single-mode task-mapping strings, such strings for multimode specifications combine the mapping strings of each operational mode into one larger task-mapping string, as shown in Fig. 7. Within this string, each number represents the PE to which the corresponding task is assigned. This encoding enables the usage of a GA to optimize the placement of tasks across the PEs that form the distributed architecture. Please note that this representation supports the implementation of multiple task types. For instance, if two tasks of the same type are mapped onto different PEs, these tasks are implemented on both PEs. Thereby, the possibility of multiple task implementations is mainly inherited into the genetic mapping algorithm which is guided by a cost function that accounts for the multiple task implementations, i.e., the GA trades off between the savings in static power consumption and the increased dynamic power.

The goal of the cosynthesis is to find a mapping of tasks that minimizes the total power consumption and obeys the performance constraints. Fig. 8 outlines the pseudocode of our cosynthesis algorithm. Starting from an initial random population of multimode task mapping strings (line 1), the optimization runs until the convergency criterion is met (line 2). The used criterion is based on the diversity in the current population and the number of elapsed iterations without producing any improved individual. To judge the quality of mapping candidates, i.e., the

Algorithm: **MULTI-MODE-SYN**

---

Input: - OMSM (FSM+TGs), technology lib,  
allocated architecture  
Output (Outer loop): - Core allocation  
- Task mapping  
Output (Inner loop): - Communication mapping  
- Scheduling  
- Scaled supply voltages

---

```

01: Pop ← InitRandPop
02: while (NotConv(Pop)) { //outer loop
03:   for all map ∈ Pop {
04:     mob ← CompMobil(map) //ASAP&ALAP based
05:     cores ← ImplementHWcores(map, mob)
06:     ap ← CalcAreaPenalty(map, cores)
07:     Pstat.PE ← CalcStaticPowPE
08:     trp ← CalcTransPenalty(cores)
09:     for all mode ∈ Ω { //inner loop
10:       CommMapping.Scheduling(mode)
11:       tp(mode) ← CalcTimingPenalty(mode)
12:       Pdyn(mode) ← CalcDynPow(mode) //+ DVS
13:       Pstat.CL ← CalcStaticPowCL
14:     } // end for
15:     FM = MappingFitness(Pdyn, Pstat.CL, Pstat.PE,
                           tp, ap, trp)
16:   } // end for
17:   RankIndividuals(FM, Pop)
18:   mat ← SelMatIndividuals(Pop)
19:   TwoPointCrossover(mat)
20:   OffspringInsertion(Pop)
21:   ShutdownImproveMut(Pop)
22:   AreaImproveMut(Pop)
23:   TimingImproveMut(Pop)
24:   TransImprovemut(Pop)
25: } // end while

```

Fig. 8. Pseudocode: Multimode cosynthesis.

fitness which guides the GA, it is necessary to estimate important design objectives, including static and dynamic power dissipation, area usage, and timing behavior (lines 3–14). The fol-

lowing explains each of the required estimations. The hardware area depends on the allocated cores on each hardware component (ASIC or FPGA). Of course, for each task type mapped to hardware at least one core of this type needs to be allocated. However, if too many cores are placed onto a single ASIC or FPGA, the available area is exceeded and an area penalty is introduced (line 6). On the other hand, if multiple tasks of the same type are mapped to the same hardware component and the hardware area is not violated, it is possible to implement cores multiple times (if helpful for the energy reduction). In the proposed approach, additional cores (line 5) are allocated for parallel tasks with low mobility (line 4); therefore, the chance to exploit application parallelism is increased. The mobility of a task is the difference between its earliest and latest possible start times [31]. Clearly, from an energy point-of-view, this is also preferable, especially in the presence of DVS, where a decreased execution time results in more slack that can be exploited. Section VI-C describes the core allocation in more detail. At this point it is possible to compute the static power consumption of the implementation (line 7), taking into account component shutdown between different modes. Components can be shut down during the execution of a certain mode whenever no tasks belonging to that mode are mapped onto these components, i.e., the component is vacant (for instance, PE0 during execution of mode  $O_2$  in Fig. 5). Another important aspect is the reconfigurability of FPGAs, which allows to exchange the implemented cores to suit the active mode. However, this reconfiguration during a mode change takes time. Hence, a transition penalty is introduced if the maximal transition times are exceeded (line 8). Having determined the cores to be implemented (line 5), it is now also possible to schedule each mode of the application and to derive a feasible communication mapping (line 10). Since the modes are mutually exclusive, we can employ a communication mapping and scheduling optimization for a single-mode system. In our approach, we utilize the technique described in [25] for this step. If timing constraints are violated by the found schedule, a timing penalty is introduced (line 11). Furthermore, based on the communication mapping and scheduling, the dynamic power consumption of the application can be computed, taking into account DVS (line 12) if voltage-scalable components are present. Similar to the shutdown of PEs, it is also possible to switch off a CL when no communications are mapped to that link (line 13), therefore, further reducing the static power consumption of the system. Based upon all estimated power consumptions and penalties, a fitness is calculated (line 15) as

$$F_M = \bar{p} \cdot \text{tp} \cdot \left( 1 + w_A \cdot \sum_{\pi \in \mathcal{P}_v} \frac{a_\pi^U - a_\pi^{\max}}{a_\pi^{\max} \cdot 0.01} \right) \cdot w_R \cdot \prod_{T \in \Theta_v} \frac{t_T}{t_T^{\max}} \quad (5)$$

where the average power dissipation  $\bar{p}$  is given by (1) and  $\text{tp}$  introduces a timing penalty if the schedule exceeds task deadlines or the repetition period. Further, an area penalty is applied for all PEs with area violation  $\mathcal{P}_v$  by relating used area  $a_\pi^U$  and area constraint  $a_\pi^{\max}$ . Similarly, a transition-time penalty is applied for all transitions  $\Theta_v$  that exceed their maximal transition time limit  $t_T^{\max}$ . Both area and transition penalties are weighted ( $w_A$  and  $w_R$ ), which allows to adjust the aggressiveness of the penalties. Having assigned a fitness to all individuals of the pop-

ulation, they are ranked using linear scaling (line 17). A tournament selection scheme is used to pick individuals (line 18) for mating (line 19). The produced offsprings are inserted into the population (line 20).

In order to improve the performance of the GA, we apply four genetic mutation strategies that add problem-specific knowledge into the optimization process (lines 21–24). This is achieved by introducing a small number of mutated individuals into the current population, whenever the optimization process occurs to be trapped. These newly injected solution candidates provide the potential to turn into a high-quality solution by mating with another solution. The mutation strategies are introduced next.

1) *Shutdown Improvement*: To increase the chances of component shutdown, which leads to a reduction of static power consumption, the genetic task-mapping algorithm employs a simple yet effective strategy during the optimization. Out of the current population, randomly picked individuals (probability 2% was found to lead to good results) are modified as follows. A single mode  $O_x$  and a nonessential PE  $\pi_a$  are selected. Nonessential PEs are considered to be PEs that implement task types that have alternative implementations on other PEs, hence, they are not fundamental for a feasible solution. Our goal is to switch off PE  $\pi_a$  during the execution of mode  $O_x$ . Therefore, all tasks of mode  $O_x$ , which are mapped to  $\pi_a$ , are randomly remapped to the remaining PEs ( $\mathcal{P} \setminus \pi_a$ ). Hence, PE  $\pi_a$  can be shut down during mode  $O_x$ . Of course, only feasible mappings are allowed, i.e., tasks are always mapped randomly to the PEs that are capable of executing this kind of task type.

2) *Area Improvement*: To avoid convergence toward area-infeasible solutions, a second strategy is employed. If only infeasible-area mappings have been produced for a certain number of generations, the search is pushed away from this region by randomly remapping hardware tasks onto software-programmable PEs.

3) *Timing Improvement*: In contrast to the area-improvement strategy, if a certain amount of timing infeasible solutions have been produced, software tasks are randomly mapped to faster hardware implementations. Thereby, increasing the chance of finding timing feasible implementations.

4) *Transition Improvement*: Cores implemented in FPGAs can be dynamically reconfigured. However, this involves a time overhead. If this overhead exceeds the imposed transition time limits, the mapping is infeasible. Hence, after generating for a certain number of generations solely solutions that violate the transition times, tasks are randomly remapped away from the FPGAs that cause the violations.

Although some of the produced genomes (strings) might be infeasible in terms of area and timing behavior, all these strategies have been found to improve the search process significantly by introducing individuals that evolve into high-quality solutions. For instance, running the synthesis process (on examples of moderate size) without the shutdown improvement strategy often results in implementations which do not exploit this energy-reduction possibility.

### C. Hardware Core Allocation

For tasks mapped to ASICs and FPGAs, it is necessary to allocate hardware cores that are capable of executing the required

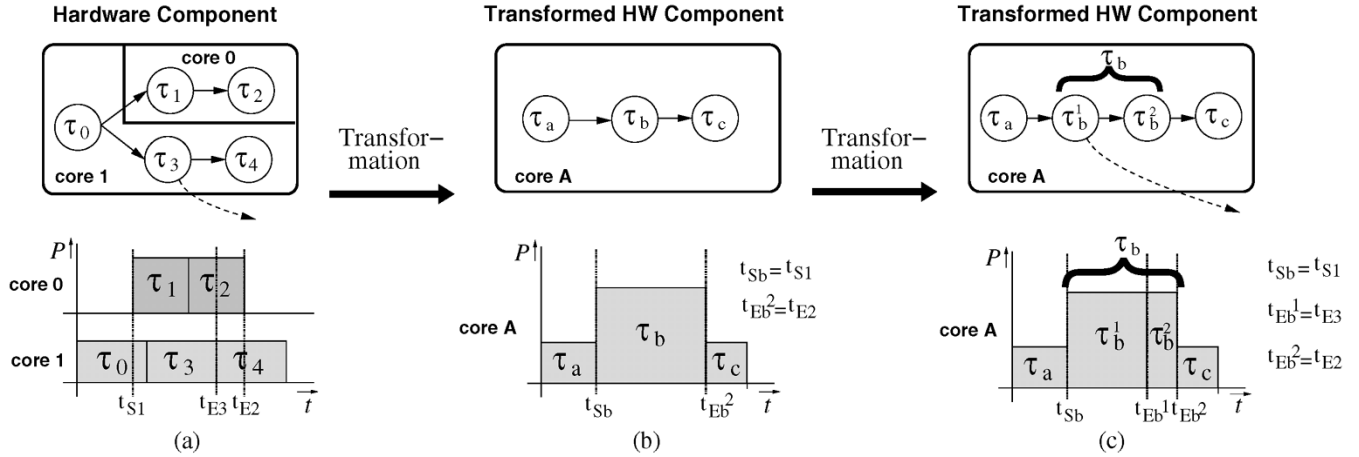


Fig. 9. DVS transformation for HW cores considering inter-PE communication. (a) Physical implementation. (b) Virtual implementation. (c) Virtual implementation.

task types. This is a trivial job as long as only tasks of different types are mapped to the same hardware component, i.e., when a single core for each task needs to be allocated. Nevertheless, if tasks of the *same* type  $\eta$  are assigned to the *same* PE more than once, it is necessary to make a decision upon how many cores of type  $\eta$  need to be implemented. This is important because hardware cores are able to execute tasks in parallel, i.e., the right quantitative choice of cores can efficiently help to exploit application parallelism, hence, improving the timing behavior as well as energy dissipation. In the proposed cosynthesis, the following approach is employed during the schedule optimization. Initially, each task type assigned to hardware is implemented only once, even if multiple tasks of this type are mapped onto the same PE. This ensures that all hardware tasks have at least one executable core implementation. If the hardware area constraints are not violated through the initial allocation, additional cores are implemented as follows. The tasks are analyzed to identify a possibly parallel-executing task, taking into account task dependencies. These tasks are then ordered according to their mobility. Clearly, tasks with low mobility are more likely to improve the timing behavior and, therefore, should be the preferred choice when implementing additional hardware cores. Cores for tasks with low mobility are implemented as long as the area constraints of the hardware components are not violated. Note that this strategy potentially improves the dynamic energy dissipation, since it is probable to result in more slack time, which, in turn, can be exploited through DVS.

#### D. DVS for Multiple Parallel-Executing Tasks

DVS is a powerful technique to reduce energy consumption by exploiting temporal performance requirements through dynamically adapting processing speed and supply voltage of PEs. The applicability of DVS to distributed embedded systems was demonstrated in [7], [17], [20], and [28]. However, these works concentrate on dynamically changing the performance of software PE's only, while parallel execution of tasks on hardware resources has been neglected. Nevertheless, in the context of energy-efficient multimode systems, where performance requirements of each operational mode can vary significantly, DVS

needs to be considered carefully. Consider, for instance, an inverse discrete cosine transformation (IDCT) algorithm implemented in fast hardware, which is used during two modes: 1) MP3 decoding and 2) JPEG image decoding. Clearly, the JPEG decoder should restore images as quickly as possible, i.e., the IDCT hardware is required to execute at maximal supply voltage (equivalent to peak performance). On the other hand, the MP3 decoder works at a fixed repetition rate of 25 ms, for which the hardware implementation operates faster than necessary, i.e., the IDCT performance can be reduced, such that this repetition rate is adequately met. By using DVS, it is possible to adapt the execution speed to suit both needs and to reduce the energy consumption to a minimum. Here, we consider that hardware components might employ DVS. However, due to the area and power overhead involved in additional DVS circuitry (dc/dc converter [14], [21]), it is assumed that all cores allocated to the same hardware component are fed by a single voltage supply, i.e., dynamically scaling the supply voltage simultaneously affects the performance of all cores on that hardware component.

The proposed technique enables an efficient usage of existing DVS approaches [7], [17], [20], [28] to handle the case of DVS on hardware components that execute tasks in parallel. To cope with this problem, the potentially parallel executing tasks on a single voltage-scalable hardware resource are transformed into an equivalent set of sequentially executing tasks, taking into account the dynamic power dissipation on each core. Note that this is done to calculate the scaled supply voltages only, i.e., this virtual transformation does not affect the real implementation. In this section, we highlight solely our transformation-based approach, while we refer the interested reader to [7], [17], [20], and [28], where different voltage scaling techniques are described in detail.

The following example illustratively outlines the proposed transformation approach. Fig. 9 shows the transformation of five hardware tasks, executing on two cores (both cores are implemented within the same hardware component), to four sequential tasks on a single core. The given schedules do not only reveal the activation times of the individual tasks, but further indicate their power dissipation over time (given by the height of the tasks). Such a power annotated schedule is referred to as

TABLE II  
CONSIDERING MODE-EXECUTION PROBABILITIES (EXCLUDING DVS)

Example (No. of modes)	Hyper- period (ms)	Mode Execution Probabilities	w/o probabilities		with probabilities		
			Average power (mW)	CPU time (s)	Average power (mW)	CPU time (s)	Reduction (%)
mul1 (4)	70,60,90,20	5:10:75:10	8.131	20.7	7.529	24.7	7.29
mul2 (4)	50,70,40,80	10:5:80:5	3.404	15.5	2.771	18.2	18.61
mul3 (5)	20,24,60,40,30	7:3:80:5:5	10.923	23.4	10.430	23.0	4.17
mul4 (5)	60,30,70,40,50	1:4:5:40:50	7.975	21.0	6.726	25.2	15.50
mul5 (4)	20,60,60,40	7:13:35:45	5.186	18.4	4.668	22.1	10.01
mul6 (4)	65,40,40,100	15:10:10:65	1.677	20.6	1.301	19.9	22.46
mul7 (4)	200,160,190,100	5:5:5:85	3.306	11.6	1.250	21.4	62.18
mul8 (4)	400,70,40,80	75:5:15:5	1.565	32.1	1.329	28.0	15.06
mul9 (4)	40,40,100,40	7:3:80:10	3.081	6.0	1.901	5.8	38.28
mul10 (5)	500,70,500,80,70	45:5:40:5:5	1.105	28.3	0.941	32.1	14.83
mul11 (3)	100,120,200	80:10:10	2.199	9.3	1.304	16.6	40.70
mul12 (4)	80,80,90,150	15:10:50:25	7.006	25.4	5.975	34.2	14.69
mul13 (3)	80,60,100	5:15:80	4.090	15.8	2.816	15.8	31.04
mul14 (5)	60,30,60,100,190	5:5:10:10:70	8.195	28.6	6.466	33.0	21.13
mul15 (5)	220,150,60,250,200	5:10:75:7:3	2.188	41.5	1.222	55.4	44.16

power-profile. The main advantage of the shown transformation lies within the fact that it results in sequentially executing tasks on a single component, which is equivalent to the behavior of software tasks. Hence, a voltage-scaling technique for software processors can be applied after the transformation, in order to exploit system idle times. The transformation is carried out in two main steps.

- 1) A single power profile is derived by adding the power profiles of both hardware cores and by splitting this power profile into individual tasks whenever the power values change. In Fig. 9(b), these points are  $t_{Sb}$  (the start time of task  $\tau_1$ ) and  $t_{Eb}$  (the end time of task  $\tau_2$ ). The power dissipation of tasks  $\tau_a$  and  $\tau_c$  are equivalent to the power of core 1, while task  $\tau_b$  dissipates a power which is the sum of the power consumptions on core 0 and core 1.
- 2) Further, for each outside-data dependency (indicated as a dashed arrow in Fig. 9), the virtual power profile is split and additional tasks are introduced. For the given example, task  $\tau_3$  communicates to an outside task. Since the execution of this task lies within the virtual task  $\tau_b$ , task  $\tau_b$  is split into  $\tau_b^1$  and  $\tau_b^2$ . In this way, the outside communication can be correctly included. Tasks with deadlines are handled in a similar fashion, in order to avoid that tasks are extended beyond their timing constraints.

## VII. EXPERIMENTAL RESULTS

Based on the techniques and algorithms presented in this work, a multimode synthesis approach has been implemented on a Pentium III 1.2-GHz Linux PC. In order to evaluate its capability of producing high-quality solutions in terms of energy consumption, timing behavior, and hardware area requirements, a set of experiments has been carried out on 15 automatically generated multimode examples (mul1-mul15<sup>1</sup>) and one real-life benchmark example (smart-phone).<sup>2</sup> All reported

results were obtained by running the optimization processes 40 times and averaging the outcomes. The average power dissipations as well as the energy consumptions have been calculated according to (1)–(4).

### A. Automatically Generated Examples

Each of the 15 generated examples (mul1-mul15) is specified by three to five operational modes, each consisting of 8 to 32 tasks (required execution cycles vary between 500–350 000). The used target architectures contain two to four heterogeneous PEs (clock frequencies are given in the range of 25–50 MHz), some of which are DVS-enabled. These PEs are interconnected through one to three communication links. The active power consumption of programmable processors was randomly chosen between 5 and 500 mW, depending on the executed task. The power dissipation of hardware components are selected to be one to two orders of magnitude lower. Further, the static power dissipation was set to be 5%–15% of the maximal active power. The execution probabilities of individual modes were randomly chosen and vary between 1% and 85%. Timing constraints have been assigned in the form of individual task deadlines as well as repetition periods to the modes (hyperperiods). The timing constraints were varied between 15 and 500 ms, such that schedulable implementations with up to 50% deadline slack could be found.

To illustrate the importance of taking mode-execution probabilities into account during the synthesis process, an execution probability neglecting approach is compared with the proposed synthesis technique, which considers the mode probabilities. The first two sets of experiments demonstrate the energy savings achievable through the consideration of mode-execution probabilities, either with or without the exploration of DVS. The third set examines the influence of the actual activation profile on the energy savings.

1) *Comparisons Excluding DVS*: To highlight the influence of mode-execution probabilities on the achievable energy saving, consider Table II which shows the multimode

<sup>1</sup>The examples were generated with the publicly available tool TGFF [13].

<sup>2</sup>The used benchmarks, including the realistic smart phone example, can be found at: <http://www.ida.liu.se/~g-marsc/benchmarks/>.

TABLE III  
CONSIDERING MODE-EXECUTION PROBABILITIES (INCLUDING DVS)

Example (No. of modes)	Hyper- period ( <i>ms</i> )	Mode Execution Probabilities	w/o probabilities		with probabilities		
			Average power ( <i>mW</i> )	CPU time ( <i>s</i> )	Average power ( <i>mW</i> )	CPU time ( <i>s</i> )	Reduction (%)
mul1 (4)	70,60,90,20	5:10:75:10	4.271	526.6	3.964	768.6	10.92
mul2 (4)	50,70,40,80	10:5:80:5	1.568	860.4	1.273	687.4	18.82
mul3 (5)	20,24,60,40,30	7:3:80:5:5	4.012	1053.5	3.344	1192.2	16.66
mul4 (5)	60,30,70,40,50	1:4:5:40:50	2.914	1135.2	2.320	1125.4	20.39
mul5 (3)	20,60,60,40	7:13:35:45	1.394	967.7	1.315	932.1	5.68
mul6 (4)	65,40,40,100	15:10:10:65	0.689	472.9	0.465	593.7	32.53
mul7 (4)	200,160,190,100	5:5:5:85	1.331	540.3	0.479	820.7	64.02
mul8 (4)	400,70,40,80	75:5:15:5	0.564	1262.1	0.436	1412.0	22.64
mul9 (4)	40,40,100,40	7:3:80:10	0.942	161.2	0.648	177.1	34.66
mul10 (5)	500,70,500,80,70	45:5:40:5:5	0.480	1456.3	0.394	1361.9	17.88
mul11 (3)	100,120,200	80:10:10	0.396	318.1	0.255	403.2	35.53
mul12 (4)	80,80,90,150	15:10:50:25	2.857	1384.7	2.460	1450.7	13.91
mul13 (3)	80,60,100	5:15:80	1.185	498.3	0.953	576.6	19.56
mul14 (5)	60,30,60,100,190	5:5:10:10:70	2.320	1512.3	1.797	1556.4	22.55
mul15 (5)	220,150,60,250,200	5:10:75:7:3	0.801	1316.7	0.324	1836.3	59.56

cosynthesis results for the 15 automatically generated benchmarks. The first three columns give the benchmark names, the hyperperiod (repetition period) of each mode, and the mode-execution probabilities. The fourth and the fifth columns present the dissipated average power and optimization time for the execution probability neglecting synthesis approach. Note that the execution probabilities are neglected during the synthesis only, while the computed power dissipations at the end of the synthesis incorporate the execution probabilities, in order to ensure a meaningful comparison. The sixth and seventh columns show the same for the proposed approach, which considers the execution probabilities throughout the synthesis process. Take, for instance, example *mul6*. When ignoring the execution probabilities during the optimization, an average power dissipation of 1.677 mW is achieved. However, optimizing the same benchmark example under the consideration that modes execute with *uneven* probabilities (e.g., 15:10:10:65—i.e., mode 1 is active for 15%, mode 2 is active of 10%, and so on), the average power can be reduced by an appropriate task mapping and core allocation to 1.301 mW. This is a significant reduction of 22.46%. Furthermore, it can be observed that the proposed technique was able to reduce the energy consumption of all examples with up to 62.18% (*mul7*). Note that these reductions are achieved without any modification of the underlying hardware architectures, i.e., the system costs are not increased. It is also important to note that the achieved energy reductions are solely introduced by taking the mode-execution probabilities into account during the cosynthesis process, i.e., both compared approaches allow the same resource sharing and rely on the same scheduling technique. When comparing the optimization times for both approaches, it can be observed that the proposed technique shows a slightly increased CPU time for most examples, which is mainly due to the more complex design space structure.

2) *Comparisons Including DVS*: The next experiments were conducted to see how the proposed technique compares to DVS and if further savings can be achieved by taking the mode prob-

abilities and DVS simultaneously into account. Table III reports on the findings. The DVS technique that was used here is based on PV-DVS [28], which has been extended to enable the consideration of DVS not only for software processors, but also for parallel executing cores on hardware PEs (see Section VI-D). As in the first experiments, two approaches are compared here. The first approach disregards the mode-execution probabilities during optimization, while the second takes them into account throughout the cosynthesis. Similar to Table II, the fourth and fifth columns of Table III show the results without consideration of execution probabilities, while the sixth and the seventh columns present the results achieved by taking execution probabilities into account. Let us consider again benchmark *mul6*. Although the execution probabilities are neglected in the fourth column, a reduced average power consumption (0.689 mW) can be observed, when compared to the results given in Table II. This clearly demonstrates the high energy reduction capabilities of DVS. Nevertheless, it is possible to further minimize the power consumption to 0.465 mW by considering the execution probabilities together with DVS. This is an improvement of 32.53%, solely due to the synthesis for the particular execution probabilities. For all other benchmarks savings of up to 64.02% (*mul7*) were achieved. Due to the computation of scaled supply voltages and the influence of scheduling on the energy consumption, the optimization times are higher when DVS is considered.

3) *Influence of Real Activation Probabilities*: The next experiment is conducted to highlight the influence of the real user behavior on the energy efficiency of a system that has been synthesized under the consideration of certain mode-execution probabilities. Certainly, the mode-execution probabilities which are used during the synthesis represent an “imaginative” user and the activation probabilities of a *real* user will differ from those. In accordance, our experiments try to answer the following question: How is the energy efficiency affected by the different activation profiles during application runtime? For experimental purposes, a simple specification with two modes is used, which contains 14 and 24 tasks (modes 1 and 2). The un-




Fig. 10. System specification consisting of two operational modes optimized for different execution probabilities (solid line: 0.1:0.9, dashed: 0.9:0.1, dotted: 0.5:0.5).

derlying architecture consists of two programmable processors and a single ASIC, all connected via a shared bus. This configuration was synthesised for three different pairs of execution probabilities (0.1:0.9, 0.9:0.1, and 0.5:0.5). These three implementation possibilities correspond to the three lines shown in Fig. 10. All implementations are based on the same hardware architecture; yet, each has a different task and communication mapping, core allocation, as well as schedule. The first solution (solid line) was synthesized under the consideration of execution probabilities 0.1:0.9, that is, it is assumed that modes 1 and 2 are active for 10% and 90% of the operational time, respectively. Similarly, the second (dashed line) and the third (dotted line) lines represent solutions that have been synthesized using execution probabilities 0.9:0.1 and 0.5:0.5, respectively. According to the real execution probabilities during runtime, i.e., the activation behavior of the user, the average power dissipations of the implemented systems vary. Consider the system optimized for execution probabilities 0.1:0.9 (solid line). If the user behavior corresponds to these probabilities (User A), the system dissipates an average power of approximately 5.5 mW (point I). However, if a different user (User B), for instance, uses mode 1 for 90% and mode 2 for 10% of the time (0.9:0.1), the system will dissipate approximately 12.5 mW (point II). Nevertheless, if the system would be optimized for this activation profile (0.9:0.1), as indicated by the dashed line in Fig. 10, a lower power dissipation of around 2.6 mW (point III) can be achieved. Similarly, if the system is optimized for execution probabilities 0.9:0.1 (dashed line) and the user runs the application 10% in mode 1 and 90% in mode 2 (User A), then a power dissipation of 12.5 mW (point IV) is given. While an optimization toward this usage profile can achieve a system implementation which dissipates only 5.5 mW (point I), i.e., extending the battery-lifetime by a factor of 2.83 times. The dotted plot in Fig. 10 represents the solution when the execution probabilities are neglected during the optimization, that is, the execution probabilities are considered to be equal for both modes. Of course, if

modes 1 and 2 are active for equal amounts of time, this solution achieves a lower power dissipation (6.5 mW, point V) than the systems optimized for execution probabilities 0.1:0.9 (9 mW, point VI) and 0.9:0.1 (7.6 mW, point VII). The figure reveals that the design for 0.1:0.9 (solid line) achieves the lowest power dissipation when the user complies to an activation profile between 0:1 and 0.21:0.79. While the designs for 0.5:0.5 (dotted line) and for 0.9:0.1 (dashed line) lead to the lowest energy dissipation in the ranges from 0.21:0.79 to 0.57:0.43 and 0.57:0.43 to 1:0, respectively. In summary, Fig. 10 clearly shows that the execution probabilities substantially influence the energy dissipation of the system. Certainly, the system should be optimized as close as possible toward the real behavior in order to achieve low energy consumptions, which, in turn, result in longer battery lifetimes.

### B. Smart Phone Benchmark

To further validate the cosynthesis technique in terms of real-world applicability, the introduced approach was applied to a smart-phone example. This benchmark is based on three publicly available applications: a GSM codec [3], a JPEG codec [4], and an MP3 decoder [18]. Accordingly, the smart-phone offers three different services to the user, namely, a GSM cellular phone, a digital camera, and an MP3-player. Of course, the used applications do not specify the whole smart phone device, however, a major digital part of it. The specification for this example, given as OMSM, has already been introduced in Fig. 2. For each of the eight operational modes, the corresponding task graphs have been extracted from the above given references. The individual applications have been software profiled to gather the necessary execution characteristics of each task. This was carried out by compiling profile information into the application [1], [2] and running the produced software on real-life input streams. On the other hand, the hardware estimations are not based on direct measurements, but have been based on typical values, such that hardware tasks execute

TABLE IV  
SMART PHONE EXPERIMENTS WITHOUT DVS

Mode	No. of Tasks/Comm.	Exec. Prob.	Hyper-period (s)	without probabilities		with probabilities	
				$E_O/HP_O$ (mJ)	$\bar{p}_O$ (mW)	$E_O/HP_O$ (mJ)	$\bar{p}_O$ (mW)
0	88/137	0.09	0.020	0.2637	1.1868	0.1272	0.5723
1	12/0	0.74	1.000	0.8263	0.6115	0.8210	0.6075
2	12/0	0.01	1.000	1.7176	0.0172	1.7110	0.0171
3	5/4	0.02	0.250	1.3004	0.1040	0.9545	0.0764
4	12/5	0.02	0.500	1.6650	0.0666	1.3761	0.0550
5	17/16	0.10	0.025	0.1231	0.4922	0.0719	0.2874
6	17/16	0.01	0.025	0.2203	0.0881	0.3971	0.1588
7	12/5	0.01	0.500	1.7884	0.0358	1.3245	0.0265
Overall					2.6022		1.8011

TABLE V  
SMART PHONE EXPERIMENTS WITH DVS

Mode	No. of Tasks/Comm.	Exec. Prob.	Hyper-period (s)	without probabilities		with probabilities	
				$E_O/HP_O$ (mJ)	$\bar{p}_O$ (mW)	$E_O/HP_O$ (mJ)	$\bar{p}_O$ (mW)
0	88/137	0.09	0.020	0.0746	0.3355	0.0786	0.3539
1	12/0	0.74	1.000	0.8190	0.6061	0.0180	0.0133
2	12/0	0.01	1.000	0.0280	0.0003	0.8110	0.0081
3	5/4	0.02	0.250	0.3355	0.0268	0.3545	0.0284
4	12/5	0.02	0.500	0.3556	0.0142	0.8412	0.0336
5	17/16	0.10	0.025	0.0513	0.2052	0.0813	0.3250
6	17/16	0.01	0.025	0.0492	0.0197	0.1975	0.0791
7	12/5	0.01	0.500	0.4917	0.0098	0.8671	0.0173
Overall					1.2176		0.8587

one to two orders of magnitude faster and dissipated one to two orders of magnitude less power than their software counterparts [9]. Depending on the operational mode, the number of tasks and communications varies between five and 88 nodes and zero and 137 edges, respectively. The hardware architecture of the embedded system within the smart phone consists of one DVS-enabled processor (execution properties are based on values given for the ARM8 developed in [8]) and two ASICs. These components are connected via a single bus. Tables IV and V give the results of the conducted experiments, distinguishing between optimizations without and with the consideration of DVS.

Similar to the previous experiments, approaches which neglect the execution probabilities are compared with the introduced cosynthesis technique that considers the uneven activation times of different modes. Table IV shows this comparison for a fixed voltage system, i.e., no DVS is applied. The table provides information regarding all eight modes of the smart phone. This mode information includes benchmark properties such as complexity, execution probability, and hyper-period. Furthermore, the table gives the achieved energy dissipation for the mode hyperperiod and average power consumption of each mode. The average power consumption can be calculated from the energy values by dividing the energy by the hyperperiod and multiplying the result with the execution probability. Synthesizing the system without consideration of execution probabilities results in an overall average power consumption of 2.6022 mW, when running the system after the synthesis according to the activation profile. Nevertheless, taking into account the mode usage

profile during the cosynthesis this can be reduced by 30.76% to 1.8011 mW. Please note that the given overall average power consumption is calculated based on (1)–(4); hence, these values are directly proportional to the battery-lifetime. The saving is achieved without the modification of the allocated hardware architecture, therefore, the system cost is the same for both solutions.

Also, DVS has been applied to this benchmark, considering that the GPP of the given architecture supports DVS functionality. The results are shown in Table V. It can be observed that the overall average power consumption of the smart phone drops to 1.2176 mW, even when neglecting mode-execution probabilities. However, the combination of applying DVS and taking execution probabilities into account results in the lowest power consumption of 0.8587 mW, a 29.5% reduction, when compared to the activation profile neglecting approach. That is, solely by considering the activation profile during the synthesis, the battery-lifetime could be extended by one third, even when using a system that employs DVS components. Overall, the average power is decreased from 2.602 to 0.859 mW, which represents a significant reduction of nearly 67%. Regarding the required cosynthesis times, the four implementations could be found in 80.1 s (without probabilities and DVS) to 4344.8 s (with probabilities and DVS). Clearly, considering DVS requires longer optimization times due to the voltage-scaling problem that needs to be solved repetitively within the innermost optimization loop of the cosynthesis algorithm. For instance, the optimization for DVS increases the runtime from 80.1 to 3754.1 s for the case without consideration of execution probabilities, and from 96.9

to 4344.8 s when execution probabilities are taken into account. On the other hand, the consideration of mode-execution probabilities increases the optimization time only moderately from 80.1 to 96.9 s in the case of no DVS, and from 3754.5 to 4344.8 s if DVS is considered.

### VIII. CONCLUDING REMARKS

We have introduced new techniques and algorithms for the energy minimization of multimode embedded systems. An abstract specification model called OMSM has been proposed. This model allows for the specification of mode interaction (top-level finite state machine) as well as mode functionality (task graph). The advantage of such a representation is the capability to express the complete functionality of the system within a single model, containing both control and data flow.

The presented cosynthesis technique not only optimizes mapping and scheduling toward hardware cost and timing behavior, but also aims at the reduction of power consumption at the same time. A key contribution has been the development of an effective mapping strategy that considers uneven mode-execution probabilities as well as important power reduction aspects, such as multiple task implementations and core allocation. For this purpose, a GA-based mapping approach has been proposed along with four improvement strategies to effectively handle the optimization of component shutdown, transition time, area usage, and timing behavior. These improvement strategies guide the mapping optimization of multimode specifications toward high quality solutions in terms of power consumption, timing feasibility, and area usage. A newly introduced transformation-based algorithm for DVS-enabled hardware components, which handles parallel task execution, allows to leverage the efficiency of existing voltage scaling algorithms. This algorithm transforms a set of potentially parallel executing tasks on a single HW component into a set of sequential executing tasks, taking into account imposed deadlines and inter-PE communications.

The proposed techniques and algorithms have been validated through extensive experiments, including a smart phone real-life example. These experiments have demonstrated that taking into account mode-execution probabilities throughout the system synthesis leads to substantial energy savings compared to conventional approaches which neglect this issue. Furthermore, DVS has been considered in the context of multimode embedded systems and it was shown that considerably high energy reductions can be achieved by combining both the consideration of execution probabilities and DVS.

### ACKNOWLEDGMENT

The authors are thankful to the anonymous reviewers for their valuable comments and suggestions that helped to improve this manuscript.

### REFERENCES

- [1] GNU CC Manual [Online]. Available: <http://gcc.gnu.org/>
- [2] GNU Gprof Manual [Online]. Available: <http://www.gnu.org/manual/gprof-2.9.1/gprof.html>
- [3] GSM 06.10, Technical University of Berlin [Online]. Available: <http://kbs.cs.tu-berlin.de/~jutta/toast.html>
- [4] Independent JPEG Group: jpeg-6b [Online]. Available: <ftp://ftp.uu.net/graphics/jpeg/jpegsrc.v6b.tar.gz>
- [5] Intel XScale Core, Developer's Manual, Dec. 2000.
- [6] Mobile AMD Athlon 4, Processor Model 6 CPGA Data Sheet, Nov. 2000. Publication No 24 319 Rev E.
- [7] A. Andrei, M. T. Schmitz, P. Eles, Z. Peng, and B. M. Al-Hashimi, "Overhead-conscious voltage selection for dynamic and leakage power reduction of time-constraint systems," in *Proc. Design, Automation, Test Eur. Conf.*, Feb. 2004, pp. 518–523.
- [8] T. D. Burd, "Energy-efficient processor system design," Ph.D. dissertation, Univ. California, Berkeley, 2001.
- [9] T. D. Burd and R. W. Brodersen, "Processor design for portable systems," *J. VLSI Signal Process.*, vol. 13, no. 2, pp. 203–222, Aug. 1996.
- [10] A. Chandrakasan, W. J. Bowhill, and F. Fox, *Design of High-Performance Microprocessor Circuits*. New York/Piscataway, NJ: Wiley/IEEE Press, 2001.
- [11] A. P. Chandrakasan and R. W. Brodersen, *Low Power Digital CMOS Design*. Norwell, MA: Kluwer, 1995.
- [12] E.-Y. Chung, L. Benini, and G. De Micheli, "Contents provider-assisted dynamic voltage scaling for low energy multimedia applications," in *Proc. Int. Symp. Low-Power Electron. Design*, Aug. 2002, pp. 42–47.
- [13] R. Dick, D. Rhodes, and W. Wolf, "TGFF: Task graphs for free," in *Proc. 5th Int. Workshop Hardware/Software Co-Design*, Mar. 1998, pp. 97–101.
- [14] J. Goodman, A. Chandrakasan, and A. P. Dancy, "Design and implementation of a scalable encryption processor with embedded variable dc/dc converter," in *Proc. IEEE 36th Design Automation Conf.*, 1999, pp. 855–860.
- [15] M. Grajcar, "Genetic list scheduling algorithm for scheduling and allocation on a loosely coupled heterogeneous multiprocessor system," in *Proc. IEEE 36th Design Automation Conf.*, 1999, pp. 280–285.
- [16] T. Grötzer, R. Schoenen, and H. Meyr, "PCC: A modeling technique for mixed control/data flow systems," presented at the IEEE Eur. Design Test Conf., Mar. 1997.
- [17] F. Gruian and K. Kuchcinski, "LEneS: Task scheduling for low-energy systems using variable supply voltage processors," in *Proc. Asia South Pacific—Design Automation Conf.*, Jan. 2001, pp. 449–455.
- [18] Mpeg3Play-0.9.6., by J. Hagman. [Online]. Available: <http://home.swipnet.se/~w-10694/tars/mpeg3play-0.9.6-x86.tar.gz>
- [19] A. Kalavade and P. A. Subrahmanyam, "Hardware/software partitioning for multifunction systems," *IEEE Trans. Computer-Aided Design Integr. Circuits Syst.*, vol. 17, no. 9, pp. 819–836, Sep. 1998.
- [20] J. Luo and N. K. Jha, "Battery-aware static scheduling for distributed real-time embedded systems," in *Proc. IEEE 38th Design Automation Conf.*, 2001, pp. 444–449.
- [21] W. Namgoong, M. Yu, and T. Meng, "A high-efficiency variable-voltage CMOS dynamic dc-dc switching regulator," in *Proc. Int. Solid-State Circuits Conf.*, 1997, pp. 380–381.
- [22] H. Oh and S. Ha, "A static scheduling heuristic for heterogeneous processors," presented at the 2nd Int. EuroPar Conf., vol. II, Aug. 1996.
- [23] —, "Hardware-software cosynthesis of multimode multi-task embedded systems with real-time constraints," in *Proc. 2nd Int. Symp. Hardware/Software Co-Design*, May 2002, pp. 133–138.
- [24] P. Pedro and A. Burns, "Schedulability analysis for mode changes in flexible real-time systems," in *Proc. Euromicro Workshop Real-Time Syst.*, Jun. 1998, pp. 17–19.
- [25] M. T. Schmitz, B. M. Al-Hashimi, and P. Eles, "Energy-efficient mapping and scheduling for DVS enabled distributed embedded systems," in *Proc. Design, Automation Test Eur. Conf.*, Mar. 2002, pp. 514–521.
- [26] —, "Synthesizing energy-efficient embedded systems with LOPOCOS," *J. Design Automation Embedded Syst.*, vol. 6, no. 4, pp. 401–424, 2002.
- [27] —, "Iterative schedule optimization for voltage scalable distributed embedded systems," *ACM Trans. Embedded Comput. Syst.*, vol. 3, no. 1, pp. 182–217, 2004.
- [28] —, "Iterative schedule optimization for voltage scalable distributed embedded systems," *ACM Trans. Embedded Syst. Design*, vol. 3, no. 1, pp. 182–217, 2004.
- [29] L. Sha, R. Rajkumar, J. Lehoczky, and K. Ramamritham, "Mode change protocols for priority-driven preemptive scheduling," *Real-Time Syst.*, vol. 1, pp. 243–265, 1989.
- [30] Y. Shin, D. Kim, and K. Choi, "Schedulability-driven performance analysis of multiple mode embedded real-time systems," in *Proc. IEEE 37th Design Automation Conf.*, Jun. 2000, pp. 495–500.
- [31] M. Wu and D. Gajski, "Hypertool: A programming aid for message-passing systems," *IEEE Trans. Parallel Distrib. Syst.*, vol. 1, no. 3, pp. 330–343, Jul. 1990.



- [32] P. Yang, P. Marchal, C. Wong, S. Himpe, F. Catthoor, P. David, J. Vounckx, and R. Lauwereins, "Managing dynamic concurrent tasks in embedded real-time multimedia systems," in *Proc. Int. Symp. Syst. Synthesis*, Oct. 2002, pp. 112–119.

**Marcus T. Schmitz** received the Dipl.-Ing. (FH) degree in electrical engineering from the University of Applied Science Koblenz, Koblenz, Germany, in 1999 and the Ph.D. degree in electronics from the University of Southampton, Southampton, UK, in 2003.

Since April 2003, he has been a Postdoctoral Researcher in the Embedded Systems Laboratory, Department of Computer and Information Science, Linköping University, Linköping, Sweden, where he is working on the development of offline and online voltage scaling techniques for embedded systems. He has published several papers in the area of dynamic voltage scaling for heterogeneous embedded systems and he has co-authored the book *System-Level Design Techniques for Energy-Efficient Embedded Systems* (Norwell, MA: Kluwer, 2004). His research interests include system-level co-design, application-driven design methodologies, energy-efficient system design, and reconfigurable architectures.

**Bashir M. Al-Hashimi** (M'99–SM'01) received the B.Sc. degree (with first-class classification) in electrical and electronics engineering from the University of Bath, Bath, UK, in 1984 and the Ph.D. degree from York University, York, UK, in 1989.

He worked in industry for six years designing high-performance chips for analog- and digital-signal processing applications, developing computer-aided design tools for simulation and synthesis of analog and digital circuits. In 1995, he joined Staffordshire University, Staffordshire, UK, where he formed the VLSI Signal Processing research group with funding from the industry and the government. In 1999, he joined the Department of Electronics and Computer Science, Southampton University, Southampton, UK, where he is currently a Professor of Computer Engineering. He has authored one book on SPICE simulation and coauthored two books, *Power Constrained Testing of VLSI Circuits* (Norwell, MA: Kluwer, 2002) and *System-Level Design Techniques for Energy-Efficient Embedded Systems* (Norwell, MA: Kluwer, 2004). He has authored and coauthored over 125 technical papers. His research and teaching interests include low-power hardware/software co-design, system-on-chip test, and VLSI computer-aided design.

Dr. Al-Hashimi is the Editor-in-Chief of the *IEE Proceedings: Computers and Digital Techniques* and is a Member of the Editorial Board of the *Embedded Computing Journal*. He has served on the IEEE/ACM Design Automation and Test in Europe (DATE) Conference in various capacities, including: Topic Chair (Simulation and Emulation, 2001), and as a Member of the Executive Committee for 2005. He is a Member of the Technical Program Committee of the IEEE European Test Symposium and the IEEE VLSI Test Symposium. He has served as the Guest Editor for a number of journal special issues, including most recently: *Design Automation for Embedded Systems Journal* and the *IEE Proceedings: Computers and Digital Techniques*. He was a corecipient of a Best Paper Award at the 2000 IEEE International Test Conference relating to low power built-in self-test for register transfer level data paths.

**Petru Eles** (M'99) received the M.Sc. degree in computer science from the "Politehnica" University of Timisoara, Timisoara, Romania, in 1979 and the Ph.D. degree in computer science from the "Politehnica" University of Bucharest, Bucharest, Romania, in 1993.

He is currently a Professor with the Department of Computer and Information Science, Linköping University, Linköping, Sweden. He has published extensively, including having coauthored *System Synthesis with VHDL* (Norwell, MA: Kluwer, 1997) and *System-Level Design Techniques for Energy-Efficient Embedded Systems* (Norwell, MA: Kluwer, 2004). His research interests include the design of embedded systems, hardware/software co-design, real-time systems, system specification and testing, computer-aided design for digital systems.

Prof. Eles is an Associate Editor of the *IEE Proceedings—Computers and Digital Techniques*. He has served as a Program Committee Member for numerous International Conferences such as DATE, ICCAD, ECRTS, CASES, EMSOFT, RTSS, and as Technical Program Chair of the IEEE/ACM/IFIP International Conference on Hardware/Software Co-design and System Synthesis. He was a corecipient of the Best Paper Awards at the 1992 and 1994 European Design Automation Conference, and a corecipient of the Best Presentation Award at the 2003 IEEE/ACM/IFIP International Conference on Hardware/Software Co-design and System Synthesis.