

Reliable and Energy-Efficient Digital Signal Processing

Naresh Shanbhag

Department of ECE & Coordinated Science Laboratory
University of Illinois at Urbana-Champaign
1308 West Main Street, Urbana, IL-61801
shanbhag@uiuc.edu

ABSTRACT

This paper provides an overview of *algorithmic noise-tolerance* (ANT) for designing reliable and energy-efficient digital signal processing systems. Techniques such as prediction-based, error cancellation-based, and reduced precision redundancy based ANT are discussed. Average energy-savings range from 67% to 71% over conventional systems. Fluid IP core generators are proposed as a means of encapsulating the benefits of an ANT-based low-power design methodology. CAD issues resident in such a methodology are also discussed.

Categories and Subject Descriptors

B.7 [Hardware]: Integrated Circuits

General Terms

DESIGN

Keywords

reliability, energy-efficiency, low-power, deep submicron, noise, noise-tolerance, broadband, communications

1. INTRODUCTION

Next generation wireless communication standards such as wireless LAN (IEEE 802.11b and a), access (IEEE 802.16) and 3G (UMTS) reflect the rapid growth in demand for portable and wireless services. Such systems demand higher functionality with extremely low levels of energy consumption. These conflicting requirements have been met to a great extent via scaling of feature sizes in modern CMOS technologies. However, technology scaling has also resulted in the emergence of deep submicron (DSM) noise [13] raising concerns regarding the ability of the semiconductor industry to extend Moore's law well into the deep submicron regime.

At the heart of the issue mentioned above is the problem of *achieving energy-efficiency in the presence of noise*, referred to as the *reliable low-power design* [12] problem. This

problem has two components: 1.) determining bounds on energy-efficiency in the presence of noise, and 2.) developing design techniques for approaching these bounds. Since 1995, our research has directly addressed these two subproblems.

First, we developed an information-theoretic paradigm for DSM VLSI systems [11, 6] and employed it to determine bounds on energy-efficiency of these systems in the presence of noise. In [11], we proposed the idea of viewing DSM VLSI systems as communication networks. Recently, other researchers have employed information-theoretic considerations to determine bounds on the switching energy of binary transition [9], bounds on energy for deep submicron busses [15], and on biological systems [1].

Second, we developed key elements of a design philosophy based on noise-tolerance, including noise-tolerant circuit design [2], and algorithmic noise-tolerance (ANT) [7] for broadband communication systems. Indeed, the recent 2001 International Technology Roadmap for Semiconductors [8] (ITRS2001) has described *error-tolerance* as a cost-effective approach for achieving reliability. The ITRS2001 further asserts the need to view systems-on-a-chip as a communication network.

In this paper, we provide an overview of ANT as a effective technique that provides significant gains in energy-efficiency in the presence of noise over conventional systems. The key idea behind ANT is to permit errors to occur in a signal processing block and then correct it via a separate error control block. This approach of error/noise-tolerance is fundamentally superior to mitigating noise and achieves energy-efficiencies beyond what is achievable via present day techniques. Note, ANT techniques are somewhat orthogonal to known low-power techniques [3, 10] in that ANT can be applied after/along with known power reduction techniques. A key difference between ANT and related work in the information-theoretic and coding communities [4] is that the latter do not account for energy-efficiency.

In this paper, we focus on ANT techniques as an attractive approach for jointly addressing reliability and energy-efficiency issues for broadband communications and digital signal processing systems.

2. ALGORITHMIC NOISE-TOLERANCE

In this section, we first define an ANT-based system in section 2.1 followed by a motivational example in section 2.2. In section 2.3, we describe the key elements of an ANT-based system and in section 2.4 we provide a rationale for the effectiveness of ANT in enhancing energy-efficiency.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

DAC 2002, June 10-14, 2002, New Orleans, Louisiana, USA.
Copyright 2002 ACM 1-58113-461-4/02/0006 ...\$5.00.

Table 1: Motivational Example

i	1	2	3	4	5	6	7	8
x	0.875	-0.875	0.875	-0.875	0.875	0.875	-0.875	0.875
η	0.000	0.000	0.875	0.000	0.000	0.875	0.000	0.000
y	0.547	0.219	0.328	-0.328	0.328	0.766	0.766	0.328
y'	0.547	0.219	1.203	-0.328	0.328	1.641	0.766	0.328
y_r	0.438	0.000	0.438	-0.438	0.438	0.438	0.438	0.438
\hat{y}	0.547	0.219	0.438	-0.328	0.328	0.438	0.766	0.328




Figure 1: ANT-based DSP System.

2.1 Formal Definition of ANT

An ANT-based DSP system (see Fig. 1) has a main DSP (**MDSP**) block that computes in an energy-efficient manner but makes intermittent errors. It accepts as its input the signal $x[n] + sn[n]$, where $x[n]$ is the input signal and $sn[n]$ is the input signal noise. The noisy output of the **MDSP** block is denoted by $y'[n]$. In essence, the **MDSP** block sacrifices noise-immunity for energy-efficiency. The **EC** block observes the noisy output $y'[n]$, and the input $x[n] + sn[n]$ and perhaps certain internal **MDSP** signals to detect and correct errors. The final corrected output of the ANT-based system is denoted as $\hat{y}[n]$. The error-control (**EC**) block operates in an error-free manner but consumes significantly more energy per operation than the **MDSP** block.

The assumption of an error-free **EC** block can be easily justified as follows. Due to energy-efficiency considerations, the **EC** block will be designed to be much simpler/smaller than the **MDSP** block. Thus, the constraints on delay and power will be significantly relaxed for the **EC** block and therefore can be easily met with a much higher noise-immunity. For example, one can use a highly noise-tolerant circuit technique [2] to design the **EC** block and pay a price in terms of extra power. As the **EC** block is designed to be much smaller than the **MDSP** block, this extra power will have minimal impact on the overall power consumption.

If the **MDSP** block is an FIR filter, then its noisy output $y'(n)$ can be written as,

$$y'[n] = \sum_{k=0}^{N-1} h_k x[n-k] + \eta[n] \\ = y[n] + \eta[n] \quad (1)$$

where $y[n]$ is the error-free output and $\eta[n]$ represents the manifestation of DSM noise at the algorithmic level on a per (output) sample basis. We now describe an example to illustrate the ANT concept.

2.2 Motivational Example

Consider the error-free output $y[n]$ of a finite-impulse re-

sponse (FIR) filter (main filter) given by

$$y[n] = 0.625x[n] + 0.875x[n-1] + 0.625x[n-2], \quad (2)$$

where $x[n]$ is a 4-bit two's complement uncorrelated input that is equally likely to take on a value of either 0.875 or -0.875. Note that the coefficients also have a 4-bit two's complement representation.

We model the impact of a DSM noise source with the signal $\eta[n]$ that takes on a value of ± 0.875 with probability 0.1 and 0 with a probability of 0.9. We denote the noisy output of the main filter as $y'[n]$, where

$$y'[n] = y[n] + \eta[n] \quad (3)$$

Note that $\eta[n]$ is non-zero quite infrequently but that the non-zero value (0.875) is quite large, i.e., comparable to the maximum value of the output $y[n]$ (1.86). Table 1 shows the values of $x[n]$, $y[n]$, $\eta[n]$ and $y'[n]$ for 8 sample periods.

In order to detect and correct these errors, we implement a half-precision filter (i.e., a filter whose coefficient precision is half of that of the main filter in (2)) whose input is $x[n]$ (i.e., same as that of the main filter), and whose output $y_r[n]$ is given by

$$y_r[n] = 0.5x[n] + 0.5x[n-1] + 0.5x[n-2] \quad (4)$$

Note that the coefficients of the half-precision filter can be obtained by truncating the 4-bit two's complement representation of the main filter by two LSBs.

We flag an error when the absolute value of the difference between the two filter outputs $|y'[n] - y_r[n]|$ is greater than a pre-specified threshold E_{th} . In this example, we set $E_{th} = \max(|y[n] - y_r[n]|) = 0.547$, where the maximization is over all input combinations. Let $\hat{y}[n]$ be the final (corrected) output of this system. If an error is flagged then we correct it via the assignment: $\hat{y}[n] = y_r[n]$. Otherwise, we set $\hat{y}[n] = y'[n]$.

Table 1 shows the values of $y_r[n]$ and $\hat{y}[n]$. A run of approximately 1000 samples shows that the output signal-to-noise ratios (SNR_s) as

$$SNR_n = 10 \log_{10} \left(\frac{\sigma_y^2}{\sigma_n^2} \right) = 8.0474 \text{ dB} \quad (5)$$

$$SNR_{ANT} = 10 \log_{10} \left(\frac{\sigma_y^2}{\sigma_{\hat{y}-y}^2} \right) = 13.2667 \text{ dB} \quad (6)$$

where σ_y^2 , σ_n^2 , and $\sigma_{\hat{y}-y}^2$ are the variances of the error-free output $y[n]$, the noise signal $\eta[n]$ and the residual noise signal $\hat{y}[n] - y[n]$, respectively. Note that SNR_n is the output SNR of the noisy main filter *before* error control and SNR_{ANT} is the output SNR of the ANT-based system, i.e., after error correction. It can be seen that a very simple error correction scheme can improve the output SNR by more than 5dB.

From this example, we conclude that the key elements of an ANT-based system are: 1.) the DSM noise model $\eta[n]$, 2.) the statistical nature of the metrics (e.g. SNR) employed to evaluate the algorithmic performance, and 3.) statistical properties of the signals being processed, and the algorithm itself. These are described next.

2.3 Elements of ANT

2.3.0.1 DSM Noise models and Architecture.

The signal $\eta[n]$ in (1) models the impact of DSM noise on an output sample. In general, noise may be random (e.g., alpha particle hits on a dynamic node) or may be systematic (e.g., critical path violations or delay faults due to numerous causes including variations in process, temperature and supply voltage). For practical reasons, in certain cases we may deliberately take $\eta[n]$ to be random even if the corresponding physical noise source is systematic. For example, supply bounce is a systematic source of noise because of its dependence on the input. This dependence is however extremely complex for large designs. Therefore, we may choose an $\eta[n]$ to be random when modelling the impact of supply bounce on the output of a DSP block.

DSM noise in general is hard to model especially at the algorithmic level as needed by ANT based systems. For this reason, in the past, we have proposed the concept of *voltage overscaling* (VOS) which extends the idea of voltage scaling and provides a 'well understood' source of noise. Assume that $V_{dd-crit}$ is the supply voltage at which the critical path delay just meets the delay constraints imposed by the application and architecture. VOS implies scaling of the supply voltage to a value below $V_{dd-crit}$, i.e., $V_{dd} = V_{dd-crit}/k_{vos}$, where $k_{vos} \geq 1$ is referred to as the voltage overscaling factor (VOSF). VOS results in one or more critical paths violating their delay constraints. This in turn will result in the output $y'(n)$ being in error when ever an appropriate input sequence is applied. We refer to such errors as soft errors because these disappear when the input sequence no longer excites the critical path.

Note that the statistics of $\eta[n]$ in a VOS scenario is a strong function not only of the input statistics but also the architecture. This is because the distribution of critical paths a function of the delay is dependent on the architecture. For example, an architecture utilizing ripple-carry structures and array multipliers will have a skewed distribution, i.e., a few paths will determine the overall delay of the system. This *delay imbalance* favors ANT because a skewed distribution implies a low frequency of errors when the architecture is voltage overscaled.

Therefore, in a VOS scenario and perhaps others, the noise model $\eta[n]$ has a dependence on the input signal $x[n]$. Later in the paper, we will present ANT techniques that exploit this dependence and those that do not.

2.3.0.2 Algorithmic Performance Metrics.

ANT techniques exploit the fact that the functional performance metrics for DSP and communication systems are statistical in nature. In reference to Fig. 1, the output SNR of a critically scaled (i.e., $V_{dd} = V_{dd-crit}$) (or noiseless) filter is given as

$$SNR_o = 10 \log_{10} \left(\frac{\sigma_y^2}{\sigma_{sno}^2} \right) \quad (7)$$

where σ_y^2 and σ_{sno}^2 are the powers of the output signal $y[n]$ and the output *signal noise* $sno[n]$ (not shown in Fig. 1), respectively. The term 'signal noise' refers to the noise that is present in the input signal $x[n]$ and which is denoted in Fig. 1 as $sn[n]$. Therefore, $sno[n]$ is a filtered version of $sn[n]$. Note, a signal noise source will exist even though the critically scaled filter is error-free thereby providing a finite value for SNR_o .

Now, the output SNR for a voltage overscaled or noisy filter is given by

$$SNR_n = 10 \log_{10} \left(\frac{\sigma_y^2}{\sigma_{sno}^2 + \sigma_n^2} \right) \quad (8)$$

where σ_y^2 , σ_{sno}^2 and σ_n^2 are the powers of the output signal $y[n]$, output signal noise $sno[n]$, and soft errors due to DSM noise $\eta[n]$, respectively.

Similarly, the output SNR of an ANT-based filter is given by

$$SNR_{ANT} = 10 \log_{10} \left(\frac{\sigma_y^2}{\sigma_{sno}^2 + \sigma_r^2} \right) \quad (9)$$

where σ_r^2 is the power in the residual soft error $r[n] = \hat{y}[n] - y[n]$.

Note that the presence of σ_{sno}^2 in the denominator in (9) provides room for imperfect error control and hence greater energy-efficiencies. This implies that in an ANT-based system, it is sufficient to reduce the DSM noise source $\eta[n]$ below the output signal noise floor represented by σ_{sno}^2 in order to minimize the impact on the output SNR .

2.3.0.3 Signal and Algorithmic Properties.

Signals being processed in DSP and communication systems can be modelled as random processes. The power spectrum of these signals are modulated by transmit filters, the physical channel and the receive filters (equalizers). Thus, one has a good idea of what the statistical properties of signals should be in a well-behaved or error-free system. DSM noise sources tend to disturb the statistical structure of signals. ANT techniques are extremely effective when the disturbance is large but correctable. Large disturbances make it easy to detect errors and, if correctable, the overall impact on SNR can be made minimal, i.e., $SNR_{ANT} \approx SNR_o$.

2.4 Energy Savings due to ANT

We now derive the conditions under which an ANT-based system leads to energy savings in a VOS scenario over a critically voltage-scaled system (defined as a system operating at $V_{dd-crit}$). The dynamic energy dissipation per clock cycle \mathcal{E}_{orig} of such a system is given by

$$\mathcal{E}_{orig} = C_{orig} V_{dd-crit}^2 \quad (10)$$

where C_{orig} is the average switching capacitance. Note that \mathcal{E}_{orig} is the minimum energy dissipation that conventional voltage scaling can achieve.

In comparison, the dynamic energy dissipation per clock cycle \mathcal{E}_{ANT} of the corresponding ANT-based system is given by

$$\mathcal{E}_{ANT} = C_{orig} \left(\frac{V_{dd-crit}}{k_{vos}} \right)^2 + C_{ANT} V_{dd-ant}^2 \quad (11)$$

where C_{ANT} represents the overhead complexity due to ANT, V_{dd-ant} is the critical supply voltage for EC block, and

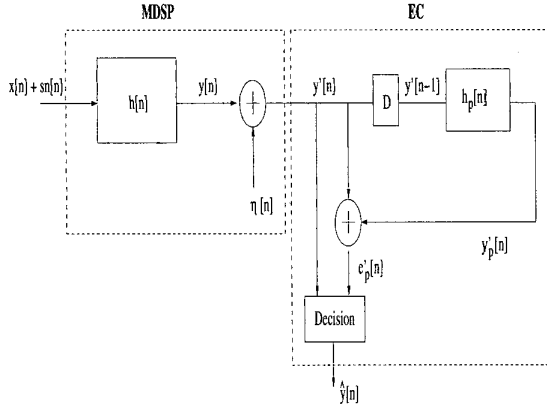


Figure 2: Prediction-based ANT technique.

$k_{vos} > 1$ is the VOS factor (VOSF). From (10)–(11), it can be easily shown that $\mathcal{E}_{ANT} < \mathcal{E}_{crit}$ provided that

$$C_{ANT} V_{dd-ant}^2 < C_{orig} V_{dd-crit}^2 \left(1 - \frac{1}{k_{vos}^2}\right) \quad (12)$$

In practice, the condition in (12) is easily satisfied by making C_{ANT} as small as possible and/or by making k_{vos} as large as possible. There is indeed an interesting trade-off between k_{vos} and C_{ANT} . When k_{vos} is increased, the performance degradation becomes larger as more critical paths and other longer paths start to fail. This requires increasingly sophisticated and hence complex ANT techniques which would increase C_{ANT} .

3. PREDICTION-BASED ANT

In this section, we describe a technique referred to as prediction-based ANT [7] (see Fig. 2) that employs the filter output $y'[n]$ to detect and correct errors.

Consider the output $y'[n]$ of a VOS filter that uses a least significant bit (LSB) first architecture. LSB first architectures are in fact the most commonly employed in practice. If the filter is sufficiently narrowband then the error-free output signal $y[n]$ is highly correlated, i.e., the MSBs do not change from one sample to the next. A linear predictor, which is another filter, therefore can be used to statistically predict the output $y[n]$. The output $y_p[n]$ of a linear predictor of $y[n]$ is given by

$$y_p[n] = \sum_{k=0}^{N_p-1} p_k y[n-k-1] \quad (13)$$

where N_p is the number of predictor taps and p_k are the predictor coefficients. The prediction error $e[n] = y[n] - y_p[n]$ can be minimized via a proper choice of the coefficients p_k .

Under VOS, we do not have access to $y[n]$ (the error-free output) but only to $y'[n]$ (the noisy output). We will assume however that the errors occur with a frequency which is sufficiently less than $1/(2N_p)$. In that case, the predictor output under VOS $y_p'[n]$ equals $y_p[n]$ (the error-free output).

Also, under VOS, the MSBs will fail first, i.e., $\eta[n]$ in (1) is large. The prediction error under VOS $e_p'[n]$ is therefore

given by

$$\begin{aligned} e_p'[n] &= y'[n] - y_p'[n] \\ &= y[n] - y_p[n] + \eta[n]. \end{aligned} \quad (14)$$

In this case, error detection is accomplished via a simple threshold on $e_p'[n]$. Error correction on the other hand is also achieved by assigning the final corrected output $\hat{y}[n]$ as

$$\hat{y}[n] = y_p[n]. \quad (15)$$

In addition, we also assign $y_p[n]$ as the next N_p corrected outputs. This is because once $y'[n]$ is in error (i.e., $y'[n] \neq y[n]$) then $y_p[n] \neq y_p'[n]$ for the next N_p samples while the incorrect output $y'[n]$ is being flushed out from the data buffer in the predictor (see (13)).

We summarize the prediction-based ANT scheme consists of the following steps:

- Error detection: if $|e_p'[n]| > E_{th}$ an error is declared.
- Error correction: If an error is declared then $\hat{y}[n] = y_p'[n]$ else $\hat{y}[n] = y'[n]$.

The assumptions made in deriving the prediction-based ANT technique are:

- magnitude of noise $\eta[n]$ is relatively large, i.e., comparable to the maximum value that $y[n]$ can achieve.
- the probability of $\eta[n] \neq 0$ is less than $1/(2N_p)$.

In addition, for obtaining large energy savings, we require that the bandwidth of the output $y[n]$ is sufficiently small so that the predictor complexity is much smaller than that of the MDSP block. We have shown through measurements conducted on an integrated circuit designed in $0.35\mu\text{m}$, 3.3V CMOS process [5] that the prediction-based ANT technique provides up to 67% energy savings over a critically scaled filter.

The prediction-based ANT technique is surprisingly effective even when the output bits are flipped randomly with a probability p_{err} . We have shown [7] that ANT technique reduces the drop in SNR over a noiseless system from 11dB to 2dB for a value of p_{err} as high as 10^{-3} . In other words, the ANT based system improves the SNR by 9dB when each filter output bit is being flipped at an average rate of once every 1000 samples independent of each other.

4. ERROR CANCELLATION-BASED ANT

The error cancellation-based ANT technique shown in Fig. 3 exploits any correlation that may exist between $\eta[n]$ and the input $x[n]$. Note that such correlation certainly exists under a VOS scenario.

The error cancellation-based ANT requires a separate filter called the error canceller $h_e[n]$ that generates an estimate of $\eta[n]$ (denoted as $\hat{\eta}[n]$) from $x[n]$. This error canceller needs to be trained first in order to learn the correlation structure between $\eta[n]$ and the input $x[n]$. Figure 3, shows how this is done. During the training phase, a known input sequence $x[n] + sn[n]$ is provided at the input to the MDSP block. At the same time, the multiplexer provides the corresponding precomputed error-free output $y[n]$. Therefore, during the training phase, the output $y''[n]$ is given by,

$$y''[n] = y'[n] - y[n] = y[n] + \eta[n] - y[n] = \eta[n] \quad (16)$$




Figure 3: Error cancellation-based ANT technique.

and therefore the final output $\hat{y}[n]$ is given by

$$\hat{y}[n] = y'[n] - \hat{\eta}[n] = e[n] \quad (17)$$

where the estimation error $e[n]$ is used to adapt the error canceller. Any of the well-known adaptive filtering algorithms such as the least-mean squared (LMS) algorithm can be employed here.

Normal operation commences once the error canceller has been trained, i.e., the variance of $e[n]$ has been minimized. In this phase, the input sequence is unknown and the multiplexer outputs a zero. In addition, the error canceller stops adapting. The output $y''[n]$ is therefore given by,

$$y''[n] = y'[n] - 0 = y[n] + \eta[n] \quad (18)$$

and hence the final output $\hat{y}[n]$ is given by

$$\begin{aligned} \hat{y}[n] &= y''[n] - \hat{\eta}[n] \\ &= y[n] + \eta[n] - \hat{\eta}[n] = y[n] + e[n] \end{aligned} \quad (19)$$

Equation (19) indicates that if the error canceller has been trained correctly then during normal mode of operation, the final output $y''[n]$ will be very close to the error free output $y[n]$.

When applied to digital FIR filters, we have shown that error-cancellation works best for broadband filters and energy savings of up to 71% can be achieved over critically scaled systems.

5. REDUCED PRECISION REDUNDANCY (RPR) BASED ANT

In the RPR technique [14], we exploit the fact that MSBs are the critical bits and hence need to be protected from noise.

The RPR technique shown in Fig. 4 has a low precision replica $h_r[n]$ of the filter $h[n]$ in the MDSP block. In this case, $h_r[n]$ takes the same input $x[n] + sn[n]$ as $h[n]$ but computes only the MSBs of the error-free output $y[n]$. Note that the output of the RPR filter $y_r[n]$ will not equal the error-free output $y[n]$ of the MDSP filter due to quantization noise. Quantization noise properties for DSP systems are well-understood. Thus, the error control scheme in case of RPR is as follows:

- compute the difference metric: $d(y', y_r) = |y'[n] - y_r[n]|$.




Figure 4: Reduced precision redundancy-based ANT technique.

- Error detection: if $d(y', y_r) > E_{th}$, where E_{th} is a predefined threshold value, then flag an error.
- Error correction: if an error has been flagged then $\hat{y}[n] = y_r[n]$ otherwise $\hat{y}[n] = y'[n]$.

Note that the threshold E_{th} will need to be set at a value above the quantization noise floor. An ideal value for E_{th} would be $\max(y'[n] - y_r[n])$, where the maximization is carried out over all noise scenarios. In case of VOS, this maximization needs to be done over all possible input combinations.

When applied to digital filtering [14], RPR provides up to 67% energy savings. We have also shown that RPR-based multipliers provide up to 44% energy savings in an FFT.

6. CAD ISSUES

The design of ANT-based DSP systems presents numerous opportunities for CAD researchers. This section describes some of the challenges.

6.1 Noise-Tolerant Design Methodology

New techniques for DSM noise analysis are required that lead to a better understanding of the energy penalty involved in controlling noise at the physical level. Also, algorithms and techniques need to be developed that propagate the impact of various noise sources to the algorithmic level, and development of algorithmic noise models. This will enable designers to evaluate the energy-efficiency benefits of ANT-based DSP systems under general noise scenarios and contribute greatly to the goals stated in ITRS2001[8].

Determining energy-optimal trade-offs between circuit level noise-tolerance and ANT is an open problem. The VOS scenario also provides a nice test case for exploring these trade-offs because the above mentioned noise modelling problem for VOS is much more tractable than for the general DSM noise case.

While any noise-tolerance based design paradigm would be effective in achieving energy-efficiency in the presence of noise, a question still remains as to what the bounds on energy-efficiency in the presence of noise are? The information-theoretic view proposed in [11] has been applied to simple gates, where ANT techniques do not apply. There is

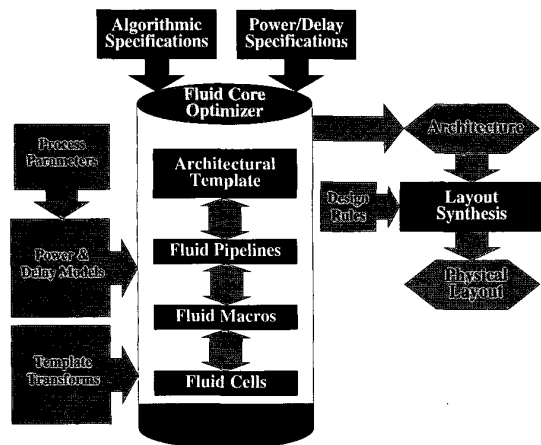


Figure 5: A fluid IP core generator.

currently a gap between our ability to compute achievable bounds on energy-efficiency and being able to compare those with the efficiencies achieved via techniques such as ANT, which by default are effective for complex systems. Computational techniques are required for determining the bounds on energy-efficiencies for complex VLSI systems.

6.2 Fluid IP Core Generators

In order to make the benefits of an ANT-based design methodology available to the designer, we propose the development of a fluid IP core generator as shown in Fig. 5.

A fluid IP core generator synthesizes a custom-quality layout of an algorithm-specific block (such as an equalizer) from high-level specifications without going through a synthesis and place & route step. The term 'fluid' refers to the fact that the core generator while optimizing the architecture can reach down into the circuit fabric and tune transistor sizes while generating layouts. Core generators are useful in broadband communication systems where energy and throughput efficiency limits need to be approached and where canonical blocks (such as equalizers, filters, FFTs, Reed-Solomon decoders etc.) are employed repeatedly though with varying parameters across multiple applications.

The key features/benefits of a fluid IP core generator are: 1.) process and device scalability, 2.) custom quality design within a synthesis quality design cycle, 3.) predictability of hard cores and flexibility of soft cores, and 4.) encapsulation of cross-domain optimization techniques such as ANT. An inherent drawback of such a core generator is that these are designed for algorithm-specific blocks.

An IP core generator has two major components: 1.) an architecture optimizer, and 2.) a layout synthesizer, and it accepts as inputs: 1.) power and delay models, 2.) library of template transforms, 3.) algorithmic specifications and 4.) power and delay specifications.

Design of a fluid IP core generator presents numerous problems. Development of process scalable power and delay models is one. Such development seems feasible given the fact that the functionality of algorithm-specific blocks do not change over process generations. Hence, models that require incremental changes from one process generation to the next would be useful.

7. ACKNOWLEDGMENTS

The author would like to acknowledge support from National Science Foundation grants CCR-000987, CCR 99-79381 and DARPA for the work reported in this paper.

8. REFERENCES

- [1] P. Abshire and A. Andreou. Capacity and energy cost of information in biological and silicon photoreceptors. *Proceedings of IEEE*, 89(7):1052–1064, July 2001.
- [2] G. Balamurugan and N. R. Shanbhag. The twin-transistor noise-tolerant dynamic circuit technique. *IEEE Journal of Solid-State Circuits*, 36(2):273–280, February 2001.
- [3] A. Chandrakasan and R. Brodersen. Minimizing power consumption in digital cmos circuits. *Proceedings of the IEEE*, 83(4), April 1995.
- [4] B. Hajek and T. Weller. On the maximum tolerable noise for reliable computation by formulas. *IEEE Transactions on Information Theory*, 37(2), March 1991.
- [5] R. Hegde and N. R. Shanbhag. A low-power digital filter ic via soft dsp. In *Proceedings of Custom Integrated Circuits Conference*, May 1991.
- [6] R. Hegde and N. R. Shanbhag. Towards achieving energy-efficiency in the presence of deep submicron noise. *IEEE Transactions on VLSI*, 8(4):379–391, April 2000.
- [7] R. Hegde and N. R. Shanbhag. Soft digital signal processing. *IEEE Transactions on VLSI*, 9(6):813–823, December 2001.
- [8] ITRS2001. International technology roadmap for semiconductors. <http://public.itrs.net/Files/2001ITRS/Home.htm>, 2001.
- [9] J. D. Meindl and J. A. Davis. The fundamental limit on binary switching energy for terascale integration (tsi). *IEEE Journal of Solid-State Circuits*, 36:1515–1516, October 2000.
- [10] J. Rabaey and M. Pedram. *Low Power Design Methodologies*. Kluwer Academic Publishers, 1996.
- [11] N. R. Shanbhag. A mathematical basis for power reduction in digital vlsi systems. *IEEE Transactions on Circuits and Systems: Part II*, 44(11):935–951, Nov. 1997.
- [12] N. R. Shanbhag and et. al. Reliable low-power design in the presence of deep submicron noise. In *Proceedings of the Intl. Symposium on Low-Power Design*, pages 295–302, July 2000.
- [13] K. L. Shepard. Conquering noise in deep submicron digital ics. *IEEE Design and Test of Computers*, pages 51–62, Jan.-Mar. 1998.
- [14] B. Shim and N. R. Shanbhag. Reduced precision redundancy for low-power digital filtering. In *Proceedings of Asilomar Conference*, 2001.
- [15] P. Sotiriadis, A. Chandrakasan, and V. Tarokh. Maximum achievable energy-reduction using coding with applications to deep sub-micron busses. In *Proc. of Intl. Symp. on Circuits and Systems*, 2002.