

Optimizing Power Using Transformations

Anantha P. Chandrakasan, Miodrag Potkonjak, Renu Mehra, Jan Rabaey, and Robert W. Brodersen, *Fellow, IEEE*

Abstract—The increasing demand for portable computing has elevated power consumption to be one of the most critical design parameters. A high-level synthesis system, HYPER-LP, is presented for minimizing power consumption in application specific datapath intensive CMOS circuits using a variety of architectural and computational transformations. The synthesis environment consists of high-level estimation of power consumption, a library of transformation primitives, and heuristic/probabilistic optimization search mechanisms for fast and efficient scanning of the design space. Examples with varying degree of computational complexity and structures are optimized and synthesized using the HYPER-LP system. The results indicate that more than an order of magnitude reduction in power can be achieved over current-day design methodologies while maintaining the system throughput; in some cases this can be accomplished while preserving or reducing the implementation area.

I. INTRODUCTION

VERY LARGE SCALE INTEGRATION (VLSI) research and development efforts have focused primarily on optimizing speed to realize computationally intensive real-time tasks such as video compression and speech recognition. As a result, many systems have successfully integrated various complex signal processing modules to meet users' computation and entertainment demands. While these solutions have provided answers to the real-time problem, they have not addressed the rapidly increasing demand for portable operation. The strict limitation on power dissipation which portability imposes, must be met by the designer while still meeting ever higher computational requirements.

The desirability of portable operation of all types of electronic systems has become clear. It seems that by merely providing the same functionality of a wired system in an untethered implementation, a significant market advantage can be obtained for virtually any electronic function. Consumer electronics companies have made available portable televisions, camcorders, compact disk players, and, recently, hand held multimedia devices; cordless phones provide limited mobile access for voice communications with virtually unlimited access now being provided by cellular phones; laptops, notebooks, and palmtops are the fastest growing market segment of computers.

Manuscript received April 12, 1993; revised December 15, 1993 and August 19, 1994. This work was supported by ARPA. This paper was recommended by Associate Editor Kurt Keutzer.

A. P. Chandrakasan is with the Department of Electrical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139 USA.

M. Potkonjak was with the EECS Department, University of California at Berkeley, Berkeley, CA 97420 USA. He is now with Computer and Communication Research Laboratories, NEC USA, Princeton, NJ 08540 USA.

R. Mehra, J. Rabaey, and R. W. Brodersen are with the EECS Department, University of California at Berkeley, Berkeley, CA 97420 USA.
IEEE Log Number 9406361.

While wireless devices are rapidly making their way to the consumer market, a key design constraint for portable operation has not been considered, namely the total power consumption of the device. Reducing the total power consumption in such systems is important since it is desirable to maximize the run time with minimum requirements on size and weight allocated to batteries.

In this paper, we address the problem of automatically finding computational structures that result in the lowest power consumption for a specified throughput given a high-level algorithmic specification. The basic approach is to scan the design space utilizing various flowgraph transformations, high-level power estimation, and efficient heuristic/probabilistic search mechanisms. While transformations have been successfully applied in high-level synthesis with the goal of optimizing speed and/or area, they have not addressed the problem of minimizing power.

II. SOURCES OF POWER DISSIPATION

In CMOS technology, there are three sources of power dissipation: switching, short-circuit and leakage currents. The switching component, however, is the only one that can not be made negligible if proper design techniques are followed. The average switching power for a CMOS gate with a load capacitor, C_L , is given by [1]:

$$\begin{aligned} P_{\text{switching}} &= p_t \bullet \text{Energy per Transition} \bullet f = (p_t C_L \\ &= \bullet V_{dd}^2 \bullet f = C_{\text{avg}} \bullet V_{dd}^2 \bullet f \end{aligned} \quad (1)$$

where f is the clock frequency, and p_t is the probability of a power consuming transition ($0 \rightarrow 1$).

The energy consumed is therefore a quadratic function of the operating voltage, as verified in Fig. 1(a), which is an experimentally derived plot of the normalized energy versus V_{dd} . This dependence on supply voltage has been verified for a number of logic functions and logic styles [2]. The average capacitance switched, $C_{\text{avg}} = \sum p_t C_L$, for a uniformly distributed set of input values has been characterized for each logic and memory element in the cell library. Similarly, the delays for the various elements in the cell-library were characterized as a function of supply voltage. Fig. 1(b) shows experimental data of normalized delay versus V_{dd} for a typical CMOS gate implemented in 2.0 μm technology. Once again, the delay dependence on supply voltage was verified to be relatively independent of various logic functions and logic styles [2].

It is clear that operating at the lowest possible voltage is most desirable, however, this comes at the cost of increased delays and thus reduced throughput. However, by modifying the architecture through a variety of transformations the




Fig. 1. (a) Plot of normalized Energy versus V_{dd} for a typical circuit implemented in 2 μm CMOS technology. (b) Plot of normalized Delay versus V_{dd} for typical circuit implemented in 2 μm CMOS technology.

throughput can be regained, and thus a power savings can be accomplished while retaining the required functionality. It is also possible to reduce the power by choosing an architecture that minimizes the effective switched capacitance at a fixed voltage: through reductions in the number of operations, the interconnect capacitance, the glitching activity, and internal bit widths and using operations that require less energy per computation.

III. PREVIOUS WORK

In this section, a summary of the previous work done in low-power design and high-level design exploration is presented along with how this research relates to the previous efforts.

A. Low-Power Design Techniques

The quadratic influence of voltage on power, makes voltage scaling the most attractive scheme for power reduction. A “technology” based approach proposes choosing the power

supply voltage based on maintaining the speed performance for a given submicron technology [3]. By exploiting the relative independence of delay on supply voltage at high electric fields, the voltage can be dropped to some extent for a velocity-saturated device with very little penalty in speed performance. This was found to achieve a 60% reduction in power for a 3.3 V system when compared to a 5 V operation [4]. [2] presents an architecture based voltage scaling strategy that results in an optimal voltage for power that is much lower (in the 1–1.5 V range) than obtained from the technology based scaling. The idea is to maintain throughput at reduced supply voltages through hardware duplication or pipelining. By using parallel, identical units, the speed requirements on each unit are reduced, allowing for a reduction in voltage. This work presents an efficient transformation based high-level synthesis approach to explore the architecture based voltage scaling strategy [2].

Power reduction can also be achieved by minimizing the capacitance switched at all levels of the design. There are a number of options available in choosing and optimizing the basic circuit approach and topology for implementing various logic and arithmetic functions. A pass-transistor logic family was found to minimize physical capacitance when compared to a conventional CMOS logic family [5]. At another level, there are various topological choices for implementing a given function. For example, an adder can be implemented using ripple-carry or carry-lookahead approaches. The power trade-off between various types of adders and multipliers were investigated in [6] and they concluded that a carry-lookahead topology was the “best” after taking into account the speed-capacitance trade-off. Optimizing transistor sizing is yet another degree of freedom available in minimizing the energy. To minimize the parasitic capacitances, it is desirable to use minimal sized devices as much as possible, which also assists in minimizing the interconnect routing lengths and their associated parasitic capacitances [2]. To minimize the total switched capacitance in random logic modules, several logic synthesis optimization techniques have been proposed to lower the power dissipation [7].

B. Transformations in High-Level Synthesis

Over the last few years, several high-level synthesis systems have incorporated comprehensive sets of transformations, coupled with powerful optimization strategies. Example systems with elaborate applications of transformations are Flamel [8], SAW [9], SPAID [10], HYPER [11], and CATHEDRAL [12].

Among the set of transformations used by the Flamel design system are loop transformations, height reduction and constant propagation. SAW uses among other transformations in-line expansion, dead code elimination, four types of transformations for conditional statements and pipelining as supporting steps during the behavioral and structural partitioning. The SPAID system’s set of transformations includes retiming and pipelining, interleaving, substitution of multiplications with constants by addition and shifts and algebraic transformations. HYPER uses more than 20 different transformations whose application is supported by several probabilistic optimization

algorithms. CATHEDRAL emphasizes on support for loop transformations. All systems provide interactive frameworks where the designer explores the influence of transformation mechanism or optimization algorithms for a specific transformation.

The systems described above use transformations to optimize design parameters such as area and throughput. This work addresses the problem of how high-level flowgraph transformations can be used to reduce the power consumption of VLSI implementation.

C. Power Estimation

Most of the work done in power estimation falls under three categories: gate-level probabilistic estimation, switch-level estimation, and circuit-level estimation. The primary trade-off between these approaches is the computational complexity versus accuracy. Circuit-level approaches result in the most accurate estimates, while being the most computationally intensive.

Gate-level probabilistic approaches estimate the internal node activities of a network given the distribution of the input signals [13]–[15]. Once the signal probability for each node in the network is determined, the total average capacitance switched is then estimated as $\sum p_{i0}(1 - p_{i0})C_i$, where p_{i0} is the probability that node i will be in the ZERO state, and C_i is the physical capacitance associated with node i . The total power is then estimated as $C_{avg}^* V_{dd}^{2*} f_{clk}$.

An approach for estimating the power consumption in CMOS circuits using a switch-level simulator is presented in [16]. The basic idea is to monitor the number of times each node in the circuit transitions during the simulation period. C_{avg} is given by $\sum N_i / N C_i$, where N_i is the total number of power consuming transitions for node i , N is the number of simulation cycles, and C_i is the physical capacitance of node i . This approach (using the IRSIM simulator [17]) was used to estimate power at the layout level. The results from a few fabricated chips, indicate that the predicted power from IRSIM is within 30% of the measured power.

At the lowest level, power can be estimated using a circuit simulator such as SPICE. The design is described at a very low-level and the simulator accounts for short-circuit and leakage components of power. The major problem with this approach is the long simulation time which limits the number of patterns that can be applied.

The approaches mentioned above estimate power consumption from a low-level of abstraction. To use these approaches in a high-level synthesis framework, the high-level representation of the algorithm has to be mapped to a low-level description (gate or transistor level), which is very time consuming. Also, the estimation time for each new topology is too long to meaningfully explore many architectures. Hence, power must be estimated efficiently from a high level of abstraction. Recently, techniques to estimate the power consumption at the architecture level (after the flowgraph has been scheduled) have been developed [18]. In this paper, power is estimated from an algorithmic level so the design space can be quickly explored.

IV. USING TRANSFORMATIONS TO OPTIMIZE POWER

Transformations are changes to the computational structure in a manner that the input/output behavior is preserved. The number and type of computational modules, their interconnection and their sequencing of operation are optimized. The use of transformations makes it possible to explore a number of alternative architectures and to choose those which result in the lowest power. We will use the control-data flow graph format to represent computation [11], in which nodes represent operations, and edges represent data and control dependencies. Transformations have primarily been used until now to optimize either the implementation area or the system throughput. The goal here is to optimize a different function, namely the power dissipation of the final circuit while meeting the functional throughput of the system. Two key approaches are used to reduce power for a fixed throughput: reducing the supply voltage by utilizing speed-up transformations (Section IV-A) and reducing the effective capacitance being switched using a variety of transformations (Sections IV-B through IV-F).

A. Control Step Reduction to Lower Supply Voltage

This is probably the single most important type of transformations for power reduction. It is not only the most common type of transformation, but often has the strongest impact on power. The basic idea is to reduce the number of control steps, so that slower control clock cycles can be used for a fixed throughput, allowing for a reduction in supply voltage [2]. The reduction in control step requirements is most often possible due to the exploitation of concurrency. Many transformations profoundly affect the amount of concurrency in a given computation. This includes retiming/pipelining, algebraic transformations and loop transformations.

To illustrate the application of speed-up transformations to lower power, consider a first order IIR filter, as shown in Fig. 2(a), with a critical path of 2. Due to the recursive bottleneck [19] imposed by the filter structure, it is impossible to reduce the critical path using retiming or pipelining. Also, the simple structure does not provide opportunities for the application of algebraic transformations and applying a single transformation is not enough to reduce power in this example. Applying loop unrolling (Fig. 2(b)) does not change the effective critical path or capacitance and therefore the supply voltage cannot be altered; as a result, the power cannot be reduced. However, loop unrolling enables several other transformations (distributivity, constant propagation, and pipelining) which result in a significant reduction in power dissipation. After applying loop unrolling, distributivity and constant propagation in a systematic way, the output samples can be represented as:

$$Y_{N-1} = X_{N-1} + A^* Y_{N-2} \quad (2)$$

$$Y_N = X_N + A^* X_{N-1} + A^{2*} Y_{N-2} \quad (3)$$

The transformed solution has a critical path of 3 (Fig. 2(c)). However, pipelining can now be applied to this structure, reducing the critical path further to 2 cycles (Fig. 2(d)). Since the final transformed block is working at half the original sample

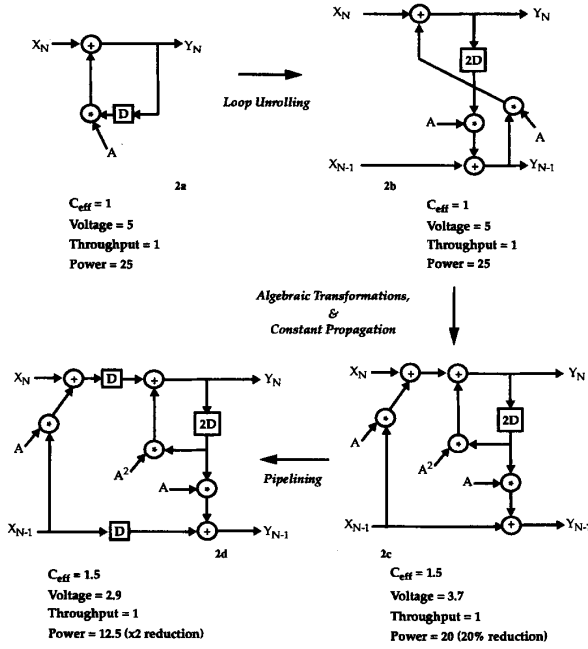


Fig. 2. Using speed-up transformations to reduce power.

rate (since we are processing 2 samples in parallel), and the critical path is same as the original datapath (2 control cycles), the supply voltage can be dropped to 2.9 V (the voltage at which the delays increase by a factor of 2, see Fig. 1(b)). However, note that the effective capacitance increases since the transformed graph requires three multiplications and three additions for processing two samples while the initial graph requires only one multiplication and one addition to process one sample, or effectively a 50% increase in capacitance. The reduction in supply voltage, however, more than compensates for the increase in capacitance resulting in an overall reduction of the power by a factor of 2 (due to the quadratic effect of voltage on power).

This simple example can be used to illustrate that optimizing for throughput will result in *different* solution than optimizing for power. For this example, arbitrary speedup can be achieved by continuing to apply loop unrolling, combined with other transformations (algebraic, constant propagation, and pipelining). The speedup grows linearly with the unrolling factor, as shown in Fig. 3. If the goal is to minimize power consumption while keeping the throughput fixed, the speedup can be used to drop the supply voltage. Unfortunately, the capacitance grows linearly with unrolling factor (since the number of operations per input sample increases) and soon limits the gains from reducing the supply voltage. This results in an “optimum” unrolling factor for power of 3, beyond which the power consumption starts to increase again.

This example brings out two very important points: First, the application of a particular transformation can have conflicting effects on the different components of power consumption. For example, a transformation can reduce the voltage component of power (through a reduction in the critical path)

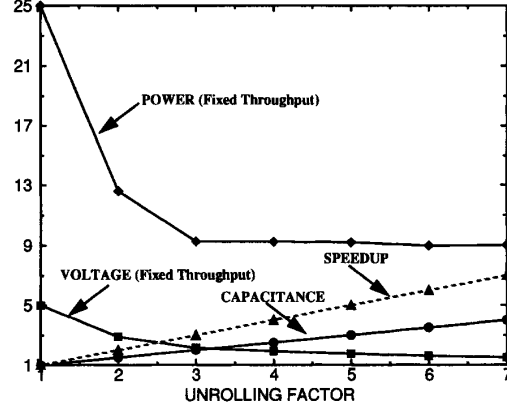


Fig. 3. Speed optimization *different* from power optimization.

while simultaneously increasing the capacitance component of power. Therefore, while speed-up transformations can be used to reduce power by allowing for reduced supply voltages, the “fastest” solution is often NOT the lowest power solution. The design environment described later can be used to automatically explore the design space using speed-up transformations (as described above) and find the “optimum” solution which trades off switched capacitance and voltage for a fixed throughput. Second, the application of transformations in a combined fashion almost always results in lower power consumption than the isolated application of a single transformation. In fact, it is often necessary to apply a transformation which may temporarily increase the power budget, in order to enable the application of transformations which will result in a more dramatic power reduction.

A simple first order IIR filter was presented above to demonstrate that the optimization for throughput is very different than optimization for power. Now consider a bigger and more widely used example, the DCT (Discrete Cosine Transform), to illustrate this point further. We will compare two implementations of the DCT: one proposed by Feig and Winograd [20] and the other, the direct maximally fast form. Feig’s DCT algorithm can be derived from the direct form using the exceptionally sophisticated application of common subexpression elimination/replication and algebraic rules. The direct form can be derived from the Feig’s DCT, by the simple application of the transformation set using the procedure for maximally fast implementation of linear computation [21]. While Feig’s DCT has a critical path of 11 cycles, the maximally fast DCT has a critical path of only 7 cycles. Therefore the V_{dd} can be reduced from 5 V to 3.25 V. While the reduction of the supply was relatively significant, the effective capacitance increased at a rate larger than an order of magnitude. Feig’s form has 54 multiplications and 462 additions while the maximally fast form has 4096 multiplication and 4096 additions. The power increases by a factor of more than 20, even after taking into account the voltage reduction due to throughput improvement.

The examples presented in this subsection clearly indicate that there is a sharp difference in the objective function for




Fig. 4. (a) Reducing capacitance while maintaining throughput. (b) Reducing capacitance at the expense of a higher supply voltage.

throughput and power. A number of key transformation for throughput enhancement have the effect of increasing the number of operations or increasing the interconnect area. For example, it has been shown that common subexpression can yield, for many classes of design implementation, arbitrarily high throughput at the expense of additional number of operations [21]. Similarly, it has been demonstrated that retiming for throughput often results in designs with exceptionally high interconnect requirements [22].

B. Operation Reduction

The most obvious approach to reduce the switched capacitance, is to reduce the number of operations (and hence the number of switching events) in the data control flow graph. While reducing the operation count typically has the effect of reducing the effective capacitance, the effect on critical path is case dependent. To illustrate this trade-off, consider evaluating second and third order polynomials. Computation of polynomials is very common in digital signal processing, and Horner's scheme (the final structure in our examples) is often suggested in filter design and FFT calculations when very few frequency components are needed [23].

First we will analyze the second order polynomial $X^2 + AX + B$. The left-hand side of Fig. 4(a) shows the straightforward implementation which requires two multiplications and two additions and has a critical path of 3 (assuming that each operation takes one control cycle). On the right side of Fig. 4(a), a transformed version having a different computational structure (obtained using algebraic transformations) is shown. The transformed graph has the same critical path as the initial solution and therefore the two solutions will have the same throughput at any given supply voltage. However, the transformed flowgraph has one less multiplication, and therefore has a lower capacitance and power.

Fig. 4(b) illustrates a situation where a significant reduction in the number of operations is achieved at the expense of a longer critical path. This example once again involves the computation of a polynomial, this time of the form $X^3 +$

$AX^2 + BX + C$. Again, by applying algebraic transformations we can transform the computation to the Horner's scheme. The number of multiplications reduces by two, resulting in a reduction of the effective capacitance. However, the critical path is increased from 4 to 5, dictating a higher supply voltage than in the initial flowgraph for the same computational throughput. Once again, this example shows that a transformation can have different effects on capacitance and voltage, making the associated power minimization task a difficult optimization problem.

Transformations which directly reduce the number of operations in a control data flow graph include: common subexpression elimination, manifest expression elimination, and distributivity.

C. Operation Substitution

Certain operations inherently require less energy per computation than other operations. A prime example of a transformation which explores this trade-off is strength reduction, often used in software compilers, in which multiplications are substituted by additions [24]. Although this situation is not as common as the ones presented in the previous section, sometimes it is possible to achieve significant savings using this type of tradeoff.

Unfortunately, this type of transformation often comes at the expense of an increase in the critical path length. This point is illustrated in Fig. 5, which shows the frequently used application of redundancy manipulation, distributivity and common subexpression in complex number multiplication. A_i and A_r are constants. While the second implementation has a lower effective capacitance, it has a longer critical path. Fixing the available time to be a constant, we see that the second implementation requires at least three control cycles (assuming each operation take one control cycle), while the first one requires only two. This implies that the voltage in the first implementation can be dropped lower than the voltage in the second implementation, since the same computational throughput can be met with a 50% slower clock rate.

A powerful transformation in this category is conversion of multiplications with constants into shift-add operations. Since multiplications with fixed coefficients are quite common in signal processing applications like DCT, filters, etc., the application scope of this transformation is large. Table I shows the power breakdown for a 11 tap FIR filter before and after this transformation. The power consumed by the execution units is reduced to one-eighth of the original power and the power consumed in the registers is also reduced. A small penalty is paid in the controller but there is a gain in the interconnect power due to a reduction in the length of the bus lines (since the final implementation has a smaller area), resulting in a total power savings of 62%.

D. Resource Utilization

It is often possible to reduce the required amount of hardware, while preserving the number of control steps [22]. This is possible because after certain transformations, operations are more uniformly distributed over the available time, resulting in a denser schedule. These transformations include retiming




Fig. 5. Trading multiplication for an addition.

TABLE I
BREAKDOWN OF POWER CONSUMED FOR A 11 TAP FILTER BEFORE AND AFTER CONSTANT MULTIPLICATIONS

Component	Before constant multiplications		After constant multiplications	
	Switched Capacitance, pF	% of total capacitance	Switched Capacitance, pF	% of total capacitance
Exu	739.65	64.80	93.07	21.63
Register/Clock	179.57	15.73	161.4	37.50
Control	65.45	5.73	83.79	19.47
Interconnect	156.69	13.74	92.10	21.40
Total	1141.36	100.00	430.36	100.00

for resource utilization, associativity, distributivity and commutativity. Fig. 6 shows the result of applying retiming on a second-order IIR filter. Both the transformed graphs, 1 and 2, are obtained from retiming and have a critical path of 3; however, the transformed graph 2 can be scheduled with only 2 multipliers while the first needs 4 multipliers (since all the four multiplications can be performed only in the third control step).

Given that there is a degree of freedom in choosing the amount of resources used for a fixed throughput (i.e, at a fixed supply voltage), the question then becomes is the solution with the minimal amount of resources, as chosen for the minimal area implementation, the “best” for low-power. On one hand, reducing the amount of resources can reduce the wiring capacitance since there are fewer interconnects and/or fewer functional elements and registers, which are obstacles during floorplanning and routing. However, the amount of multiplexers and control logic circuitry will typically increase with more time-sharing of resources. Therefore the optimization strategy (and hence the power estimation model) must take into account the trade-off between interconnect capacitance and control circuitry (which determines the effective capacitance being switched).

E. Reducing the Transition Activity

Designs using static CMOS logic can exhibit spurious transitions due to finite propagation delays from one logic

block to the next (called critical races or dynamic hazards), i.e., a node can have multiple transitions in a single clock cycle before settling to the correct logic value. The amount of extra transitions is a complex function of logic depth, statistics of the input patterns, and skew. To minimize the “extra” transitions and power in a design, it is important to balance all signal paths and reduce the logic depth. For example, consider the two implementations for adding four numbers shown in Fig. 7 (assuming a cascaded or nonpipelined implementation). Assume that all primary inputs arrive at the same time. Since there is a finite propagation delay through the first adder for the chained case, the second adder is computing with the new *C* input and the previous output of *A+B*. When the correct value of *A + B* finally propagates, the second adder recomputes the sum. Similarly, the third adder computes three times per cycle. In the tree implementation, however, the signal paths are more balanced and the amount of extra transitions is reduced. The capacitance switched for a chained implementation is a factor of 1.5 larger than the tree implementation for a four input addition and 2.5 larger for an eight input addition. The above simulations were done on layouts generated by the LagerIV silicon compiler [25] using the IRSIM [17] switch-level simulator over 1000 uncorrelated random input patterns.

The results presented above indicate that increasing the logic depth (through more cascading) will increase the capacitance due to glitching while reducing the logic depth will increase register power. Hence the decision to increase or decrease logic depth is based on a trade-off between glitching capacitance




Fig. 6. Retiming for minimizing bus capacitance: second-order IIR filter.




Fig. 7. Reducing the glitching activity.

versus register capacitance. Also note that reducing the logic depth can reduce the supply voltage while keeping throughput fixed.

F. Wordlength Reduction

The number of bits used strongly affects all key parameters of a design, including speed, area and power. It is desirable to minimize the number of bits during power optimization for at least three reasons:

- fewer bits result in fewer switching events and therefore lower capacitance.
- fewer bits imply that the functional operations can be done faster, and therefore the voltage can be reduced while keeping the throughput constant.
- fewer bits not only reduce the number of transfer lines, but also reduce the average interconnect length and capacitance.

The influence of various transformation on numerical stability (and therefore the required wordlength) varies a lot. While some transformations, for example retiming, pipelining and commutativity, do not affect wordlength, associativity and distributivity often have a dramatic influence [26]. In some cases, it is possible to reduce both the number of power expensive operations and the required wordlength. In other cases, however, a reduction in wordlength comes at the expense of an increased number of operations.

To illustrate the importance of wordlength optimization, consider the direct form and parallel form implementations of an eighth order Avenhaus band-pass filter (as shown in Fig. 8(a) and (b)). The critical path of the direct form, after multiplication to shift and addition substitution is 20 clock cycles while the critical path of the parallel form is 28 cycles, assuming that each operation, in both cases, takes one cycle. However, the numerical stability of the parallel form is significantly higher, resulting in wordlength requirements of only 11 bits, while the direct form solution requires 23 bits. The effective critical path for the direct form solution is 980 ns, while parallel form has a critical path of only 610 ns. Therefore, taking wordlength requirements into consideration, changes the situation from one where the critical path in the parallel form is almost 50% longer, to one where the critical path in the direct form is more than 50% longer. Similarly, a detailed analysis of the final implementation of the parallel filter is reduced by a factor of four for a given throughput. Note that the direct form can be transformed to the parallel form (or vice versa) using a specific ordered set of transformations [27].

The results from Section IV can be summarized in the following requirements for the application of transformations for power reduction:

- Efficient implementation of known and new transformations so that power is an explicit part of the cost function and the development of the cost function itself.
- Development of a transformation framework and search mechanism which will determine the order and extent to which the transformations are applied so that final result is globally optimal.

V. COST FUNCTION

The goal is to develop an objective function that is highly correlated to the final (and unknown) power dissipation of the circuit. The objective function should be very easy to compute since it has to be evaluated many times during the optimization process. A detailed estimation, while being accurate, will require hardware mapping and compilation steps to convert a flowgraph to layout, making it impractical during the optimization process. Hence a model correlated to the power must be developed strictly from the flowgraph level. This involves a statistical study of the effects of various high level parameters on interconnect capacitance, control capacitance, etc.

The power consumption of a circuit (represented in a control data flowgraph) is given by:

$$P_{\text{total}} = C_{\text{total}}^* V^2 * f_{\text{sampling}} \quad (4)$$




Fig. 8. Eighth order Avenhaus filter. (a) Direct form. (b) Parallel form.

The goal of power optimization in this work is to keep the throughput constant (i.e., f_{sampling} is fixed) by allowing the supply voltage to vary. For a fixed sample period, power optimization is equivalent to minimizing the total energy switched, $C_{\text{total}} V^2$, where V is appropriate voltage required to meet the throughput rate.

Fig. 9 shows the hardware model used in HYPER. A datapath module contains execution units, register files, multiplexors, and buffers. Data stored in the register files is read out to the execution unit whose result is written back into register files in the same datapath or in other datapaths. The register file may be preceded by multiplexors depending on whether the execution unit receives data from different sources at different control cycles in the sample period. The execution units on different datapaths communicate through global busses and the outputs of the execution units are tri-stated if two different units share a common global bus. A distributed control approach (which contains a global and local controller) is used to sequence through the operations in the control cycles. This section addresses

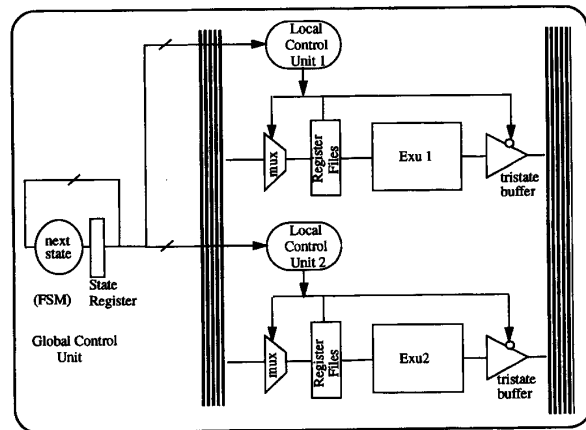


Fig. 9. Hardware model used in HYPER.

the estimation of the capacitance components and the supply voltage.

TABLE II
CAPACITANCE MODELS FOR SOME OF THE LIBRARY HARDWARE UNITS FOR A 1.2 μm CMOS TECHNOLOGY

Component	Capacitance Model (in fF)	Parameters
ripple adder	$-46 + 151 N$	N: bitwidth
carry select	$-158 + 214 N$	N: bitwidth
comparater	$181 N$	N:bitwidth
multiplier	$253 * N1 * N2$	N1,N2: input bitwidths
shifter	$N * 28.7 + \log(M+1) * (43.8 + 12.2N + .06 N^2 + 0.24MN - 0.18 SN)$	M: Maximum shift allowed N: bitwidth S: Shift
register file	$87 + 51 R + 35 N + 8 R N$	N: bitwidth R: number of register

A. Switched Capacitance Estimation

The total capacitance switched depends on four components:

$$C_{\text{total}} = C_{\text{exu}} + C_{\text{registers}} + C_{\text{interconnect}} + C_{\text{control}} \quad (5)$$

The capacitance estimation is built on top of an existing estimation routine in HYPER that determines bounds and activity of various execution, register and interconnect components as well as the implementation area [28], [29]. The details of the capacitance estimation routines are described below.

Execution Units: The capacitance switched by the execution units is estimated by multiplying (over all types of operations) the number of times the operation is performed per sample period and the average capacitance per access of the operation. The total capacitance is hence given by:

$$C_{\text{exu}} = \sum_{i=1}^{\text{numtypes}} N_i \cdot C_i \quad (6)$$

where *numtypes* is the total number of operation types, N_i the number the times the operation of type i is performed per sample period (or the activity), and C_i is average capacitance per execution of operation type i . The average capacitance per execution has been characterized for the various modules (through SPICE and IRSIM simulations, using models that were calibrated with results from experimental measurements) for a uniformly distributed set of inputs. In general the probabilities are not uniform, however, this assumption is made to simplify the cost evaluation. The capacitance values are parameterized as a function of bit-width and are accessed when computing the power contributed due to the execution units. For example, the capacitance per access for a ripple carry and a carry select adder as a function of the bitwidth (using a linear least squares fit) is shown in Fig. 10. Table II shows the high-level capacitance models for various modules in the hardware library.

The model presented above assumes the capacitance contribution due the execution units is relatively independent of allocation since the required number of operations must be performed within the sample period. For example, if

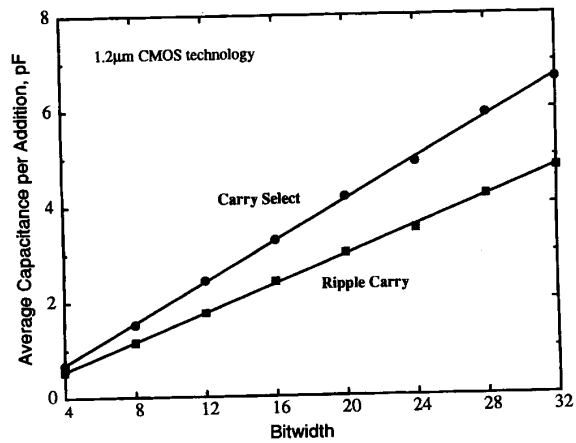


Fig. 10. Average capacitance per addition for two different modules.

a flowgraph has five additions, it is assumed that it does not matter whether the final implementation has one adder performing all the five additions within the sample period or has five adders each performing one addition within the sample period. This is not completely correct since the capacitance switched is a function of signal correlation. Techniques to model signal correlation at a high-level have been recently developed [18].

Registers: Registers are treated the same way as the execution units. The existing estimation program gives information about the total number of register accesses (read/write) within a given sampling period. This is essentially the "activity" of the registers. For the purposes of calculating the register energy, it is enough to know this "activity" and the actual number of physical registers is not required. The number of register accesses is multiplied with the average capacitance per register access to yield a register contribution given by:

$$C_{\text{registers}} = N_{\text{registers}} \cdot C_{\text{registers}} \quad (7)$$

While the total number of registers is not important in calculating the register switching capacitance, it will affect floorplanning and chip area and therefore the interconnect

capacitance. Gated clocks are used and the clock capacitance is taken into account during the characterization of the registers. Each register has a control slice that locally buffers the incoming clock.

Interconnect: While estimating the power consumed by the execution units and registers is quite simple and accurate, estimating the interconnect component is a very difficult and challenging task. Driven by yield, floorplanning and synthesis considerations for throughput and area optimization, several elaborate prediction models for total chip and interconnect area have been built and successfully used [30], [31]. However, high-level synthesis adds additional requirements on the prediction tools next to accuracy; during the optimization process in high level synthesis, it is necessary to estimate the final cost frequently and therefore computationally intensive models are prohibited, regardless of their precision.

Estimating the interconnect component of the final implementation should not only take into account the effects of a wide spectrum of high level synthesis tools, such as assignment, allocation, scheduling and partitioning into macro blocks, but also the effects of many low-level CAD tools, such as placement, floorplanning and global and detailed routing. Of course, an accurate model for such a complex system can be built only when a particular set of design tools is targeted. As mentioned earlier, we targeted the HYPER high-level synthesis tools and the Lager IV silicon assembler [25]. We used the scalable CMOS design rules provided by Mosis and targeted feature sizes of 1.2 μm and 2 μm .

The selection of this particular suite of design tools, enabled us to somewhat simplify the estimation process. We concentrated our attention only on the interblock (between macro blocks, e.g., different datapaths) routing capacitance. The effect of intrablock (between logic modules inside a datapath) routing capacitance is already taken into account during the calculation of execution units and register contribution, by incorporating an average loading capacitance (determined for the datapath compiler used in the Lager IV silicon assembler).

Building a model which will take into account the effects of the various tools mentioned above is a formidable task. An extensive experimental study, followed by in-depth statistical analysis and verification is the only viable solution which will satisfy the contradictory requirements of modeling a complex system, with high accuracy in a computationally efficient manner.

The model for interconnect capacitance was built using fifty examples which were mapped from their Silage descriptions to layout using the HYPER synthesis system and the Lager IV silicon assembler. The selected examples cover a wide variety of DSP applications including linear and nonlinear filters (including FIR, direct form, cascade, parallel, continuous function, ladder, wave digital, Rao-Kailath IIR, polynomial and homomorphic filters), fast transformations algorithms (FFT and DCT), several video and image processing algorithms and audio examples. Selecting the set of examples for building the model was guided by the goal of including as diverse and as typical examples as possible. While the smallest example had only 12 operations, the largest one had more than 400 operations. One half of examples was pipelined to various

extents, and a subset of the rest were transformed using several transformation in different orders. The examples cover a wide variety in the ratio of critical path to available time, amount of parallelism, types of used operations, level of multiplexing, size of the final implementations and other parameters.

After assembling the results for half the examples (for 25 examples), we started building the statistical model. It immediately became apparent that the best correlation is one between the total interconnect capacitance and the implementation area predicted by HYPER. It is widely recognized that the quality of the prediction model is inversely proportional to number of parameters used during the prediction model building [32]. The number of parameters is equal to the sum of the number of input variables in the model, and amount of data needed to describe the model. We built the interconnect capacitance model using only one variable as the predictor (estimated area of the chip) and the complexity of the curve which fits data was minimized as much as possible. Although it appeared that both the line and quadratic polynomial, and in particular the third order polynomial fit the data well, none of these models passed the strict statistical test for goodness of fit and resubstitution validation procedure. However, a piecewise linear least square fit showed both excellent accuracy and robustness.

The three segment piece-wise linear fit used to model the interconnect is characterized by four points, which are (2.95, 16.5), (10, 77.6), (23.5, 217) and (80, 1257). The model is valid from a predicted area of 2 mm^2 (which is equal to the size of a chip with one execution unit, with a few registers and interconnects) to a predicted area of 90 mm^2 (which is equivalent to a chip can that can accommodate more than 30 execution units, with more than hundred registers and multiplexers).

The measurement on all 50 examples, showed that the average error of the piecewise linear model is less than 18% percent, and the maximum error is smaller than 33%. Robustness of the model is illustrated by the fact that during 10 random resubstitution of 25 different values for prediction, the largest average error was below 20%, and the maximum error did not exceed 40%.

During the development of prediction tools, often consistency (for two different randomly selected instances, the one with lower predicted value indeed has lower measured value) is more important than accuracy. Consistency of our single prediction variable (predicted area) was, despite the very simple and robust model, higher than 96%. Fig. 11 shows both the measured (for all 50 examples) and predicted values of the interconnect capacitance as a function of the estimated chip area.

Once the physical interconnect capacitance is accurately estimated, it is easy to establish a good interconnect power consumption model. The interconnect capacitance component is then given by:

$$C_{\text{interconnect}} = \alpha * C_{\text{total}} / N \quad (8)$$

where α is the average activity (the total number of interconnect accesses multiplied by an average signal transition probability), C_{total} is the total estimated interconnect capacitance of the chip and N is an estimate of the number of physical

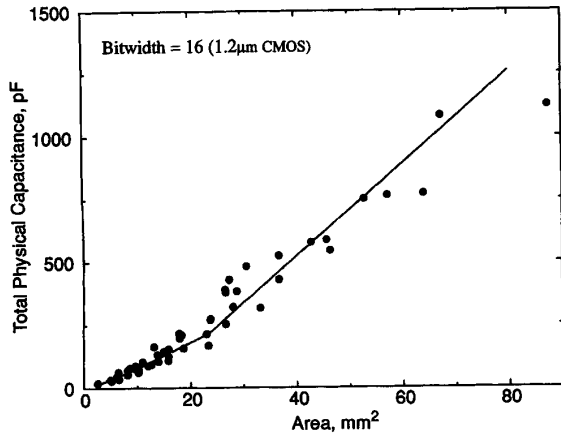


Fig. 11. Statistical estimation of interconnect capacitance.

interconnects (after bus-merging). The HYPER system provides accurate estimates of the number of interconnects and activity.

The interconnect model was built using the automatic place and route features of the Lager IV placement and layout tool (Flint). In the future, a model should be developed based on hand-optimized floorplanning so the user can get feedback on the efficiency of their layout.

Control Logic: High level estimation of the power consumed by the controller, like that of interconnect, is complicated because the control is not defined until the hardware mapping process. Neither the number of control blocks nor their function/size is known at the estimation stage, making it impossible to estimate any properties of the control theoretically. After scheduling, the control is defined and optimized by the *hardware mapper* and further by the *logic synthesis process* before mapping to layout. Like interconnect, therefore, the control needs to be estimated statistically.

The distributed control model used by HYPER is especially suited for low power. The control model incorporates a central finite state machine from which the state information is distributed to local controllers. Control signals (e.g., LOAD for the register file) for the datapath are generated in the local controllers. Bus capacitance on the control lines is reduced by placing the local controllers close to the datapaths. Thus, only the global state lines need to be distributed globally across the chip.

A large amount of data from several different DSP algorithms were analyzed in order to model the power consumed by the global and local controllers. Both initial and transformed versions were used to get a complete description of the sample space. Fully pipelined versions were not considered since there is only one control cycle per sample period, eliminating the need for a controller. Each of the benchmark examples were mapped to SDL (structural description) using HYPER and then into layout using the LAGER IV silicon compiler. IRSIM was used to extract the switching capacitances required to build the statistical model. This process has been automated for characterizing new libraries or the effects of new tools.




Fig. 12. Total capacitance switched versus number of states for the global controller.

Global Control Model: The amount of capacitance switched in the global controller was found by simulating each benchmark circuit for a whole sample period after the initial conditions were set up. The capacitance switched per sample period in the global controller was found to be directly related to the number of states, N_{states} , and is given for a 1.2 μm technology by:

$$C_{FSM} = \alpha_1 N_{states} + \alpha_2 \quad (9)$$

For a 1.2 μm technology, α_1 is 4.9 fF and α_2 is 22.1 fF. Fig. 12 shows the total capacitance switched by the global controller as a function of the number of states. Note that the number of states not only determines the number of cycles the FSM goes through, but also the number of output bits it drives. Therefore, the total number of transitions and hence the capacitance switched, is strongly dependent on the number of states.

Local Control Model: Since there are several local controllers in each design, the local controllers account for a larger percentage of the total capacitance than the global controller. A set of fifty examples were used to build the capacitance model for the local controller. Table III shows the data obtained for four different examples. The capacitance switched in the local controllers were determined by applying the state input sequence to each of the local controllers and simulating using IRSIM over a whole sample period. Correlations between the capacitance switched and several parameters like the number of states in the control, the number of outputs of the controller, and the total number of transitions on the output control signals were measured. The capacitance of the control was found to be highly correlated to the number of transitions on the output control signals. Note that, in this case, unlike the global controllers, the number of states gives no information about the number of transitions on the output nodes which depends on the glue-logic to be implemented. The transitions must therefore be separately accounted for. Using statistical tools to fit the data to a polynomial function, it was found that the total capacitance switched, in one sample period, for any local

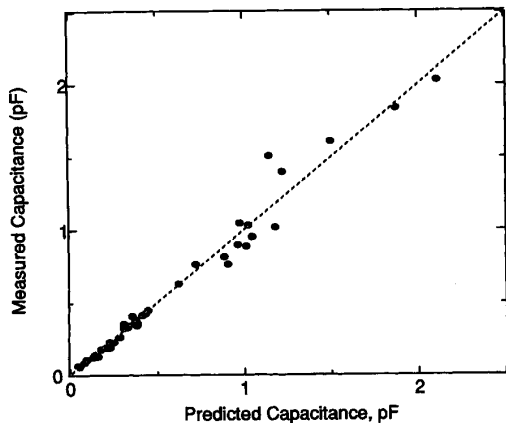


Fig. 13. Total capacitance switched versus predicted capacitance for different local controller.

controller is a linear function of the number of transitions, the number of states and the bus factor given by (10).

$$C_{lc} = \beta_0 + \beta_1 N_{trans} + \beta_2 N_{states} + \beta_3 B_f \quad (10)$$

where N_{trans} is the number of transitions, N_{states} is the number of states, B_f is the bus factor (explained below), and C_{lc} is the capacitance switched in any local controller in one sample period. B_f represents the activity factor on busses and is defined as the ratio of the number of bus accesses to the number of busses. It is a measure of the average number of times busses are accessed. It indicates the number of multiplexors, and each multiplexor requires control signals from the controller. For a 1.2μ technology, β_0 , β_1 , β_2 and β_3 are 72, 0.15, 8.3, and 0.55 respectively. Fig. 13 shows the correlation between the actual and predicted capacitance switched per sample period for several different examples.

This model, though accurate to within 20% (as indicated by resubstitution validation procedure) is not sufficient since it assumes that the number of transitions is known. The number of transitions depends on assignment, scheduling, optimizations performed by the hardware mapper and the logic optimization tool (misII), the standard cell library used, the amount of glitching, and the statistics of the inputs. It is impossible to determine the combined effects of all these elements apriori. A statistical model relating the number of transitions to high level parameters was therefore derived. The important high level parameters that affected the number of transitions are:

- Size/complexity of the graph, i.e., the number of nodes and the number of edges. This is because the number of reads from and write to registers is dependant on the number of edges and nodes in the graph respectively. Control signals need to be generated for every read and write process.
- The number of execution units times the number of states. This is because each execution unit receives the clock and clock inverse every control cycle. In the hardware model used, the clock inverse is generated in the local controllers and fed to the execution units.

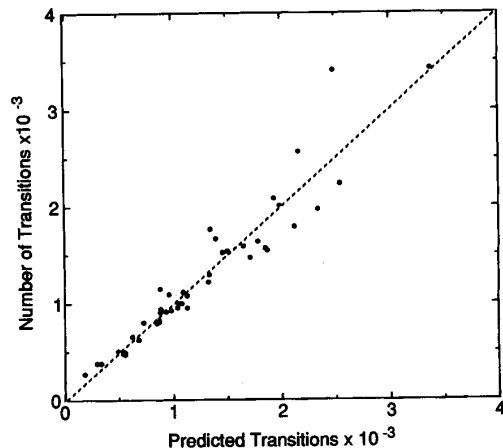


Fig. 14. Actual number of transitions versus predicted transitions for various local controllers.

The HYPER high level synthesis system provides a lower bound on the number of execution units within an accuracy of 10% [28]. The lower bound on the busses tracks the actual number of busses closely and the maximum number of busses tracks the total number of bus accesses. The numbers predicted by HYPER were therefore used in building the correlation model. The fit function obtained for the total number of output transitions on local controllers based on the three high level parameters mentioned above is given in (11).

$$N_{trans} = \gamma_1 + \gamma_2(N_{nodes} + N_{edges}) + \gamma_3(S \times N_{Exu}) \quad (11)$$

where N_{trans} is the number of transitions on the outputs of the local controllers, S is the number of control cycles per sample period, N_{edges} and N_{nodes} are the number of edges and nodes respectively in the CDFG and N_{Exu} is an estimate for the total number of execution units. For a 1.2μ technology γ_1 , γ_2 and γ_3 are 178.7, 7.2, and 2.0, respectively. Fig. 14 shows the total number of transitions in one sample period for the different examples versus the fit function for the transitions based on high level parameters.

This statistical estimate takes into account the effects of logic synthesis and optimization. Effects of undetermined factors such as glitching are also included. Each datapath element in the cell library has a built in control-slice to locally buffer the incoming control signals. The internal gate and routing capacitance is taken into account when characterizing the leafcells. For example, the multiplexor select signal for a 16-bit datapath is buffered in the control slice. Therefore, the controller only drives a single (typically minimum sized) gate. This information was used to determine the loading capacitance during simulations.

B. Supply Voltage Estimation

The power supply voltage at which the flowgraph implementation will meet the timing constraints is estimated. The initial flowgraph which meets the timing constraints is typically assumed to be operating at a supply voltage of 5 V with a critical path of $T_{initial}$ (the initial voltage will be

TABLE III
RESULTS FROM VARIOUS LOCAL CONTROLLER UNITS OF THE SELECTED SAMPLE SET

Name (example)	Name (exu-unit)	# of outputs	Total Switched Capacitance (pF)	# of transistors	length, width (in λ)	Area in μ^2 .
Wavelet (xlarge) 3 States 2 Inputs	add160	5	0.928	14	195, 94	6598.80
	add1610	5	0.943	14	219, 86	6780.24
	add1611	5	0.721	14	76, 46	1258.56
	shr160	11	1.787	34	227, 166	13565.52
	shr1610	9	1.509	30	211, 174	13217.04
	shr1611	5	0.943	14	211, 86	6532.56
	sub160	5	0.691	14	203, 94	6869.52
	sub1610	7	1.129	22	203, 118	8623.44
	sub1611	7	0.980	18	219, 118	9303.12
	transfer160	5	0.845	18	203, 78	5700.24
	transfer161	5	0.930	14	203, 126	9208.08
	transfer162	5	0.808	18	195, 70	4914.00
IIR 22 States 5 Inputs	adder120	53	82.738	544	451, 1504	244189.44
	adder121	47	76.981	528	427, 1232	189383.04
	shr120	30	42.553	296	403, 1048	152043.84
	shr121	25	55.556	302	363, 976	127543.68
	shr122	24	43.253	312	403, 1096	159007.68
	sub120	59	49.855	364	411, 1232	182286.72
	sub121	45	55.556	380	427, 1088	167247.36
	transfer	19	14.706	64	307, 504	55702.08
Discrete Cosine transform; 36 states; 6inputs	adder80GL	51	128.8085	644	539, 1726	334913.04
	shr80GL	35	101.891	428	475, 1566	267786.00
	subtractor80GL	53	145.349	704	547, 1806	355637.52
Avenhaus-Ladder Filter 79 States 7inputs	adder140GL	7	27.68	174	251,344	31083.84
	adder1410GL	7	24.916	170	243,256	22394.88
	adder1411GL	7	13.410	174	235,232	19627.20
	adder1412GL	7	27.635	174	251,280	25300.80
	shr140GL	12	97.841	420	395,1128	160401.60
	shr141GL	13	119.757	560	355,1008	128822.40
	shr142GL	12	102.044	426	403,1000	145080.00
	shr143GL	12	143.883	542	395,1120	159264.00
	subtractor140GL	25	97.104	430	87,1000	139320.00
	subtractor141GL	17	78.879	406	371,952	127149.12
	subtractor142GL	17	81.948	426	355,896	114508.80
	subtractor143GL	11	71.729	374	339,744	90797.76

lower if $T_{\text{initial}} < T_{\text{sampling}}$, where T_{sampling} is the throughput constraint). After each move, the critical path is re-estimated, and the new supply voltage at which the transformed flow-graph still meets the time constraint, T_{sampling} , is determined. For example, if the initial solution requires 10 control steps (and let's assume that this is the same as the sampling period) running at a supply voltage of 5 V, then a transformed solution

that requires only 5 control steps can run at a supply voltage of 2.9 V (where the delay increases by a factor of 2, Fig. 1(b)) while meeting the same constraints as the initial graph.

A model for delay as a function of V_{dd} is derived from the curve shown in Fig. 1(b). Let the speedup of a transformed solution be defined as:

$$\text{Speedup} = T_{\text{sampling}}/T_{\text{criticalpath}} \quad (12)$$

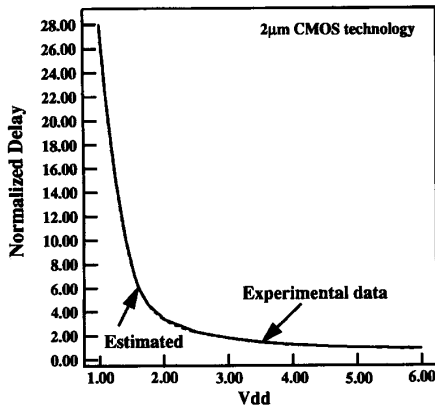

Fig. 15. Accuracy of V_{dd} estimation.

Fig. 16. Overview of power estimation.

where $T_{criticalpath}$ is the critical path of the transformed solution.

Since, the major power reduction during CDFG optimization using transformations is attributed to a reduction of the supply voltage and since the supply voltage has to be constantly evaluated (every time the objective function is called), it is important to model delay- V_{dd} relationship using an accurate and computationally efficient procedure. This relationship (of delay- V_{dd}) was modeled using Neville's algorithm for rational function interpolation and extrapolation [33]. Neville's algorithm provides an indirect way for constructing a polynomial of degree $N-1$ so that all of the used points are exactly matched.

Fig. 15 also shows the accuracy of the interpolated data (plotted for speedup values ranging from 0.86 to 26 with increments of 0.1) when compared to the experimental data. Fig. 16 shows an overview of the power estimation routine. It takes a control dataflow graph as input and computes the power using the existing estimation routines (critical path, bounds on resources and activity) along with the newly developed capacitance and voltage estimation routines (as discussed in this section).

VI. OPTIMIZATION ALGORITHM

The algorithm development for power minimization using transformations was greatly driven by several considerations and constraints. The basic strategy involved:

- Efficient tailoring of existing library transformations for power optimization.
- Fully utilizing the information about the design solution space obtained through extensive experimentation and theoretical insights (so that the algorithm can quickly come out of deep local minima).
- Efficient use of computer resources (CPU time and memory) so that large real-life examples can be addressed.
- Developing the modular software so that future enhancements or changes can easily be incorporated.

The computational complexity analysis of the power minimization problem showed that even highly simplified versions of the power optimization tasks using transformations are NP-complete. For example, we have shown that retiming for power minimization is a NP-complete problem. Computational complexity of retiming for power optimization is particularly interesting and somewhat surprising, since retiming for critical path reduction has several algorithms of polynomial-time complexity [34].

The computational complexity of the power minimization problem implies that it is very unlikely, even when the set of applied transformations is restricted, that a polynomial time optimal algorithm can be designed. Two widely used alternatives for design of high quality suboptimal optimization algorithms are probabilistic and heuristic algorithms. Both heuristic and probabilistic algorithms have several distinctive advantages over each other. While the most important advantage of heuristic algorithms is a shorter run time, probabilistic algorithms are more robust and have stronger mechanisms for escaping local minima.

An extremely wide spectrum of probabilistic algorithms are commonly used in various CAD and optimization problems. Among them, the most popular and successful are simulated annealing, simulated evolution, genetic algorithms, and various neural network algorithms (e.g., Boltzmann machine). Although recently there has been a considerable effort in analyzing these algorithms and the type of problems they are best suited for (for example a deep relationship between simulated annealing and solution space with fractal topology have been verified both experimentally and theoretically [35]), algorithm selection for the task at hand is still mainly an experimental and intuitive art.

In order to satisfy all major considerations for the power minimization problem, we decided to use a combination of heuristic and probabilistic algorithms so that we can leverage on the advantages of both approaches. Before we get into the algorithmic details, we will briefly outline the set of transformations used.

The transformation mechanism is based on two types of moves, global and local. While global moves optimize the whole DCFG simultaneously, local moves involve applying a transformation only on one or very few nodes in the DCFG. The most important advantage of global moves is, of course, a higher optimization effect; the advantages of local moves is their simplicity and small computational cost. We used the following global transformations i) retiming and pipelining for critical path reduction, ii) associativity, iii) constant elimination, and iv) loop unrolling. In the library

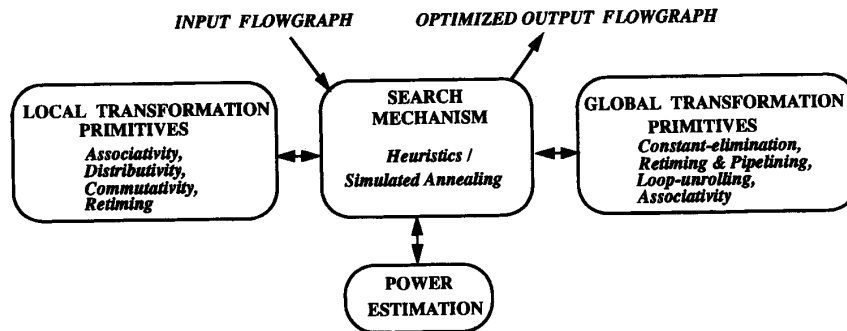


Fig. 17. Overview of the HYPER-LP system.

of local moves we have implemented three transformations: associativity (generalized to include properties of inverse elements as introduced in [22]), commutativity, and local retiming. All global moves minimize the critical path using polynomial optimal algorithms. Their use is motivated by the need for shorter run-times. For example, the global pipelining move for a seventh order IIR filter with 33 nodes took 1.8 CPU seconds on a SPARC2 workstation. Extensive experimentation showed that although the best solution in semi-exhaustive searches is rarely one with the minimum critical path, it is very often topologically very close (a small number of moves is needed to reach it starting from a solution with the shortest critical path).

The Leiserson-Saxe algorithm for critical path minimization using retiming is used. The algorithm is modified in such a way that it automatically introduces the optimal number of pipeline stages needed to minimize the cost function for power. Pipelining is accomplished by allowing the simultaneous addition of the equal number of delay elements on all inputs or all outputs but not both. The Leiserson-Saxe algorithm is strictly based on minimizing the critical path and inherently does not optimize the capacitance being switched for a given level of pipelining. The modified algorithm determines the level of pipelining where the power is minimized and this is used as a "good" starting point for optimization using local moves (as will be discussed below). The global move for associativity involves the minimization of the critical path using dynamic programming. Loop unrolling does not involve any optimization, instead it enables transformations which reveal large amount of concurrency for other transformations. For each local move, we defined an inverse local move which undoes the effect of the initial move.

As previously mentioned, the algorithm for power minimization using transformations has both heuristic and probabilistic components. While the heuristic part uses global transformations, the probabilistic component uses local moves. The heuristic part applies global transformations one at the time in order to provide good starting points for the application of the probabilistic algorithm. The probabilistic algorithm conducts a search in a vicinity of the solution provided by the heuristic part with the goal of minimizing the effective capacitance (trading-off control, interconnect and activity). The underlying search mechanism of the probabilistic part

is simulated annealing with the goal that local minima can be escaped by using uphill moves and the effects of various transformations can be efficiently combined so as to minimize the capacitance component of power.

Initially constant propagation is applied to reduce both the number of operations and the critical path, followed by the modified Leiserson-Saxe algorithm and the dynamic programming algorithm for the minimization of the critical path using associativity. This step is followed by simulated annealing which tries to improve the initial solution by applying local moves probabilistically. For simulated annealing we used the most popular set of parameters [36]. For example, we used a geometric cooling schedule with a stopping criteria which terminates the probabilistic search when no improvement was observed on three consecutive temperatures. After each move, a list of all possible moves is generated for each local transformation. The transformation and particular move to be applied is then selected randomly. In addition to the above algorithm, loop-unrolling can be used as preprocessing step.

HYPER-LP also allows optimization using a user specified sub-set of transformations (for example, optimizing with only pipelining). Fig. 17 shows an overview of the HYPER-LP system.

VII. EXAMPLES AND RESULTS

The techniques described in the previous sections will now be applied to "real-world" examples to demonstrate that a significant improvement in power can be achieved. Three examples of distinctly different computational structures are presented in this section and are optimized for power using the HYPER-LP system. Multiplications with constants were converted to shift-add operations. The improvement from this transformation was not taken into account since this is a standard transformation performed for area or throughput optimization. The run time for all examples were less than 1 minute on a SPARC2 workstation. The bitwidth requirement was established using the HYPER simulator tool [11].

A. Example #1: Wavelet Filter

The first example is a wavelet filter, used in speech and video applications. The implementation contains high-pass and low-pass filters that are realized as 14th order FIR modules

TABLE IV
RESULTS FOR THE WAVELET CHIP (1.2 μ m CMOS TECHNOLOGY)

Version	f_{samling}	# of Control Cycles / Clock Period	Supply Voltage	Total Average Capacitance	Power, mW	Area, mm ²
Initial	1.5Mhz	22 / 30ns	5V	2870pF	107mW	8.5 mm ²
Optimized for Power	1.5Mhz	3 / 220ns	1.5V	1735pF	5.8mW	62.9mm ²

using a wordlength of 16-bits (determined by high-level simulation). This example is representative of a wide class of important signal processing applications such as FFT, DCT (Discrete Cosine Transform), matrix multiplication, cyclic convolution and correlation that have no feedback loops in the signal flowgraph. For this class of applications, retiming and pipelining provides an efficient and straight forward way to reduce critical paths (and thus voltage) while preserving throughput and the number of operations.

The initial reference design is a fully time-multiplexed area-optimized realization of the filter, with minimal amount of resources (1 adder, 1 subtractor, 1 shifter, and 4 global busses). The number of global busses were minimized using global bus-merging (i.e busses that are never accessed simultaneously can be collapsed) with the goal of minimizing the area. This solution has a critical path of 22 clock cycles and can be clocked at a maximum frequency of 1.5 Mhz (since the critical path for this 16-bit datapath is around 30 ns). Using HYPER-LP, it was found that retiming by itself was sufficient to reduce the critical path of this design to three control cycles. Since the throughput requirement is 1.5 Mhz, the clock period can be made approximately 7 times as long ($= 22/3$) while meeting the timing constraint and therefore the supply voltage can be reduced to 1.5 V (Fig. 1(b)).

Since the amount of resources increases significantly (since the allocation is done with available time = critical path = 3 cycles), we expect the chip area and hence the interconnect component of power to increase. This is seen from Fig. 18, which shows a plot of the average capacitance per bit obtained from layout extraction as a function of the bus number for the initial 22 cycle implementation (which contained four busses) and for the final three cycle implementation (which contained 25 busses). The figure also shows the capacitance obtained from the estimation of interconnect using the model presented in Section VI-A-3.

The global bus capacitance switched increased by a factor of 3 from 400 pF to 1200 pF. However, the total capacitance switched was actually *smaller* than the minimal area version! There are several factors contributing to this decrease in capacitance:

- The amount of multiplexor, tri-state and control switches (which contributed to 45% of the total power of the minimal area design) were reduced significantly.
- The initial design uses bus-merging to significantly reduce the interconnect area. However, this does not come for free as the loading for the tristate buffers driving the bus increases. Also all of the multiplexors whose inputs are

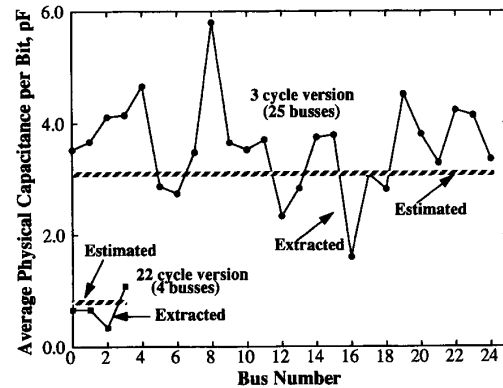


Fig. 18. Bus capacitance for the wavelet filter (estimated and extracted).

connected to the shared bus will switch every time the bus value transitions, resulting in extra power. The problem is not as severe in the final design since there is little room for bus-sharing since there are fewer control cycles.

- The average transition activity (which is not modeled by the estimator presented in Section VI) and hence the capacitance per access for the logic modules (adder, subtractor, and shifter) was significantly lower in the final version. The reason is that in the initial version, since the modules were fully time-multiplexed, the inputs to the modules were coming from different sources (and hence uncorrelated) and had a higher probability of transitioning. In the final version (in which many units were allocated), there are longer periods of inactivity during which the inputs to the modules do not change and hence the units do not switch, resulting in lower activity.

The total power was reduced by a factor of 18 (these results are obtained from the IRSIM switch level simulator using a linear RC model using real input data). Table IV shows the statistics of the initial and final designs. Note that area has been traded for lower power.

B. Example #2: IIR Filter

The second example is a 16 bit seventh order IIR filter with a fourth order equalizer designed using the filter design program Filsyn [37]. This example is representative of the largest class of examples where feedback is an inherent but not particularly limiting part of the computation structure.

Unlike the previous example, when the transformation set of the HYPER-LP system was limited only to retiming, no

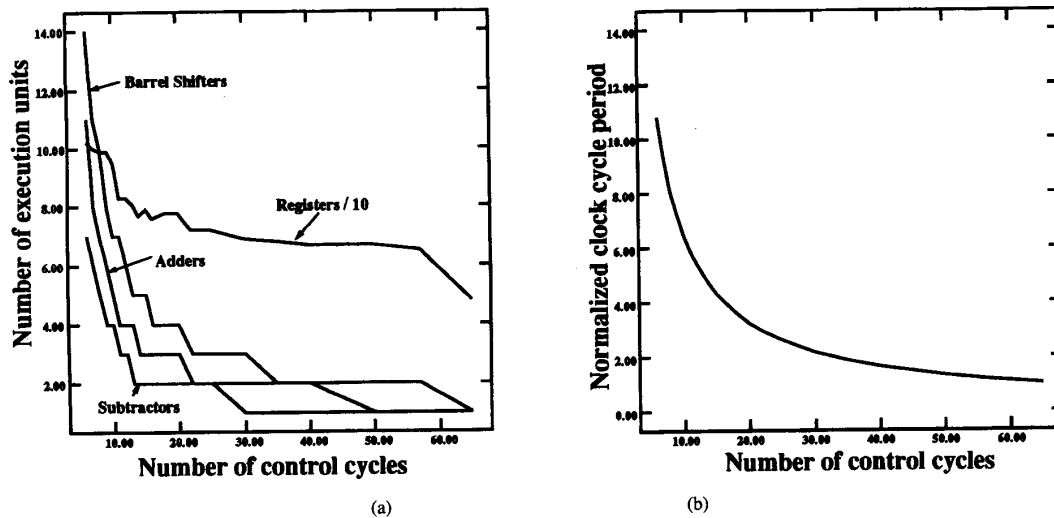


Fig. 19. Plot of number of units (a) and clock cycle period (b) versus number of control cycles for the IIR example.

reduction in the power consumption was observed. However, when the transformation set was enhanced to include both retiming and pipelining, the tool was very effective and achieved a reduction in the number of control steps used from 65 to 6 cycles. The implementation area increased as the number of control cycles decreased. Fig. 19(a) shows a plot of the number of execution units versus the number of control cycles. The number of registers increased from an initial number of 48 (for 65 control cycles) to 102 (for 6 control cycles). This is a typical example where the area is traded for lower power. It is immediately evident that the price paid for reducing critical path is the increased registers and routing overhead. Fig. 19(b) shows a plot of the normalized clock cycle period (for a fixed throughput) as a function of number of control cycles. We see that reducing the number of control cycles through transformations allows for an increase in clock cycle period or equivalently a reduction in supply voltage. However, at very low voltages, the overhead due to routing increases at a very fast rate and the power starts to increase with further reduction in supply voltage. The power reduced approximately by a factor of 8 as a result of dropping the voltage from 5 V to approximately 1.5 V. Fig. 20 shows a plot of Power versus V_{dd} for a fixed throughput. Note that the low-power solution is somewhere between the minimal area solution and the maximum throughput solution.

C. Example #3: Volterra Filter

The third example is a second-order Volterra filter. This is a particularly challenging example since pipelining cannot be applied due to a recursive bottleneck imposed by long feedback loops (optimizing with pipelining as the only transformation results in power reduction only by a factor of 1.5). For this case, loop unrolling, retiming and associativity transformations proved to be important in alleviating the recursive bottleneck, allowing for a reduction in critical path and voltage.

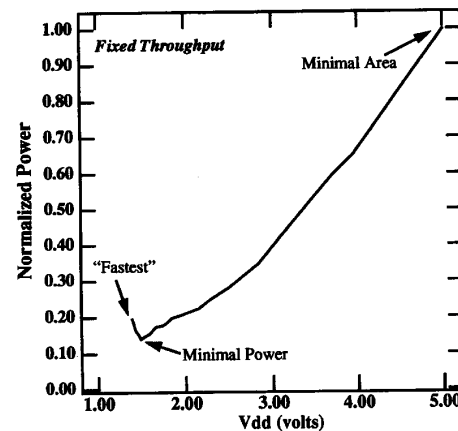


Fig. 20. Power versus V_{dd} for the IIR example.

Fig. 21 shows the final layouts of optimizing the Volterra filter for power as well as for area while keeping the throughput fixed. This final implementation has approximately nine times better power characteristic and the same area as the initial implementation.

Table V shows a summary of power improvement after applying transformations relative to initial solutions that met the required throughput constraint at 5 V for the representative examples described in this section. The results indicate that a large reduction in power consumption is possible (at the expense of area) compared to present-day methodologies. Also interesting was the fact that the optimal final supply voltage for all the examples was much lower than existing standards and was around 1.5 V. A couple of overall comments on the results obtained:

- It is possible, after optimization, that the final solution will end up with a supply voltage that is non-standard (e.g., 2 V). In this case, high efficiency (> 90%) switch-

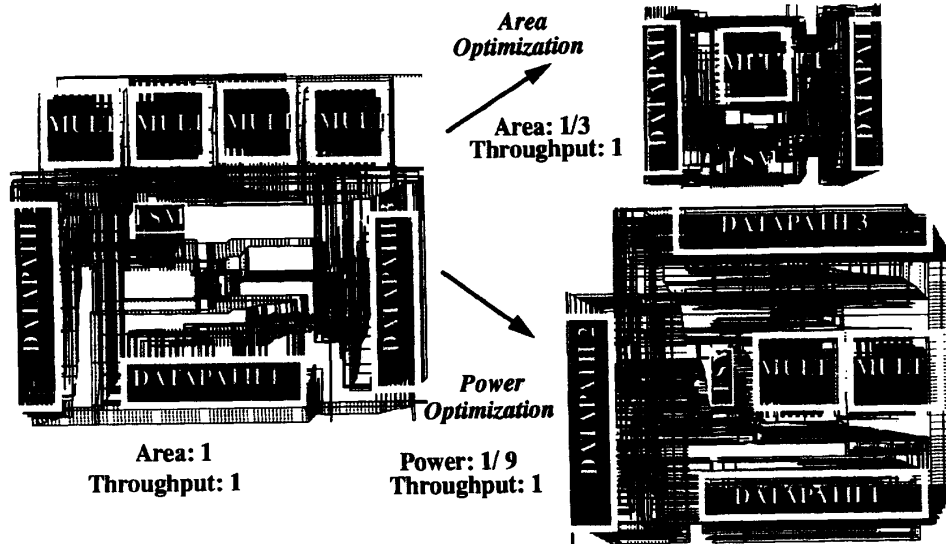


Fig. 21. Optimizing for area and power result in different solutions.

ing regulators (as opposed to linear regulators) can be used to provide the required supply voltage and voltage conversion circuits can be used at the system level to communicate between subsystems that are working at different voltages. Voltage level conversion circuitry can be implemented with very low area and power overhead [38]. The use of multiple supply voltages (in which each part of the system operates at its own “optimum” voltage) can result in significant system power reduction compared to a solution which uses a single power supply voltage.

- Reducing the supply voltage raises some issues on noise immunity. Inductive noise is a result of the voltage drops ($L \cdot di/dt$) across the inductors in the supply lines. Ignoring the threshold effects on current, we see that the di/dt scales as V^3 (even faster when threshold effects are taken into account), and hence drastically reducing the voltage noise. However, as mentioned earlier, reducing the voltage will require more parallelism (to meet throughput requirements) and will hence increase the number of simultaneously switching elements. Even so, the overall noise reduces in approximately a quadratic fashion with respect to the supply voltage. Also, using advanced packaging can reduce the lead inductances dramatically over conventional packaging (e.g., a MCM package has less than 2 nH of lead inductance while a typical wirebond pin grid array package has about 10 nH). A similar analysis shows that the resistive noise (proportional to the current) scales at least linearly.

VIII. CONCLUSION

Power minimization is becoming a very important problem with the increasing demand for portable computing. An automated high-level synthesis system, HYPER-LP, was presented for optimizing power consumption in algorithm specific datapath intensive circuits using a variety of architectural and

TABLE V
SUMMARY OF RESULTS

Example	Power Reduction	Area Increase
Wavelet Filter	18	7.4
IIR Filter	8	6.4
Volterra Filter	9	1

computational transformations. The synthesis approach consisted of applying transformation primitives in a well defined manner in conjunction with efficient high-level estimation of power consumption. The experimental results indicate that more than an order of magnitude reduction in power is possible over current-day design methodologies (that tend to minimize area for a given throughput) while maintaining the system throughput. It was found that the optimal supply voltage for minimizing power was much lower than existing standards (present-day 5 V and emerging 3.3 V) and was around 1.5 V for most of the examples investigated. This work has addressed some key problems in the automated design of low-power systems, and provides a good starting point for addressing other research problems like detailed power estimation, module selection, partitioning, and scheduling for power optimization.

ACKNOWLEDGMENT

We would like to thank several members of the HYPER team for their support and suggestions; including S.-H. Huang for coding and supporting the library of local transformation primitives, O. Bentz for supporting the back-end HYPER code (HardwareMapper), I. Verbaauwhede for many useful suggestions and the design and high-level simulation of the wavelet filter, and S. Wu for the functional simulation of the wavelet filter at the IRSIM level.

REFERENCES

- [1] N. Weste and K. Eshragian, *Principles of CMOS VLSI Design: A Systems Perspective*. Reading, MA: Addison-Wesley, 1988.
- [2] A. Chandrakasan, S. Sheng, and R. Brodersen, "Low-power CMOS digital design," *IEEE J. Solid-State Circuits*, pp. 473-484, Apr. 1992.
- [3] M. Kakumu and M. Kinugawa, "Power-supply voltage impact on circuit performance for half and lower submicrometer CMOS LSI," *IEEE Trans. Electron Devices*, vol. 37, pp. 1902-1908, Aug. 1990.
- [4] D. Dahle, "Designing high performance systems to run from 3.3 V or lower sources," *Silicon Valley Personal Computer Conf.*, 1991, pp. 685-691.
- [5] K. Yano *et al.*, "A 3.8 ns CMOS 16x16 multiplier using complementary pass transistor logic," *IEEE J. Solid-State Circuits*, pp. 388-395, Apr. 1990.
- [6] T. Callaway and E. Swartzlander, "Optimizing arithmetic elements for signal processing," *VLSI Signal Process. Wkshp.*, pp. 91-100, 1992.
- [7] A. Shen, A. Ghosh, S. Devedas, and K. Keutzer, "On average power dissipation and random pattern testability of CMOS combinational logic networks," *Proc. IEEE ICCAD*, 1992, pp. 402-407.
- [8] H. Trickey, "Flamel: A high-level hardware compiler," *IEEE Trans. Computer-Aided Design*, vol. 6, no. 2, pp. 259-269, 1987.
- [9] R. A. Walker and D. E. Thomas, "Behavioral transformation for algorithmic level IC design" *IEEE Trans. Computer-Aided Design*, vol. 8, no. 10, pp. 1115-1127, 1989.
- [10] B. S. Haroun and M. I. Elmasry, "Architectural synthesis for DSP silicon compilers," *IEEE Trans. Computer-Aided Design for IC*, vol. 8, no. 4, pp. 431-447, 1989.
- [11] J. Rabaey, C. Chu, P. Hoang, and M. Potkonjak, "Fast prototyping of data path intensive architecture," *IEEE Design and Test*, vol. 8, no. 2, pp. 40-51, 1991.
- [12] F. H. M. Franssen, F. Balasa, M. F. X. B. van Swaaij, F. V. M. Catthoor, and H. J. De Man, "Modeling multidimensional data and control flow," *IEEE Trans. VLSI Syst.*, vol. 1, no. 3, pp. 319-327, 1993.
- [13] M. Cirit, "Estimating dynamic power consumption of CMOS circuits," *IEEE Int. Conf. Computer Aided Design*, 1987, pp. 534-537.
- [14] F. Najm, "Transition density, a stochastic measure of activities in digital circuits," *DAC*, pp. 644-649, 1991.
- [15] A. Ghosh, S. Devedas, K. Keutzer, and J. White, "Estimation of average switching activity," in *Proc. DAC*, 1992, pp. 253-259.
- [16] N. Kimura and J. Tsujimoto, "Calculation of total dynamic current of VLSI using a switch level timing simulator (RSIM-FX)," *CICC*, 1991.
- [17] A. Salz and M. Horowitz, "IRSIM: An incremental MOS switch-level simulator," *Proc. 26th ACM/IEEE Design Automation Conf.*, June 1989, pp. 173-178.
- [18] P. E. Landman and J. M. Rabaey, "Power estimation for high level synthesis," *Proc. Euro. Conf. Design Automation '93*, Paris, France, Feb. 1993, pp. 361-366.
- [19] K. K. Parhi, "Algorithm transformation techniques for concurrent processors," *Proc. IEEE*, vol. 77, no. 12, pp. 1879-1895, 1989.
- [20] E. Feig and S. Winograd, "Fast algorithms for discrete cosine transform," *IEEE Trans. Signal Process.*, vol. 40, no. 9, pp. 2174-2193, 1992.
- [21] M. Potkonjak and J. Rabaey, "Maximally fast and arbitrarily fast implementation of linear computations," *IEEE ICCAD*, pp. 304-308.
- [22] M. Potkonjak and J. Rabaey, "Optimizing the resource utilization using transformations," *Proc. IEEE ICCAD Conf.*, Nov. 1991, pp. 88-91.
- [23] G. Goertzel, "An algorithm for the evaluation of finite trigonometric series," *Amer. Math. Monthly*, Vol. 65, No. 1, pp. 34-35, 1968.
- [24] A. V. Aho and J. D. Ullman, *Principles of Compiler Design*. Reading, MA: Addison-Wesley, 1977.
- [25] C. S. Shung *et al.*, "An Integrated CAD System for Algorithm-Specific IC Design," *IEEE Trans. Computer-Aided Design*, vol. 10, no. 6, pp. 447-463, April 1991.
- [26] D. Goldberg, "What every computer scientist should know about floating-point arithmetic," *ACM Computing Surveys*, vol. 23, no. 1, pp. 5-48, 1991.
- [27] M. B. Srivastava and M. Potkonjak, "Transforming linear systems for joint latency and throughput optimization," *EDAC-94*, pp. 267-271, 1994.
- [28] J. Rabaey and M. Potkonjak, "Estimating implementation bounds for real time application specific circuits," *Euro. Solid-State Circuits Conf.*, Milano, Italy, Sept. 11-13, 1991, pp. 201-204.
- [29] D. Schultz, "The influence of hardware mapping on high-level synthesis," *U.C. Berkeley M.S. Report*, ERL M92/ 54.
- [30] F. J. Kurdahi and A. C. Parker, "Techniques and area estimation of VLSI layouts," *IEEE Trans. Computer-Aided Design*, vol. 8, no. 1, pp. 81-92, Jan 1989.
- [31] W. R. Heller *et al.*, "Prediction of wiring space requirements for LSI," *IEEE/ACM Proc. 14th Design Automation Conf.*, pp. 20-22, 1977.
- [32] L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone, *Classification and regression trees*. Monterey, CA: Wadsworth & Brooks, 1984.
- [33] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical recipes in C*. Cambridge: Cambridge Univ., 1988.
- [34] C. E. Leiserson and J. B. Saxe, "Retiming synchronous circuitry," *Algorithmica*, vol. 6, no. 1, pp. 5-35, 1991.
- [35] G. B. Sorkin, "Efficient simulated annealing on fractal energy landscapes," *Algorithmica*, vol. 6, no. 3, pp. 367-418, 1991.
- [36] *Algorithmica, Special Issue, Simulated Annealing*. A. Sangiovanni-Vincentelli, Ed., vol. 6, no. 3, pp. 295-478, 1991.
- [37] Comsat General Integrated Systems, *Super-Filsyn Users Manual*, Mar. 1982.
- [38] A. P. Chandrakasan, A. Burstein, and R. W. Brodersen, "A low-power chipset for portable multimedia applications," *IEEE Int. Solid-State Circuits Conf.*, pp. 82-83, Feb. 1994.



Anantha P. Chandrakasan received the B.S. degree (with highest honors) and the M.S. and Ph.D. degrees in electrical engineering and computer sciences from the University of California, Berkeley, in 1989, 1990, and 1994 respectively.

He received a Fellowship in Engineering from IBM in 1993. Since September 1994, he has been an Assistant Professor of Electrical Engineering at the Massachusetts Institute of Technology, Cambridge. He has been awarded the Analog Devices Career Development Chair. His research interests include

design methodologies for low power portable systems, computer-aided design tools for VLSI, and system level integration.




Miodrag Potkonjak received the B.S. and M.S. degrees in electrical engineering from the University of Belgrade, Yugoslavia, and the Ph.D. degree in electrical engineering and computer science from the University of California, Berkeley.

He has been a research staff member with Computer and Communication Research Laboratories, NEC USA, Princeton, NJ since 1991. His main research interests include computer-aided analysis, synthesis and evaluation of system level designs, high level synthesis, and interaction between high performance application specific computations and communications.



Renu Mehra (S'93) received the B.Tech. degree in electrical engineering from I.I.T., Kanpur, India, in 1991, and the M.S. degree from the University of California, Berkeley, in 1993.


She is currently working toward the Ph.D. degree at the University of California, Berkeley. Her research interests include design automation techniques for low-power systems, high-level synthesis, and DSP system design.



Jan Rabaey received the E.E. and Ph.D. degrees (*summa cum laude*) in applied sciences in 1978 and 1983, respectively, from the Katholieke Universiteit Leuven, Belgium.

In 1987, he joined the faculty of the University of California, Berkeley, where he is now a Professor. His current research interests are the study of architectures, the computer-aided analysis and the automated design of digital signal-processing systems. He has authored or co-authored more than 100 papers in the area of signal processing.

Dr. Rabaey received the 1985 IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN Best Paper Award (Circuits and Systems Society) in 1986. In 1989, he received the Presidential Young Investigator award. He also was a recipient of the Analog Devices Career Development Professorship award. He has served as Associate Editor of the IEEE JOURNAL OF SOLID-STATE CIRCUITS and he is/has been on the program committee of the ISSCC, EDAC, ICCD, ICCAD, High-Level Synthesis, and VLSI Signal Processing conferences.



Robert W. Brodersen (F'82) received the B.S. degree in electrical engineering and mathematics from California State Polytechnic University, Cambridge. He received the Engineers, M.S., and Ph.D. degrees (the latter in 1972) from the Massachusetts Institute of Technology.

From 1972 to 1976, he was with the Central Research Laboratory at Texas Instruments. In 1976, he joined the Electrical Engineering and Computer Science Faculty of the University of California at Berkeley, where he is now a Professor. In addition

to teaching, he is involved in research on new applications of integrated circuits, which is focused in the areas of low-power design and wireless communications, and the CAD tools necessary to support this activity.

Dr. Brodersen has won best paper awards for a number of journal and conference papers in the areas of integrated circuit design, CAD, and communications, including the W. G. Baker Award for the Best IEEE Publication in all areas, in 1979. In 1983 he was co-recipient of the IEEE Morris Liebmann Award. In 1986 he received the Technical Achievement Award of the IEEE Circuits and Systems Society, and in 1991 from the Signal Processing Society. In 1988 he was elected a member of the National Academy of Engineering.