

On-Line Global Energy Optimization in Multi-Core Systems Using Principles of Analog Computation

Zeynep Toprak Deniz, Yusuf Leblebici, Eric Vittoz
Ecole Polytechnique Fédérale de Lausanne
Microelectronic Systems Laboratory
Lausanne, Switzerland
(zeynep.toprak, yusuf.leblebici, eric.vittoz)@epfl.ch

Abstract— This work presents the design and the silicon implementation of an on-line energy optimizer unit, which is capable of dynamically adjusting power supply voltages and operating frequencies of multiple processing elements. The optimized voltage/frequency assignments are tailored to the instantaneous workload information and fully adaptive to variations in process and temperature. The optimizer unit has a fast response time of 50 μ s, occupies a silicon area of 0.021mm² / task and dissipates 2 mW / task.

I. INTRODUCTION

The significance of the energy management problem is underlined by the increasing prominence of multi-core systems that must operate under strict energy budget constraints in mobile applications. In multi-processing element (PE) systems, due to the diversity of the applications that run within the system and their different degrees of parallelism, the workloads imposed on the system components are non-uniform over time. This introduces slack times during which the system can reduce its performance to save energy. The key to energy-efficient designs is the ability to tune PE performance to the non-uniform workload.

Dynamic voltage scaling (DVS) is based on reducing the performance level of the component during periods of low utilization so that the task is always completed just-in-time, consuming minimum energy. While the *local* energy dissipation of each PE can be minimized using DVS techniques based on workload predictions, it can be shown that these *local* minima usually do not represent the *global* energy minimum, which can only be reached by considering the relative timing dependencies of all tasks running in the system. This problem of minimizing the overall dissipated energy in a multi-PE system under timing constraints and subject to DVS, has already been formulated in a rigorous fashion, yet a compact real-time implementation has not been offered [1,2,3].

Our approach demonstrates the solution to the problem of on-line optimization of the dissipated energy in multi-PE

systems with interrelated tasks under timing constraints using the basic principles of analog computation by converging on the global minima of the constrained optimization problem which are represented as stable operating points of a simple resistive network (RN). The input set of the circuit consists of individual workload estimates for each task and for each PE, while the output consists of assigned supply voltage/frequency values for each PE as well as the allocated time duration for each task as illustrated in Fig. 1.

The remainder of this paper is organized as follows. In Sections 2 and 3 we concentrate on demonstrating an on-line solution to complex multi-variable energy optimization problem. The closed loop operation principle of the proposed analog optimizer block is described in Section 4. In Section 5 experimental results is discussed in comparison with the simulation results. Conclusions are provided in Section 6.




Figure 1. Block diagram representation of the proposed on-line global energy management unit.

II. FROM TASK GRAPH TO RESISTIVE NETWORK

The authors have previously demonstrated the clear *analogy* between the problem of *minimizing energy consumption on a complex system under timing constraints*, and the problem of *minimizing power dissipation in a resistive network under Kirchoff's Current Law (KCL) constraints* [4]. According to Maxwell's Heat theorem the RN will consume the *lowest possible power* (P_{total}), at steady-state for a given driving current [5]. The equivalence,

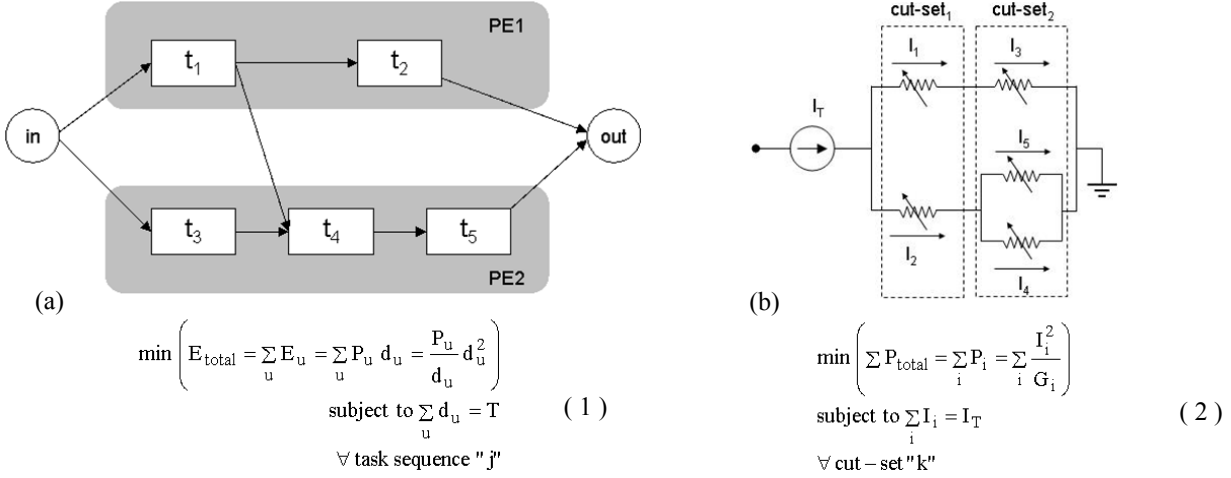


Figure 2. (a) Task graph of five tasks mapped on two processing elements, and (b) the resistive network equivalent of the given TG.

between the two analogous minimization problems is summarized in (1) and (2). Here, individual tasks are modeled with branch conductances (G_i), controlled by the ratio (d_u / P_u) where P_u is the average power consumption of the PE during task t_u , and d_u is the task duration.

Consider the example illustrated in Fig. 2(a), where five tasks are mapped and scheduled on two PEs. The total dissipated energy in the system can be written as the summation of all the task energies (1). Since each task (t_u) requires a given number (N_u) of cycles, and each cycle consumes an amount of energy, this amount can be reduced if the supply voltage (V_{PEi}) of the PE is reduced under the cost of cycle time (CT_u) increase. The formal algorithmic solution of (1) is certainly possible in real time, yet the computational overhead that is needed may become prohibitive especially when taking into account realistic timing/delay models and secondary effects such as leakage dissipation. Fig. 2(b) shows the equivalent RN of the given task graph (TG), in which duration d_u of each task t_u corresponds to the current I_i in a resistor R_i . The TG period T corresponds to the total current I_T driving the circuit. Due to KCL, I_T will be split into parallel branch currents that are inversely proportional to branch resistances. Hence, it can be seen that the simple RN actually realizes the solution to the dissipated power minimization problem under KCL constraints (2). It is important to emphasize that the mapping of a given TG to its equivalent RN is based on converting the time domain relation between tasks into equivalent RN currents. Hence, we do not consider this procedure to be equivalent to finding the dual of a given TG.

III. IMPLEMENTATION OF THE ANALOG OPTIMIZER

The total energy E_{total} required by the system to execute the whole set of tasks within a fixed duration T , is emulated by the power P_{total} dissipated in the equivalent RN, driven by a current I_T . Each resistor in the equivalent RN is implemented as a pseudo-resistor (R_i^*), so that its value can

be adjusted proportionally to the ratio given by (d_u / P_u) by means of a feedback loop that includes a calculation of P_u .

Fig. 3 shows the simplified block diagram implementation of the feedback loop for one branch conductance, where a *current-based* approach is used to represent key loop variables. A key element of the loop is the dynamic Ghost Circuit (GC) that emulates the maximum operation frequency of the processing element operated at the same supply voltage (V_{PEi}). This GC is essentially a ring oscillator replicating the critical path of the PE that is used in each loop to continuously determine the minimum supply voltage and the supply current that correspond to a target operation frequency for the PE. It is forced to run at frequency $f_i = 1 / CT_i$ which is imposed by its supply current ($I_{FTi} \propto N_i / I_i$). The predicted workload information (N_i) is injected into each loop in the form of a 4-bit external control variable. Any change in N_i influences the current corresponding to the target operation frequency (I_{FTi}) in the feedback loop. Hence, the simple GC determines the supply voltage level to be applied to the PE for achieving the target frequency as well as the resulting dynamic current consumption (I_{gi}). The voltage V_i and the frequency f_i are transmitted to the PE. They are also converted to current representations I_{Vi} and I_{Fi} in order to calculate the pseudo-resistor controlling currents (I_{Gi}). Current-mode processing in each feedback loop is carried out by single quadrant current multiplier/dividers labeled as TLL_i. Each current operator is implemented by the simple alternating topology translinear loop of four transistors operated in *weak inversion* as shown in the inset of Fig. 3.

Each pseudo-resistor is realized as a single MOS transistor operating in weak inversion where the equivalent conductance value of each transistor is controlled independently by a current by means of a control transistor (Fig. 3) – thus, utilizing only a few transistors. Note that the linear pseudo-Ohm's law is still valid and the network of controlled resistors remain linear with respect to currents [6].

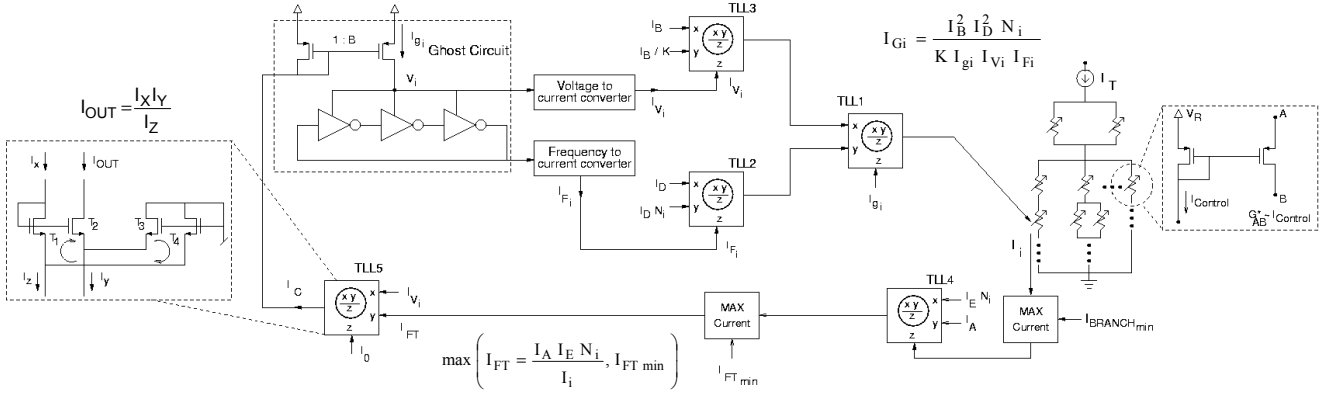


Figure 3. Block diagram implementation of the optimization feedback loop.

The result of current mode processing in each loop is the current I_{G_i} (3) that drives the corresponding pseudo-resistor as illustrated in Fig. 3. The factor K introduced by TLL₃ is proportional to the equivalent switching capacitor, that may be different for different processing elements. Here, I_{S_i} represents the modeled static current consumption of the PE (proportional to the total number of gates), with a static GC which is added to the loop. This current is added to the dynamic current consumption (I_{g_i}), resulting in (3).

$$G_i = \frac{1}{R_i} \propto I_{G_i} = \frac{I_B^2 I_D^2 N_i}{K (I_{g_i} + I_{S_i}) I_{V_i} I_{F_i}} \propto \frac{I_i}{P_u} \quad (3)$$

Consequently, the corresponding branch conductance value changes according to I_{G_i} (3). This change in the value of branch conductance forces all the branch currents in the RN to be adjusted by means of KCL. As the system settles to its new operating point, the new branch currents in the pseudo-RN are determined by KCL, dictating the optimum task duration with the prescribed supply voltage and operating frequency for each PE and for each task to minimize system-wide energy dissipation.

IV. CLOSED LOOP OPERATION OF THE OPTIMIZER

It is important to highlight that the feedback loop responsible for updating each G_i value operates in continuous time (based on GC response), rather than in a discrete-time iteration. The stability behavior of the feedback loops taking into account the coupling between loops through the RN has been thoroughly analyzed. It was shown that the dynamic behavior of each resistive element control loop is governed by a single-dominant-pole transfer function. Therefore, it was shown analytically that the entire system always converges to a stable and unique operating point for a given set of workloads. Also, note that the GC can effectively capture the actual frequency-voltage-power relationship of the PEs, reflecting the actual operating conditions on-chip (inherently taking into account the local variations of temperature, as well as process-related fluctuations of device parameters) eliminating any analytical approximation of the physical behavior that is inherently prone to inaccuracies.

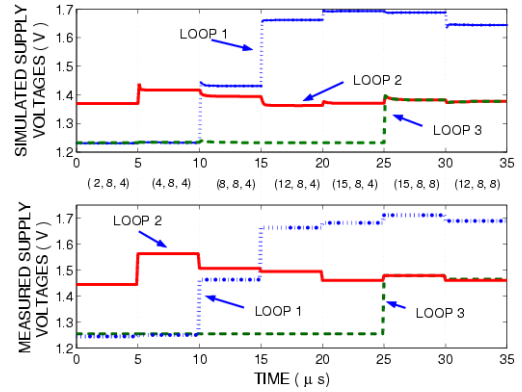


Figure 4. Simulated and measured supply voltages of the three-parallel loop optimizer circuit for an arbitrary sequence of workload combinations.




Figure 5. The corresponding branch currents (task durations) of the three-parallel loop optimizer circuit for the same workload conditions as in Fig. 4.

Fig. 4 shows the simulated vs. measured operation of a three-loop optimizer network which is used to model the behavior of a TG comprising three sequential tasks. Here, the supply voltages resulting in the optimum system energy dissipation are shown for various workload combinations indicated as (N_1, N_2, N_3) for each simulation interval. Similarly, Fig. 5 shows the corresponding simulated and measured task durations (branch currents) for the same set of

workload conditions. The available time is shared among the three tasks for all workload conditions; guaranteeing timing constraints and optimizing the dissipated energy in the system by means of optimally utilizing the available time. The comparison of measured and simulated branch currents as well as the GC supply voltages shows a good agreement between simulated and measured values.

The comparison of the simulated supply voltages (V), operation frequencies (MHz) and task durations (branch currents- μ A) of the same system has been made for the proposed global optimization approach versus local energy optimization applied to each task. When using the proposed global optimization approach, any change in workload condition of any of the tasks influences all task durations (hence, supply voltages and operation frequency) corresponding to a minimization of the total system energy dissipation by optimally using the overall available time (T). It was demonstrated that the additional energy savings is larger than 11%, even in the worst case [4].

V. MEASUREMENT RESULTS

The three-loop demonstrator circuit of the proposed analog optimizer architecture has been implemented using a 0.18 μ m standard digital CMOS process (Fig. 6). The overall circuit area of the optimizer is (250 μ m x 700 μ m) excluding decoupling capacitors, while each loop circuit occupies only (180 μ m x 120 μ m). The circuit is capable of supporting the desired frequency range of 170 MHz – 290 MHz, as well as the voltage range of 1.2 V – 1.8 V (Fig. 5). The measured worst-case settling time for supply voltages is less than 50 μ s. The average power consumption of the entire three-loop optimizer is 6.5 mW.




Figure 6. Chip microphotograph of the three-parallel loop optimizer.

Fig. 7 shows the variation of the overall energy dissipation of the same system composed of three sequential tasks as a function of changing workload conditions, calculated from measured voltage/frequency and task duration values. To test the optimality of this solution, the branch current values were perturbed from their actual values (while keeping the sum constant) and the energy surface has been re-calculated. The resulting energy surface is clearly *higher* than the original solution for all workload combinations and for all branch current perturbations, demonstrating that the original solution indeed is the minimum energy surface.

The mismatch in loops of the analog optimizer is equivalent to the relative error in predicted (estimated) workload levels. Consequently, the accuracy of the system can be modeled with the precision of the estimated workload conditions. Furthermore, since the workload of a given task is represented with 4-bit coded value, an error of approximately 6% in the predicted workload of each task is inevitable due to the quantization. Still, this can be calibrated with pre-correction per loop after fabrication.

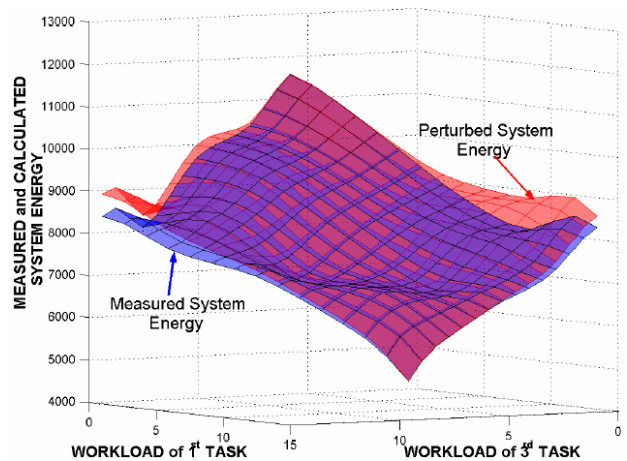


Figure 7. Comparison of the measured and the perturbed system energies.

VI. CONCLUSIONS

In this work, the analogy between the energy minimization problem under timing constraints in a general TG and the power minimization problem under KCL constraints in an equivalent RN is exploited. A novel fully analog, current-based solution to implement on-line energy minimization in complex multi-core systems under varying workload conditions is demonstrated. It is shown that the proposed approach achieves significant overall energy savings compared to the local energy minimization approach.

REFERENCES

- [1] Y. Zhang, X.S. Hu, D.Z. Chen, "Energy minimization of real-time tasks on variable voltage processors with transition energy overhead", ASPDAC, pp.65-70, 2003.
- [2] A. Andrei, M. Schmitz, P. Eles, Z. Peng, B.M. Al-Hashimi, "Overhead-Conscious Voltage Selection for Dynamic and Leakage Energy Reduction of Time-Constrained Systems", DATE 2004, pp.105-118.
- [3] Y. Zhang, X. Hu, D. Chen, "Task scheduling and voltage selection for energy minimization", DAC 2002, pp.183-188.
- [4] Z. Toprak, Y. Leblebici, E. Vittoz, "On-Line Global Energy Optimization in Multi-Core Embedded Systems Based on Analog Computation", submitted to 43rd Design Automation Conference.
- [5] James C. Maxwell, A Treatise on Electricity and Magnetism, 3rd ed., vol. 1. Oxford: Clarendon, 1892, pp.399-410.
- [6] E.A. Vittoz, "Analog VLSI for collective computation", IEEE International Conference on Electronics, Circuits and Systems, pp.3-6, 1998.