

System-Level Exploration for Pareto-Optimal Configurations in Parameterized System-on-a-Chip (December 2002)

Tony Givargis, Frank Vahid, and Jörg Henkel, *Member, IEEE*

Abstract—In this work, we provide a technique for efficiently exploring the power/performance design space of a parameterized system-on-a-chip (SOC) architecture to find all Pareto-optimal configurations. These Pareto-optimal configurations will represent the range of power and performance tradeoffs that are obtainable by adjusting parameter values for a fixed application that is mapped on the SOC architecture. Our approach extensively prunes the potentially large configuration space by taking advantage of parameter dependencies. We have successfully applied our technique to explore Pareto-optimal configurations of our SOC architecture for a number of applications.

Index Terms—Design Space Exploration, Low Power Design, Pareto-optimal Configurations, Platform Based Design, and System-on-a-chip Design.

I. INTRODUCTION

THE growing demand for portable embedded computing devices is leading to new system-on-a-chip (SOC) architectures intended for embedded systems. Such SOC architectures must be general enough to be used across several different applications, in order to be economically viable, leading to recent attention to parameterized SOC architectures. Different applications often have very different power and performance requirements. Therefore, these parameterized SOC architectures must be optimally configured to meet varied power and performance requirements of a large class of applications.

A typical SOC architecture will have a processor core, one or more caches, on-chip bus hierarchy, on-chip memory, and a large number of peripheral cores that provide application specific functionality such as multi-media and communication processing. Each of these SOC cores is likely to be

parameterized, enabling a designer to tune a core's settings for a specific application that is to be mapped on the SOC architecture. For example, the on-chip buses may be configured to use bus-invert [1] coding for low power, or the caches may be configured to use a greater or lesser degree of associativity for increased performance [2][3]. An assignment of a value to each of these parameters will impact the overall performance and power consumption of the SOC architecture. However, such impacts are highly dependent on the application running on the SOC. Therefore, a designer must have a method for finding a feasible set of parameter values, referred to as a configuration of the SOC, that meets the specification requirements. We outline an exploration approach that efficiently searches the entire configuration space and outputs Pareto-optimal configurations providing the designer with only the interesting configurations that result in a tradeoff between power and performance.

Our exploration algorithm fits in the SOC design flow as follows. As depicted in Fig. 1, the SOC provider provides a parameterized architecture in HDL or configurable IC format along with all the traditional development tools such as compilers, debuggers, emulators, etc. In addition, the SOC provider provides a system-level model and a tuning environment. This tuning environment enables the SOC user to search the parameter space of the SOC and to find a configuration that meets power and performance requirements of the target application. This tuning application is the focus of this work.

The remainder of this paper is organized as follows. In Section 2, we describe related work. In Section 3, we state the parameterized SOC exploration problem and outline our

Manuscript received October 30, 2000, revised August 23, 2001. This work was supported by the National Science Foundation (CCR-9811164, CCR-9876006), NEC, and a Design Automation Conference Graduate Scholarship.

Tony Givargis is with the Department of Information and Computer Science and member of the Center for Embedded Computer Systems, University of California, Irvine, CA 92697 USA (e-mail: givargis@ics.uci.edu).

Frank Vahid is with the Department of Computer Science and Engineering, University of California, Riverside, CA 92521 USA (e-mail: vahid@cs.ucr.edu). He is also a member of the Center for Embedded Computer Systems, UCI.

Jörg Henkel is with C&C Research Laboratories, NEC, 4 Independence Way, Princeton, NJ 08540 USA (email: henkel@ccl.nj.nec.com).

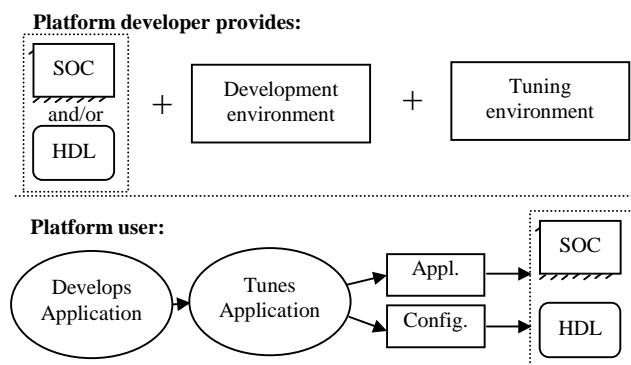


Fig. 1. SOC design flow.




Fig. 2. Y-chart approach for the design of configurable architectures.

approach for solving it. In Section 4, we give some experimental results. In Section 5, we state our conclusions.

II. PREVIOUS WORK

Much previous work has focused on power evaluation of SOC architectures at various levels of abstraction. Circuit-level approaches simulate the circuit at the transistor level while monitoring supply current [4][5]. Logic-level, or gate-level, approaches simulate a gate-level design, and calculate power by considering switching activity of nodes in the design [6][7], executing orders of magnitude faster than circuit-level approaches at the expense of some accuracy. RTL (Register-Transfer Level) power evaluation operates at an even higher-level of abstraction, modeling power consumption of more abstract circuit components, such as adders and multipliers etc. Simulation is performed at the RT-level and power is obtained by using these power models, also known as macro-models. RTL power evaluation, in some publications such as [8], is shown to be accurate to within 5% of actual power consumption. Behavioral-level approaches seek to estimate power of a behavioral HDL description before a synthesized design is obtained. An abstract notion of physical capacitance and switching activity is used. Switching is estimated using entropy from circuit input to circuit output by quadratic or exponential degradation [9][10].

Work has been done to evaluate power consumption of a particular type of core, like microprocessors. One approach, instruction-level power modeling, is proposed by [11]. Given a program execution trace, energy is computed as the sum of the energy consumed by each instruction that is executed, circuit state energy consumed when a particular instruction is followed by another, and energy consumed by other effects such as stalls and cache misses. This approach is sped up in [12], by deriving a shorter program trace that results in equal power dissipation when compared to the original trace. In [13] a mathematical generic power model for 32-bit microprocessors is proposed. The approach classifies the instruction set into classes, like branches. The model has been applied to various 32-bit processors. Other researchers have focused on fast system-level models for cache, memory and

bus power consumption [14][15][16], consisting mostly of simulators coupled with equations that compute power consumption as a function of usage/traffic and core parameters. Further approaches aim at estimating the power consumption of whole embedded systems. In [17], a cycle-accurate power simulation tool, for an embedded system using a strong ARM architecture as CPU, is introduced. The reported results are accurate within 5% compared to measurements conducted on a hardware board. A trace-based approach deploying a mix of analytical models (for instruction cache, data cache and main memory) and instruction set simulators (ISS) is introduced in [18].

Other system-level approaches have been proposed for application-driven design of core-based systems [19]. Here, given a fixed application, heuristics are used to determine cache size and organization in order to minimize cache misses while also minimizing chip area. Likewise in [20], an analytical approach is provided for exploring the on-chip memory architecture given a fixed application.

A methodology, closely related to ours, named SPADE (System level Performance Analysis and Design space Exploration), is proposed in [21]. This work defines a general scheme for the design of programmable architectures, referred to as the Y-chart and shown in Fig. 2. Here, target applications are mapped onto the architecture, and their performance is analyzed to obtain performance numbers. (The architecture, applications and performance numbers represent the dimensions along the Y shaped chart, hence the name Y-chart.) After analysis, the architecture or applications are tuned and the process is repeated until a desired system is obtained. In our work, we outline an approach to automate the exploration.

Previous work has focused on techniques that quickly and accurately simulate SOC architectures in order to obtain power and performance metrics. Our technique combines this work with an approach for efficiently exploring the configuration space of SOC architectures by pruning configurations that are guaranteed to be inferior to others already evaluated.

III. APPROACH OVERVIEW

We will next state the problem and outline our solution. Our solution will be given by first looking at an exhaustive method. Then we state the key observation that makes our approach more efficient, followed by our efficient solution.

A. Problem Formulation

We are given a system-on-a-chip architecture composed of numerous interconnected parameterized computational, communication, and memory elements. Each of these parameters can be assigned a value from a finite set of values. A complete assignment of values to all the parameters is a configuration. We are also given a parameterized system-level model of the SOC that when executed can yield the power and performance of the SOC for a configuration. Such parameterized simulation models have been outlined in [17][22][23]. The problem is to efficiently compute, with the

aid of a system-level model, the Pareto-optimal configurations, with respect to power and performance, for a fixed application executing on the SOC. In our problem, a configuration is Pareto-optimal if no other configuration has better power for a given performance.

The algorithms described in this paper can utilize any available power and performance measuring approach such as in-circuit emulation, gate-level simulation, RTL simulation or a system-level behavior approach. Hence, the exploration and evaluation approaches are absolutely orthogonal. In this work a system-level model is used for evaluation [22][23]. This system-level model can achieve simulation speeds that are 4 to 5 orders of magnitude faster than gate-level simulation, while maintaining performance/power estimation accuracy of 5% to 10% when compared to gate-level estimations. Here, for the CPU, cache and memory cores of the system an instruction-level power model is used that is based on [11]. For other cores, such as UART and CODECS, the core provider selects a set of appropriate instructions covering the possible actions of each core in the architecture. Then the provider performs gate-level power analysis to construct lookup tables for each instruction, and creates a system-level core model that utilizes the lookup-tables for power evaluation through an executable specification. The core user connects the system-level peripheral and CPU core models, executes the system and thus obtains power and performance data. The simulation speed of this system is approximately 134K processor-instructions per second running on an 800 MHz Pentium PC.

B. An Exhaustive Solution

We start by outlining an exhaustive algorithm to solve the exploration problem. In this exhaustive algorithm, first, power and performance are evaluated for all configurations. Then, configurations are sorted by non-increasing execution time (i.e., higher performance). Then, in the sorted order, a walk through the space is performed while all configurations that result in power consumption above the minimum seen thus far are eliminated. The remaining configurations are Pareto-optimal. The algorithm is given below.

Algorithm 1:

```

list compute_Pareto_configurations(space s) {
  list all, pareto;
  float min_power = 1e100; /* infinity */
  for each configuration c in space s {
    simulate_SOC(c); all.push(c);
  }
  all.sort( /* key is execution time */ );
  while( !all.empty() ) {
    c = all.pop();
    if( c.power < min_power ) {
      min_power = c.power; pareto.push(c);
    }
  }
  return pareto;
}

```

The problem with this approach is that the configuration space is likely to be very large, making the approach impractical in many cases. The exhaustive approach is practical when applied to a small subset of the solution space consisting of one or two varying parameters while all others held constant. We have

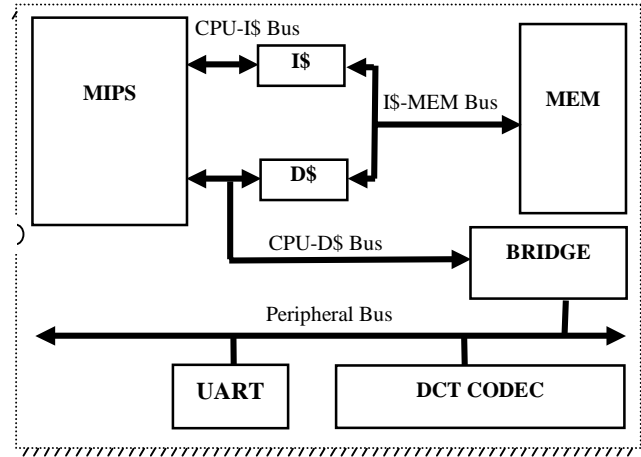



Fig. 3. Parameterized SOC platform.

found that many parameters in an SOC platform have little interdependency among each other. Two parameters are interdependent if changing the value of one of them impacts the optimal parameter value of the other.

C. Parameter Interdependency Model

We have used a directed graph model to capture the parameter interdependencies. Such a graph is constructed with its nodes representing parameters and edges representing interdependencies between parameters. Generally, a path from a node A to a node B indicates that the Pareto-optimal configurations of B should be calculated once the Pareto-optimal configurations of all the nodes from A to B, residing on the path, are calculated. During that calculation all other parameters not on the path are fixed to some arbitrary value. A path from a node A to a node B and back to A, which forms a cycle, indicates that the Pareto-optimal configurations of all the parameters on the cycle need to be calculated simultaneously. During that calculation all other parameters not on the path are fixed to some arbitrary value. The Pareto-optimal configurations of an isolated node is computed by setting all other parameters to some arbitrary value.

Fig. 3 shows the parameterized SOC architecture used in our experiments. The parameter description and interdependency graph of this architecture is depicted in Fig. 4. All interdependencies are manually determined. The cache size, associativity, and line parameters (B/C/D and E/F/G) are interdependent. The bus width and data encoding parameter pairs (H/I, J/K, L/M, N/O, P/Q, R/S, T/U, and V/W) are interdependent. The UART's buffer parameters (X/Y) are interdependent. Since the cache configuration will affect the amount of data transferred to and from the main memory bus, the I/D\$-memory-bus parameters are dependent on the cache parameters. To keep the graph from being cluttered, we have only shown a single edge from an arbitrary cache parameter to an arbitrary bus parameter (F to R for instance), since, by the transitive property, the bus parameters will be dependent on all the cache parameters. Finally, the MIPS voltage scale and DCT CODEC's parameters are independent of the remaining parameters. Our parameterized SOC architecture is composed



Node	Core	Parameter	Node	Core	Parameter
A	MIPS	Voltage scale	L	CPU-D\$-bus	Data bus width
B	I\$	Total size	M		Data bus code
C		Line size	N		Addr bus width
D		Associativity	O		Addr bus code
E	D\$	Total size	P	I/D\$-Mem-bus	Data bus width
F		Line size	Q		Data bus code
G		Associativity	R		Addr bus width
H	CPU-I\$-bus	Data bus width	S	Peripheral-bus	Addr bus code
I		Data bus code	T		Data bus width
J		Addr bus width	U		Data bus code
K		Addr bus code	V		Addr bus width
X	UART	Tx buffer size	W	DCT CODEC	Addr bus code
Y		Rx buffer size	Z		Pixel resolution

* Note that we have used double-arrowed lines in place of two single arrowed-lines for clarity.

Fig. 4. Target SOC interdependency graph and parameter descriptions.

of a total of 26 parameters. The total design space is composed of 1014 configurations. The parameters of our SOC architecture are hardware parameters and assumed to be set prior to fabrication. However, the approach given in this paper is general and applicable to other types of parameters, such as post fabrication parameters, compiler assigned parameters, and software tunable parameters.

We assume that the designer of the SOC platform determines the interdependencies among the parameters. However, our current research is focused on developing automated approaches for computing such interdependencies. These automated approaches are characterized as either being conservative or non-conservative. The conservative approaches start assuming that all parameters are interdependent and then remove dependencies if it's determined (through exhaustive simulation or sampling) that two parameters are clearly not interdependent. The non-conservative approaches start assuming that all parameters are independent and then introduce interdependencies if it's determined that two parameters are likely dependent. In both cases, the characterization of parameter interdependencies is a one-time effort.

D. An Efficient Exploration Algorithm

Given an interdependency graph, our algorithm works as shown below.

Algorithm 2:

```
list compute_Pareto_configurations_2(graph g) {
  list sub_graphs, pareto;
  sub_graphs = strongly_connected_components(g);
  // part 1
  for each sub-graph g in sub_graphs {
    pareto=compute_Pareto_configurations(g.space);
    eliminate configs. in g.space not in pareto;
  }
  // part 2
  while( !sub_graphs.size() != 1 ) {
    g1 = sub_graphs.pop_front();
    g2 = sub_graphs.pop_front();
    g = g1 union g2;
    sub_graphs.push_back(g);
    pareto=compute_Pareto_configurations(g.space);
  }
}
```

```
eliminate configs. in g.space not in pareto;
}
return pareto;
}
```

The algorithm can be broken down into two phases. The first phase performs a local search for Pareto-optimal configurations. The second phase iteratively expands the local search to discover global Pareto-optimal configurations.

Part 1 of our algorithm performs clustering of interdependent nodes in the graph. This is the same problem as finding strongly connected components of a graph. This can be computed by performing two depth-first searches in linear time. In addition, if two clusters are connected (but not strongly) then they are topologically ordered. Here, each cluster represents a disjoint sub-space of the overall configuration space. We use our exhaustive algorithm for calculating Pareto-optimal configurations for each of the clusters. Then, we restrict possible configurations of that cluster to the Pareto-optimal configurations only. This pruning is justified since if a configuration is not Pareto-optimal within a cluster, it cannot be part of a Pareto-optimal configuration for the entire configuration space. Conversely, if a configuration is Pareto-optimal within a cluster, it may or may not be Pareto-optimal given the entire configuration space, and thus must remain. Our exhaustive approach applied to clusters is usually feasible since these clusters represent only a small sub-space of the total configuration space. Nevertheless, heuristics such as probabilistic exploration techniques or genetic algorithms can be used to search within a cluster when the exhaustive method is too slow.

Part 2 of our algorithm combines pairs of clusters into a single cluster and computes Pareto-optimal configurations within it. It does this by defining the configuration space of this new cluster to be the cross product of the Pareto-optimal configurations of the two merged clusters. This procedure is repeated until all the clusters have been merged and a single cluster remains. The Pareto-optimal configurations within this last cluster represent Pareto-optimal configurations of the

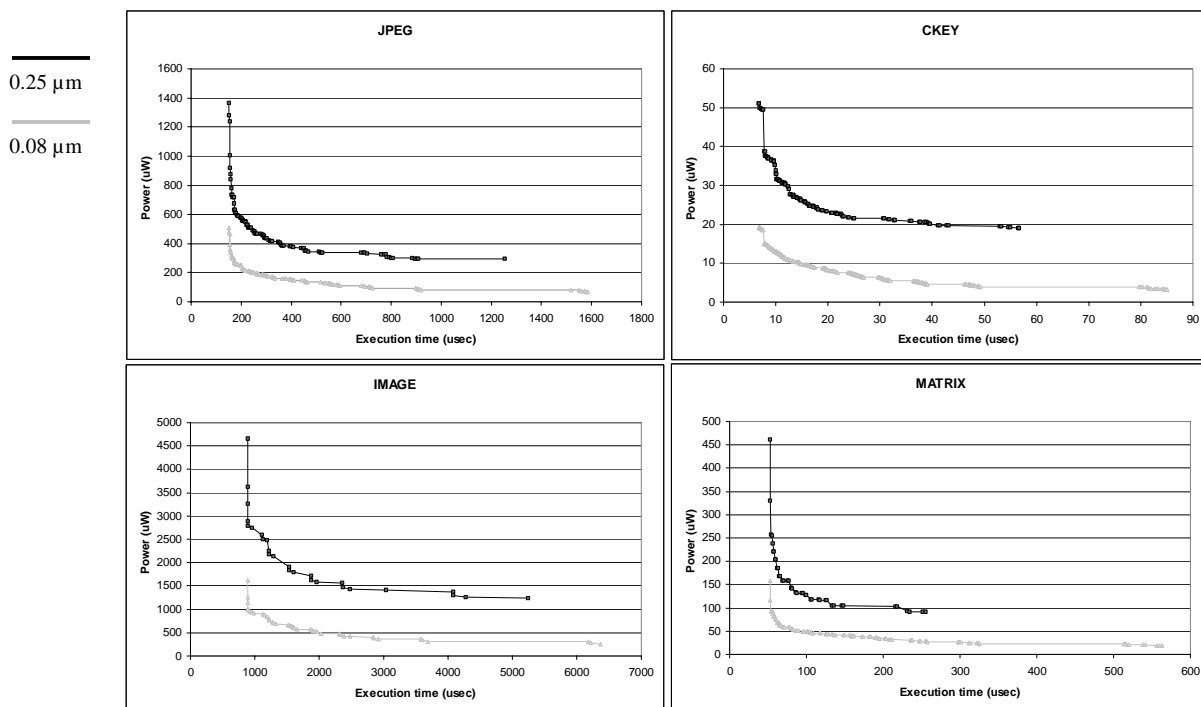


Fig. 5. Pareto-optimal configurations of *jpeg*, *ckey*, *image*, and *matrix* examples. Results shown for two technologies.

entire configuration space. Part 2 of the algorithm combines clusters in no particular order. However, by combining cluster pairs that result in the smallest merged configuration space, the total number of configurations that are examined is sometimes minimized. However, such optimization does not change the time complexity of the algorithm.

The worst case time complexity of the algorithm is bounded by $O((K + \log(K)) \times 2^{N/K})$, where K denotes the number of initial strongly connected components (i.e., clusters) computed in part 1, and N denotes the number of parameters. Here, the $2^{N/K}$ factor¹ bounds the running time of the exhaustive computations of the Pareto-optimal points. The K in the first factor is a bound on the number of times that the first part of the algorithm iterates, while the $\log(K)$ is a bound on the number of times the second part of the algorithm iterates². In the worst case, when $K=1$ (all parameters are interdependent) the running time is exponential, namely 2^N . In the best case, when $K=N$ (all parameters are independent) the running time is linear, namely N . For most practical cases the running time will be closer to the best case since the factor $2^{N/K}$ will decrease very rapidly as K increases.

IV. EXPERIMENTS

As stated in the previous section, we have a parameterized simulation model of the SOC architecture shown in Fig. 4. We

¹ For the purpose of time complexity analysis and presentation brevity, we assume each parameter can take on two values, however, in reality, parameters can be assigned larger number of values.

² The maximum number of nodes in a cluster is assumed to be N/K . This assumption holds in the worst and best case analysis. In any other case, the N/K ratio is the expected value but not the exact value.

explored the configuration space for 4 application programs and 2 different technologies representing 8 different examples. Our applications are named *jpeg*, *ckey*, *image*, and *matrix*. The *jpeg* application implements a JPEG compression algorithm using the on-chip DCT CODEC. The *ckey* application implements a complex chroma-key algorithm. The *image* application rotates an image by 90 degrees and converts the image colors to grayscale. The *matrix* application performs a matrix invert operation on a large matrix.

For each of the 4 examples, we simulated both a version of the SOC that used power models for an older technology (0.25 μm) and a version that used power models for a newer technology (0.08 μm).

Among the 8 runs, on the average, the time to explore and find Pareto-optimal configurations for any design was 34.5 minutes. On the average, our algorithm returned 93 Pareto-optimal configurations and simulated 6852 distinct configurations. Among the 8 runs, the pruning ratio was 99.7%, meaning 997 out of 1000 configurations were pruned. Our results are summarized in Table I. The power/performance tradeoffs of the Pareto-optimal configurations for all 4 applications are presented in Fig. 5. The average performance tradeoff is 8.0 times. The average power tradeoff is 5.0 times. The average energy tradeoff is 2.9 times. We make the following further observations based on the Pareto-optimal data that we gathered:

- The Pareto-optimal configurations are highly dependent on the applications. A configuration that resulted in the lowest power consumption while meeting some performance constraint for one application did not result in the lowest power consumption in the other

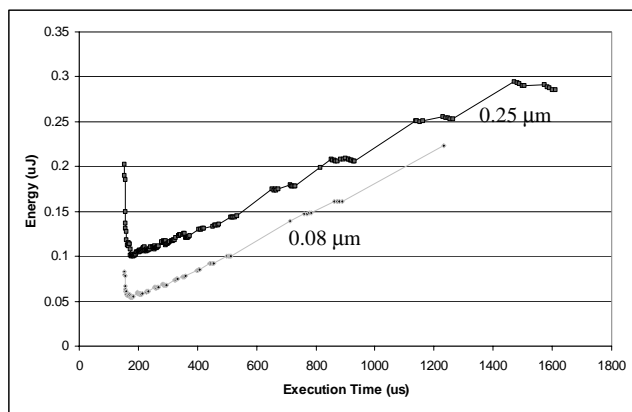


Fig. 6. Energy tradeoffs of the *jpeg* example.

applications.

- For the same application, most pairs of configurations, one from an old technology and one from the new technology that resulted in the equal performance or power, were different. This means that an optimal configuration in the old technology may or may not be an optimal configuration in the new technology.
- With respect to energy, the optimal configurations were those that lay in middle of the power/performance tradeoff curves. The energy plot for the *jpeg* example is given in Fig. 6. The configurations ordering is identical to those depicted in Fig. 5. Here, the rate of decrease in power consumption starts to become smaller compared to the rate of increase in execution time, resulting in a net increase in energy consumption.

Our exploration of the 4 examples revealed all the configurations of interest to a designer. The Pareto-optimal configurations were obtained in reasonable amount of time. Pareto-optimal configurations differed for different applications as well as technologies. An efficient simulation and exploration tool is necessary to achieve the best performance when mapping an application to a parameterized architecture.

V. CONCLUSION

We have presented an approach for efficiently finding all Pareto-optimal configurations of parameterized SOC architectures. Our approach relies on our knowledge about the interdependencies among parameters of the SOC. We use a directed graph to capture this interdependency and give

algorithms that search the configuration space, incrementally, and prune inferior configurations. Our experiments with several examples mapped onto our target SOC architecture demonstrate the feasibility of the approach.

REFERENCES

- [1] M.R. Stan, Wayne P. Burtleson. Bus-Invert Coding for Low Power I/O. IEEE Transactions on Very Large Scale Integration Systems, 1995.
- [2] A. Malik, B. Moyer, D. Cermak. A Programmable Unified Cache Architecture for Embedded Applications. International Conference on Compilers, Architecture, and Synthesis for Embedded Systems, 2000.
- [3] D.H. Albonesi. Selective Cache Ways: On-Demand Cache Resource Allocation. MICRO, 1999.
- [4] S.M. Kang. Accurate Simulation of Power Dissipation in VLSI Circuits. IEEE Journal of Solid-State Circuits, vol. CS21, no. 5, 1986.
- [5] G.Y. Yacoub, W.H. Ku. An Accurate Simulation Technique for Short-Circuit Power Dissipation Based on Current Component Isolation. International Symposium on Circuits and Systems, 1989.
- [6] R. Tjarnstrom. Power Dissipation Estimate by Switch Level Simulation. International Symposium on Circuits and Systems, 1989.
- [7] T.H. Krodel. PowerPlay - Fast Dynamic Power Evaluation Based on Logic Simulation. International Conference on Computer Aided Design, 1991.
- [8] E. Macii, M. Pedram. High-Level Power Modeling, Evaluation, and Optimization. IEEE Transactions on Computer Aided Design, vol. 17, no. 11, 1998.
- [9] D. Marculescu, R. Marculescu, M. Pedram. Information Theoretic Measures for Power Analysis. IEEE Transactions on Computer Aided Design, vol. 15, no. 6, 1996.
- [10] M. Nemani, F. Najm. Toward a High Level Power Evaluation Capability. IEEE Transactions on Computer Aided Design, vol. 15, no. 6, 1996.
- [11] V. Tiwari, S. Malik, A. Wolfe. Power Analysis of Embedded Software: A First Step Toward Software Power Minimization. IEEE Transactions on Very Large Scale Integration Systems, vol. 2, no. 4, 1994.
- [12] C.T. Hsieh, M. Pedram, H. Mehta, F. Rastgar. Profile Driven Program Synthesis for Evaluation of System Power Dissipation. Design Automation Conference, 1997.
- [13] C. Barndolese, W. Fornaciari, F. Salice, D. Sciuto. Energy Evaluation for 32-bit Microprocessor. International Workshop on Hardware/Software Co-Design, 2000.
- [14] R.J. Evans, P.D. Franzon. Energy Consumption Modeling and Optimization for SRAMs, IEEE Journal of Solid-State Circuits, vol. 30, no. 5, 1995.
- [15] T. Givargis and F. Vahid. Interface Exploration for Reduced Power in Core-Based Systems, International Symposium on System Synthesis, 1998.
- [16] T. Givargis, J. Henkel, and F. Vahid. Interface and Cache Power Exploration for Core-Based Embedded System Design. International Conference on Computer Aided Design, 1999.
- [17] T. Simunic, L. Benini, G. De Micheli. Cycle-accurate Evaluation of Energy Consumption in Embedded Systems. Design Automation Conference, 1999.
- [18] Y. Li, J. Henkel. A Framework for Estimating and Minimizing Energy Dissipation of Embedded HW/SW Systems. Design Automation Conference, 1998.
- [19] D. Kirovski, C. Lee, M. Potkonjak, W. Mangione-Smith. Application-

TABLE I
SUMMARY OF RESULTS

APPL.	EXPL. TIME (.08/.25μm)	Configs. Visited (.08/.25μm)	Pareto configs. (.08/.25μm)	Exe-time tradeoff (.08/.25μm)	Power tradeoff (.08/.25μm)	Energy tradeoff (.08/.25μm)
<i>jpeg</i>	6.3/25 min	2672/10611	97/118	10/8.3 x	7.4/4.7 x	2.7/3.4 x
<i>ckey</i>	24/68 min	3522/10050	125/213	13/8.4 x	6.0/2.7 x	2.6/3.6 x
<i>image</i>	9.5/47 min	1927/9531	26/42	7.1/5.9 x	3.1/3.8 x	2.1/2.6 x
<i>matrix</i>	15/54 min	2535/9245	41/92	11/4.9 x	8.3/5.03 x	2.3/2.2 x

Driven Synthesis of Core-Based Systems. International Conference on Computer Aided Design, 1997.

- [20] P.R. Panda, N.D. Dutt, A. Nicolau. Local memory Exploration and Optimization in Embedded Systems. IEEE Transactions on Computer Aided Design, vol. 18, no. 1, 1999.
- [21] P. Lieverse, P. van der Wolf, E. Deprettere, K. Vissers. A Methodology for Architecture Exploration of Heterogeneous Signal Processing Systems. IEEE Workshop on Signal Processing Systems, Taipei, October 1999.
- [22] T. Givargis, F. Vahid, J. Henkel, Instruction-based System-level Power Evaluation of System-on-a-chip Peripheral Cores. Asia and South Pacific Design Automation Conference, 2000.
- [23] UCR Dalton Group. The Platune Project. www.cs.ucr.edu/~dalton/Platune.



Dr. Tony Givargis received a B.S. in Computer Science from the University of California, Riverside in 1987. He received a Ph.D. degree from the University of California, Riverside in 2001, where he was a GAANN (Graduate Assistance in the Area of National Need) Fellow. Dr. Givargis received the department's best Thesis Award. He is currently an Assistance Professor in the Department of Information and Computer Science at the University of California, Irvine. His is also a member of the Center for Embedded

Computer Systems at UC Irvine. Dr. Givargis is a co-author of the textbook "Embedded System Design" (J. Wiley and Sons 2002). His research focuses on Platform based system design, real time resource management of embedded computing systems, and design space exploration.



Dr. Frank Vahid received a B.S. in Computer Engineering from the University of Illinois in 1988, and M.S. and Ph.D. degrees from the University of California, Irvine in 1990 and 1994, respectively, where he was an SRC Fellow. He is currently an Associate Professor in the Department of Computer Science and Engineering at the University of California, Riverside, where he received the Outstanding Teacher of the College of Engineering award in 1997. His is also a faculty member at the Center for Embedded Computer Systems at UC Irvine.

He was program and general chair for the IEEE/ACM International Symposium on System Synthesis in 1996 and 1997, respectively, and for the IEEE/ACM International Workshop on Hardware/Software Codesign in 1999 and 2000. He received the best paper award from IEEE Transactions on VLSI in 2000. He is co-author of the textbooks "Embedded System Design" (J. Wiley and Sons 2002) and "Specification and Design of Embedded Systems (Prentice Hall 1994). His current research focuses on architectures and design methods for low-power embedded systems, with an emphasis on tuning system-on-a-chip platforms to their executing programs.



Dr. Jörg Henkel is currently a Senior Research Staff Member at the Computers & Communications Research Laboratories, NEC USA, Princeton, NJ, where he is leading several projects in research and development of low power SOC architectures and EDA tools. Dr. Henkel is currently the General Co-Chair of the IEEE/ACM International Symposium on Hardware/Software Co-Design (2002) and has served for the same

event as a Program Co-Chair in 2001. In addition, Dr. Henkel currently serves as a Program Co-Chair for the IEEE International Workshop on Rapid System's Prototyping 2002. Dr. Henkel is involved in the following program committees: IEEE/ACM International Symposium on Hardware/Software Co-Design (Codes), IEEE/ACM Design Automation and Test in Europe Conference (DATE), IEEE International Symposium on Low Power Electronics and Design (ISLPED). Dr. Henkel is Senior Member of the IEEE.