VLSI DESIGN '96

VLSI IN MOBILE COMMUNICATION

PROCEEDINGS OF
THE 9th INTERNATIONAL CONFERENCE ON VLSI DESIGN

Bangalore, India
January 3–6, 1996

Sponsored by

VLSI SOCIETY OF INDIA (VSI)

DOE, GOVERNMENT OF INDIA

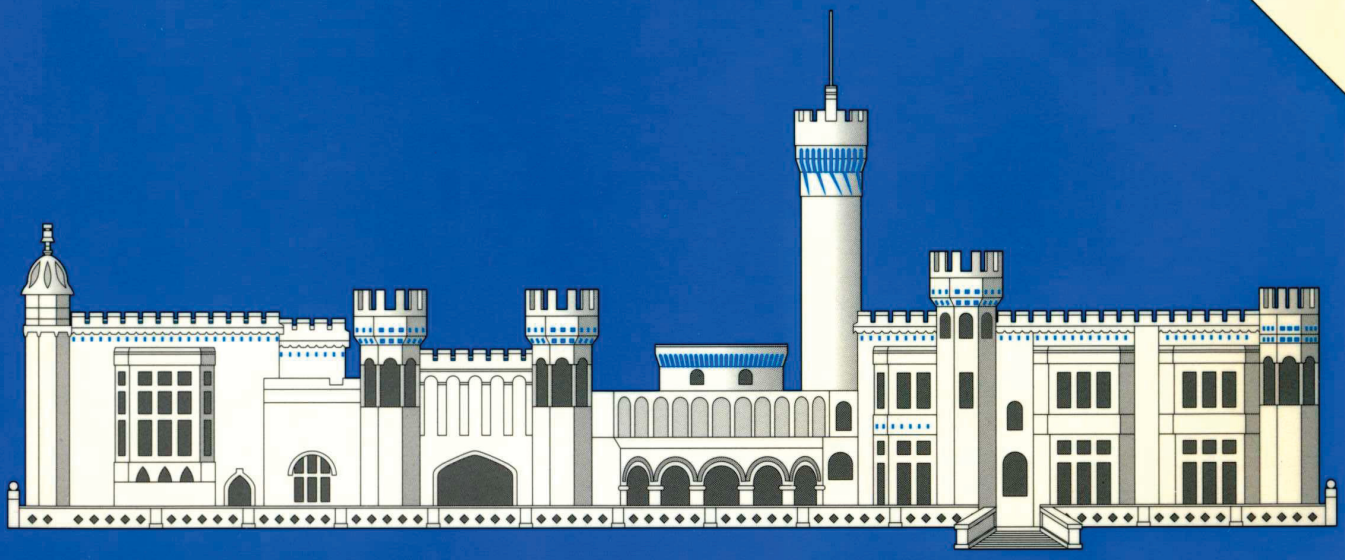
In cooperation with

ASSOCIATION FOR COMPUTING MACHINERY
SPECIAL INTEREST GROUP ON DESIGN AUTOMATION

IEEE CIRCUITS AND SYSTEMS SOCIETY

IEEE COMPUTER SOCIETY
DATC , VLSI-TC

IEEE



IEEE COMPUTER SOCIETY PRESS

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

SAGEMCOM-1053
IPR2026-00263

Comcast, Ex. 1017

Proceedings

Ninth International Conference on VLSI Design

January 3-6, 1996

Bangalore, India

Sponsored by

VLSI Society of India
Department of Electronics, Government of India

In cooperation with

IEEE Computer Society
IEEE
ACM SIGDA
IEEE Circuits & Systems Society
IEEE Computer Society Technical Committee on Design Automation
IEEE Computer Society Technical Committee on VLSI



IEEE Computer Society Press
Los Alamitos, California

Washington • Brussels • Tokyo

TK7874
.C17
1996

IEEE Computer Society Press
10662 Los Vaqueros Circle
P.O. Box 3014
Los Alamitos, CA 90720-1264

Copyright © 1995 by The Institute of Electrical and Electronics Engineers, Inc.
All rights reserved.

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries may photocopy beyond the limits of US copyright law, for private use of patrons, those articles in this volume that carry a code at the bottom of the first page, provided that the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

Other copying, reprint, or republication requests should be addressed to: IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331.

The papers in this book comprise the proceedings of the meeting mentioned on the cover and title page. They reflect the authors' opinions and, in the interests of timely dissemination, are published as presented and without change. Their inclusion in this publication does not necessarily constitute endorsement by the editors, the IEEE Computer Society Press, or the Institute of Electrical and Electronics Engineers, Inc.

IEEE Catalog Number 96TB100010
ISBN 0-8186-7228-5 (softbound)
ISBN 0-8186-7230-7 (microfiche)
ISSN: 1063-9667

Additional copies may be ordered from:

IEEE Computer Society Press
Customer Service Center
10662 Los Vaqueros Circle
P.O. Box 3014
Los Alamitos, CA 90720-1264
Tel: +1-714-821-8380
Fax: +1-714-821-4641
Email: cs.books@computer.org

IEEE Service Center
445 Hoes Lane
P.O. Box 1331
Piscataway, NJ 08855-1331
Tel: +1-908-981-1393
Fax: +1-908-981-9667

IEEE Computer Society
13, Avenue de l'Aquilon
B-1200 Brussels
BELGIUM
Tel: +32-2-770-2198
Fax: +32-2-770-8505

IEEE Computer Society
Ooshima Building
2-19-1 Minami-Aoyama
Minato-ku, Tokyo 107
Japan
Tel: +81-3-3408-3118
Fax: +81-3-3408-3553

Editorial production by Regina Spencer Sipple
Cover design by Alexander Torres
Printed in the United States of America by KNI, Inc.

The Institute of Electrical and Electronics Engineers, Inc.

7-23-96 JA

Table of Contents

Ninth International Conference on VLSI Design: VLSI in Mobile Communication

General Chair's Message xiii
Program Co-Chairs' Message xv
1995 IEEE Fellowship Awards xviii
1995 Awards xix
Conference Committee xx
Program Committee xxii
Steering Committee xxiv
1995 Fellowship Recipients xxiv
List of Reviewers xxviii
Conference History xxxiv

Tutorial Pages

T1: Low Power Design 2
K. Roy and R.K. Roy
T2: Register Transfer Level Synthesis: From Theory to Practice 2
K. Keutzer and S. Malik
T3: Mixed-Signal Design for Test 2
B. Vinnakota and R. Harjani
T4: VLSI Implementation of DSP Architectures 3
A. Ratan Gupta and V. Visvanathan
T5: Hardware Software Co-Design of Embedded Systems 3
R. Gupta
T6: Practical Test and DFT for Next Generation VLSI 3
J. Abraham and G. Ganapathy

1995 Keynote Address

Science, Technology, and the Indian Society 6
Vishwani D. Agrawal, AT&T Bell Laboratories, New Jersey, USA

1996 Keynote Address

The New Electronics Industry 10
Thomas J. Engibous, Texas Instruments, Inc., Dallas, Texas, USA

Session 1: Plenary Session: Invited Address

VLSI in Mobile Communication 11
Rajeev Jain, University of California, Los Angeles, USA
Robert W. Brodersen, University of California, Berkeley, USA

Session 2: VLSI for Mobile Communication

Chair: Pradeep Shah, Texas Instruments, Inc., Dallas, Texas, USA
2.1 Medium Access Control and Air-Interface Subsystem for an Indoor Wireless ATM Network 14
M.B. Srivastava

LINDA HALL LIBRARY
KANSAS CITY, MISSOURI

2.2 Behavioral Modeling of an ATM Switch Using SpecCharts 19
A. Sriram and F.J. Kurdahi

2.3 A Novel Allocation Strategy for Control and Memory Intensive Telecommunication Circuits 23
B. Svantesson, A. Hemani, P. Ellervee, A. Postula, J. Öberg, A. Jantsch, and H. Tenhunen

2.4 Design Tradeoffs in High-Speed Multipliers and FIR Filters 29
C. Nagendra, R.M. Owens, and M.J. Irwin

Session 3: Placement and Routing

Chair: Jun-Dong Cho, Sung Kyun Kwan University, Korea

3.1 A Simple Yet Effective Genetic Approach for the Orientation Assignment on Cell-Based Layout 33
J.-T. Yan

3.2 Parallel Simulated Annealing Strategies for VLSI Cell Placement 37
J.A. Chandy and P. Banerjee

★ **3.3** Channel Routing in Manhattan-Diagonal Model 43
S. Das and B.B. Bhattacharya

3.4 Routing Using Implicit Connection Graphs 49
S.Q. Zheng, J.S. Lim, and S.S. Iyengar

Session 4: Built-In Self-Test and Diagnosis

Chair: Asoke K. Laha, Cadence Design Systems, Noida, India

4.1 Design for High-Speed Testability of Stuck-At Faults 53
T.J. Chakraborty and V.D. Agrawal

4.2 A Novel BIST Architecture with Built-In Self Check 57
M.F. Abdulla, C.P. Ravikumar, and A. Kumar

4.3 Programmable Cellular Automata-Based Testbed for Fault Diagnosis in VLSI Circuits 61
S. Nandi, S. Chattopadhyay, and P.P. Chaudhuri

4.4 Hierarchical Probablistic Diagnosis of MCMs on Large-Area Substrates 65
K. Sasidhar and A. Chatterjee

Session 5: Hardware/Software Co-Design

Chair: Anupam Basu, IIT, Kharagpur, India

5.1 A Novel VLSI Concurrent Dual Multiplier - Dual Adder Architecture for Image and Video Coding Applications 69
D.V. Poornaiah and P.V.A. Mohan

5.2 Instruction-Set Matching and GA-Based Selection for Embedded-Processor Code Generation 73
J. Shu, T.C. Wilson, and D.K. Banerji

5.3 Instruction-Set Modeling for ASIP Code Generation 77
R. Leupers and P. Marwedel

5.4 CoDe-X: A Novel Two-Level Hardware/Software Co-Design Framework 81
R.W. Hartenstein, J. Becker, R. Kress, and H. Reinig

Session 6: Analog Circuits

Chair: A.N. Chandorkar, IIT, Bombay, India

6.1 A Micropower Analog Hearing Aid on Low-Voltage CMOS Digital Process 85
A.B. Bhattacharyya, R.S. Rana, S.K. Guha, R. Bahl, S. Anand, M.J. Zarabi, P.A. Govindacharyulu, V. Gupta, V. Mohan, J. Roy, and A. Atri

6.2 A 20MHz CMOS Variable Gain Amplifier 90
C. Srinivasan and K.R. Rao

6.3 Low-Power, Low-Voltage BiCMOS Comparators for Approximately 200MHz, 8bit Operation 94
A. Boni and C. Morandi

6.4 A Very High Gain Bandwidth Product Fully Differential Amplifier 99
J. Weiss, B. Majoux, and G. Bouvier

Session 7: Automatic Test Pattern Generation

Chair: James Jacob, IISc, Bangalore, India

7.1 Fast Algorithms for Computing I_{DDQ} Tests for Combinational Circuits 103
P. Thadikaran and S. Chakravarty

7.2 On More Efficient Combinational ATPG Using Functional Learning 107
R. Mukherjee, J. Jain, M. Fujita, J.A. Abraham, and D.S. Fussell

✶ **7.3** Sequential Circuits with Combinational Test Generation Complexity 111
A. Balakrishnan and S.T. Chakradhar

7.4 Genetic Algorithms for Scan Path Design 118
C.P. Ravikumar and R. Rajarajan

Session 8: High-Level Synthesis I

Chair: Shashi Kumar, IIT, New Delhi, India

8.1 Allocation and Binding in Data Path Synthesis Using a Genetic Algorithm Approach 122
C.A. Mandal, P.P. Chakrabarti, and S. Ghose

8.2 A Hierarchical Register Optimization Algorithm for Behavioral Synthesis 126
S. Katkoori, J. Roy, and R. Vemuri

8.3 A Rule-Based Approach for Improving Allocation of Filter Structures in HLS 133
J. Öberg, J. Isoaho, P. Ellervee, A. Jantsch, and A. Hemani

✶ **8.4** Representation and Synthesis of Interface of a Circuit for its Reuse 140
S. Sarkar, A. Basu, and A.K. Majumdar

Session 9: High-Performance Circuits

Chair: Rajeev Jain, University of California, Los Angeles, USA

9.1 Clock-Skew Constrained Cell Placement 146
V. Natesan and D. Bhatia

9.2 On Moment-Based Metrics for Optimal Termination of Transmission Line Interconnects 150
R. Gupta, B. Krauter, and L.T. Pileggi

★ **9.3** Combined Effect of Grain Boundary Depletion and PolySi/Oxide Interface Depletion on Drain Characteristics of a p-MOSFET 156
P.R. Suresh, P. Venugopal, S.T. Selvam, and S. Potla

9.4 An Analytical Delay Model Based on Boolean Process 162
Y. Min, Z. Zhao, and Z. Li

Session 10: Sequential Automatic Test Pattern Generation

Chair: M.M. Hasan, IIT, Kanpur, India

10.1 Methods for Dynamic Test Vector Compaction in Sequential Test Generation 166
T.J. Lambert and K.K. Saluja

10.2 Dynamic Test Sequence Compaction for Sequential Circuits 170
A. Raghunathan and S.T. Chakradhar

10.3 An Efficient Test Generation Technique for Sequential Circuits with Repetitive Sub-Circuits 174
D.R. Chakrabarti and A. Jain

★ **10.4** Synchronous Test Generation Model for Asynchronous Circuits 178
S. Banerjee, S.T. Chakradhar, and R.K. Roy

Session 12: High-Level Synthesis II

Chair: Bhargab Bhattacharya, Indian Statistical Institute, Calcutta, India

★ **12.1** Behavioral Synthesis of Complex Parallel Controllers 186
K. Bilinski, E.L. Dagless, and J.M. Saul

12.2 A Graph-Based Approach to the Synthesis of Multi-Chip Module Architectures 192
R.V. Cherabuddi, J. Chen, and M.A. Bayoumi

12.3 Synchronous Controller Models for Synthesis from Communicating VHDL Processes 198
N. Narasimhan, J. Roy, and R. Vemuri

12.4 Synthesis of Testable Pipelined Datapaths Using Genetic Search 205
C.P. Ravikumar and V. Saxena

Session 13: Field-Programmable Gate Arrays

Chair: Mahesh Mehendale, Texas Instruments, Bangalore, India

13.1 Bipartitioning for Hybrid FPGA-Software Simulation 211
A. Singla and T.M. Conte

13.2 LUT-Based FPGA Technology Mapping Using Permissible Functions 215
T. Suyama, H. Sawada, and A. Nagoya

13.3 VaWiRAM: A Variable Width Random Access Memory Module 219
L.K. John

13.4 Node-Covering Based Defect and Fault Tolerance Methods for Increased Yield in FPGAs 225
F. Hancheck and S. Dutt

Session 14: Mixed-Signal Design and Test

Chair: Radhakrishna Rao, IIT, Madras, India

14.1 Low-Cost DC Built-In Self-Test of Linear Analog Circuits Using Checksums 230
A. Chatterjee, B. Kim, and N. Nagi

14.2 Design of High-Performance Two-Stage CMOS Cascode Op-Amps with Stable Biasing 234
P. Mandal and V. Visvanathan

14.3 Testing Analog Circuits by AC Power Supply Voltage 238
A.K.B. A'ain, A.H. Bratt, and A.P. Dorey

✶ **14.4** Test Generation for Mixed-Signal Devices Using Signal Flow Graphs 242
R. Ramadoss and M.L. Bushnell

Session 15: Logic Design and Synthesis

Chair: Sujit Dey, NEC, Princeton, New Jersey, USA

15.1 A Study of Composition Schemes for Mixed Apply/Compose Based Construction of ROBDDs 249
A. Narayan, S.P. Khatri, J. Jain, M. Fujita, R.K. Brayton, and A. Sangiovanni-Vincentelli

15.2 On Finding Functionally Identical and Functionally Opposite Lines in Combinational Logic Circuits 254
I. Pomeranz and S.M. Reddy

15.3 Does Retiming Affect Redundancy in Sequential Circuits? 260
D.K. Das and B.B. Bhattacharya

15.4 Cubical CAMP for Minimization of Boolean Functions 264
N.N. Biswas, C. Srikanth, and J. Jacob

Session 16: Architecture

Chair: K. Neelakantan, Advanced Num. Res. & Anal. Group, Hyderabad, India

16.1 A VLSI Architecture for Cellular Automata-Based Parallel Data Compression 270
S. Bhattacharjee, J. Bhattacharya, U. Raghavendra, D. Saha, and P.P. Chaudhuri

16.2 VLSI/WSI Designs for Folded Cube-Connected Cycles Architectures 276
M.P. Sebastian, P.S.N. Rao, and L. Jenkins

16.3 A Tree Matching Chip 280
V. Krishna, A. Ejnioui, and N. Ranganathan

16.4 A Systolic Architecture for LMS Adaptive Filtering with Minimal Adaptation Delay 286
S. Ramanathan and V. Visvanathan

Session 17: Logic and Fault Simulation

Chair: Debashis Bhattacharya, Texas Instruments, Dallas, Texas, USA

17.1 Statistical Path Delay Fault Coverage Estimation for Synchronous Sequential Circuits 290
L. Pappu, M.L. Bushnell, V.D. Agrawal, and M.K. Srinivas

17.2 Retiming with Logic Duplication Transformation: Theory and an Application to Partial Scan 296
A. Balakrishnan and S.T. Chakradhar

17.3 Is Compiled Simulation Really Faster than Interpreted Simulation? 303
P.M. Maurer

17.4 Distributed Logic Simulation: Time-First Evaluation vs. Event-Driven Algorithms 307
S. Sundaram and L.M. Patnaik

Session 18: VLSI in Communications and Applications

Chair: P.V. Ananda Mohan, Indian Telephone Ind., Bangalore, India

18.1 Designing Systems on Silicon: A Digital Spread Spectrum Pager 311
S. Samel, B. Gyselinckx, I. Bolsens, and H. De Man

18.2 A High-Speed 32-bit Parallel Correlator for Spread Spectrum Communication 313
S. Kulkarni, P. Mazumder, and G.I. Haddad

18.3 Architecture of a VLSI Chip for Modeling Amino Acid Sequence in Proteins 316
S. Mitra, S. Das, P.P. Chaudhuri, and S. Nandi

18.4 VLSI Implementation of Artificial Neural Network-Based Digital Multiplier and Adder 318
R. Ranade, S. Bhandari, and A.N. Chandorkar

18.5 Cellular Automata-Based Architecture of a Database Query Processor 320
S. Chattopadhyay, S. Mitra, and P.P. Chaudhuri

18.6 A Multiplier Generator for Xilinx FPGA's 322
J.P. Singh, A. Kumar, and S. Kumar

Session 19: Low-Power and Analog Design

Chair: Anantha Chandrakasan, MIT, Cambridge, Massachusetts, USA

19.1 Estimation of Power from Module-Level Netlists 324
C.P. Ravikumar, M.R. Prasad, and L.S. Hora

19.2 Instruction-Level Power Analysis and Optimization of Software 326
V. Tiwari, S. Malik, A. Wolfe, and M.T.-C. Lee

19.3 Challenges in Low-Power Microprocessor Design 329
S. Rajgopal

19.4 An Enhanced Macromodel for a CMOS Operational Amplifier for HDL Implementation 331
S. Aggarwal

19.5 KANSYS: A CAD Tool for Analog Circuit Synthesis 333
S.K. Gupta and M.M. Hasan

Session 20: Test and Logic Synthesis

Chair: Robert Grafton, National Science Foundation, Arlington, Virginia, USA

20.1 Dual Rail Static CMOS Architecture for Wave Pipelining 335
G.E. Fernández and R. Sridhar

20.2 Syndrome Signature in Output Compaction for VLSI BIST 337
S.R. Das, N. Goel, W.B. Jone, and A.R. Nayak

20.3 A Multilevel Factorization Technique for Pass Transistor Logic 339
A. Jaekel, G.A. Jullien, and S. Bandyopadhyay

20.4 Characteristic Polynomial Method for Verification and Test of Combinational Circuits 341
V.D. Agrawal and D. Lee

20.5 SUBGEN: A Genetic Approach for Subcircuit Extraction 343
N. Vijaykrishnan and N. Ranganathan

Session 21: Low-Power Design

Chair: N.S. Murty, Semiconductor Complex Ltd., Bangalore, India

✶ **21.1** Wiresizing with Buffer Placement and Sizing for Power-Delay Tradeoffs 346
J.C. Shah and S.S. Sapatnekar

21.2 Ultra Low-Power Digital Signal Processing 352
A.P. Chandrakasan

21.3 A Low-Power Video Encoder with Power, Memory and Bandwidth Scalability 358
N. Chaddha and M. Vishwanath

✶ **21.4** Maximum-Power Estimation for CMOS Circuits Using Deterministic and Statistic Approaches 364
C.-Y. Wang and K. Roy

✶ **21.5** Low-Power Realization of FIR Filters Using Multirate Architectures 370
M. Mehendale, S.D. Sherlekar, and G. Venkatesh

Session 22: Asynchronous Circuits, Retiming, and Parititioning

Chair: S.P. Uttam, Dept. of Electronics, Government of India, New Delhi, India

22.1 Self-Timed Division and Square Root Extraction 376
A. Guyot, M. Renaudin, B. El Hassan, and V. Levering

22.2 Elimination of Dynamic Hazards from Signal Transition Graphs 382
R. Nagalla and G. Hellestrand

22.3 Automatic Synthesis of Speed-Independent Circuits from Signal Transition Graph Specifications 389
S.-B. Park and T. Nanya

22.4 Multi-Way Partitioning of VLSI Circuits 393
P. Agrawal, B. Narendran, and N. Shivakumar

22.5 Geometric Bipartitioning Problem and its Applications to VLSI 400
P.S. Dasgupta, S.C. Nandy, A.K. Sen, and B.B. Bhattacharya

Session 23: Delay Testing

Chair: C. Shekhar, Central Electronic Eng. Research Institute, Pilani, India

23.1 Identifying Redundant Path Delay Faults in Sequential Circuits 406
R. Tekumalla and P.R. Menon

23.2 Diagnosis of Parametric Path Delay Faults 412
M. Sivaraman and A.J. Strojwas

23.3 On Test Coverage of Path-Delay Faults 418
A.K. Majhi, J. Jacob, L.M. Patnaik, and V.D. Agrawal

23.4 Improving Accuracy in Path Delay Fault Coverage Estimation 422
K. Heragu, J.H. Patel, and V.D. Agrawal

✶ **23.5** Parallel Concurrent Path-Delay Fault Simulation Using Single-Input Change Patterns 426
M.A. Gharaybeh, M.L. Bushnell, and V.D. Agrawal

Panel Discussion:

Mobile Communications: Demands on VLSI Technology, Design and CAD 432
Moderator:Srinivasan Rajam, Texas Instruments (India) Ltd., Bangalore, India
Panelists: A. Chandrakasan, MIT, Cambridge, Massachusetts, USA
K. Keutzer, Synopsys, Inc., Mountain View, California, USA
A. Khandekar, Motorola India Electronics Ltd., Bangalore, India
S.L. Maskara, Indian Institute of Technology, Kharagpur, India
B.D. Pradhan, Center for the Development of Telematics, New Delhi, India
M.B. Srivastava, AT&T Bell Labs, Murray Hill, New Jersey, USA

1997 Call for Papers 434
1997 Fellowship Form 435
Author Index 437

✶ Best Paper Award Candidate Based on Rankings
Best Student Paper Award Winner

This material may be protected by Copyright law (Title 17 U.S. Code)

Ultra Low Power Digital Signal Processing

Anantha P. Chandrakasan
Massachusetts Institute of Technology, Cambridge, MA

ABSTRACT

The explosive growth of portable wireless devices has elevated power consumption to be one of the most critical design parameters. This paper presents several techniques to implement DSP functions with the lowest possible power consumption. Since power is consumed only when capacitance is being switched, power can be reduced by minimizing this capacitance through operation reduction, choice of data representation, exploitation of spatial and temporal signal correlations, re-synchronization to minimize glitching, and logic, circuit and physical design. Aggressive voltage scaling to 1V and below through circuit and architecture optimization is the key to low-power design. A systems approach to low-power design can result in orders of magnitude power reduction.

1. Introduction

The weight and size of batteries, which are major factors in portable systems, are primarily determined by the power dissipation of electronic circuits. Until recently, there has not been a major focus on a design methodology of digital circuits which directly addresses power reduction, with the focus rather on ever faster clock rates and logic speeds. The emphasis of this work is on minimizing the power consumption in DSP applications, which have two important attributes that make then *different* from general purpose computation: throughput constrained computing and correlation in the data being processed.

The throughput constraint implies that once the required sample rate is met by the hardware, there is no advantage in making the computation faster (unlike general purpose computation where is goal is to compute as fast as possible). This enables a variety of voltage scaling techniques that lower power while maintaining a constant throughput [1]. The second attribute of signal processing systems that can be exploited for low-power is the correlation present in the data being processed. Unlike general purpose computation, the data being processed in many signal processing applications can exhibit high temporal correlation. An important goal of low-power architecture selection is to preserve naturally occurring signal correlations and minimize switching activity.

The focus of this paper is on CMOS circuits and it is well known that power is primarily dissipated in charging and discharging parasitic capacitors and is given by $\alpha_{0->1} \bullet C_L \bullet V_{dd}^2 \bullet f_{clk}$, where C_L is the load capacitance, V_{dd} is the power supply voltage, f_{clk} is the clock frequency, and $\alpha_{0->1}$ is the node transition activity factor (or the fraction of time the node makes a power consuming transition inside a clock period). For example, Figure 1 shows the activity factor, $\alpha_{0->1}$, for a simple 2-input NOR gate as a function of the input signal probabilities, p_a and p_b.

$$\alpha_{0->1} = p_0\, p_1 = (1-(1-p_a)\,(1-p_b))\,(1-p_a)\,(1-p_b) \quad \textbf{(EQ 1)}$$

The strong influence of signal statistics on power implies that the data statistics must be carefully considered during power optimization. Signal statistics should be considered at all levels of the design including the selection of data representation, resource allocation, operation sequencing, logic structuring, and place and route.

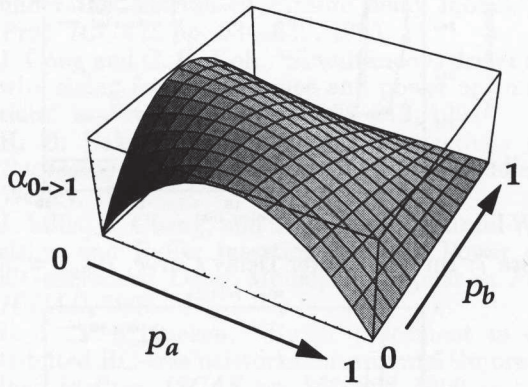


Figure 1: Dependence of power on signal statistics.

2. Optimizing Data Representation

Two's complement is typically chosen to represent numbers in signal processing applications since arithmetic operations are easy to perform. One of the problems with two's complement representation is sign-extension, which causes the msb sign-bits to toggle when a signal changes sign (for example, going from

1063-9667/95 $04.00 © 1995 IEEE

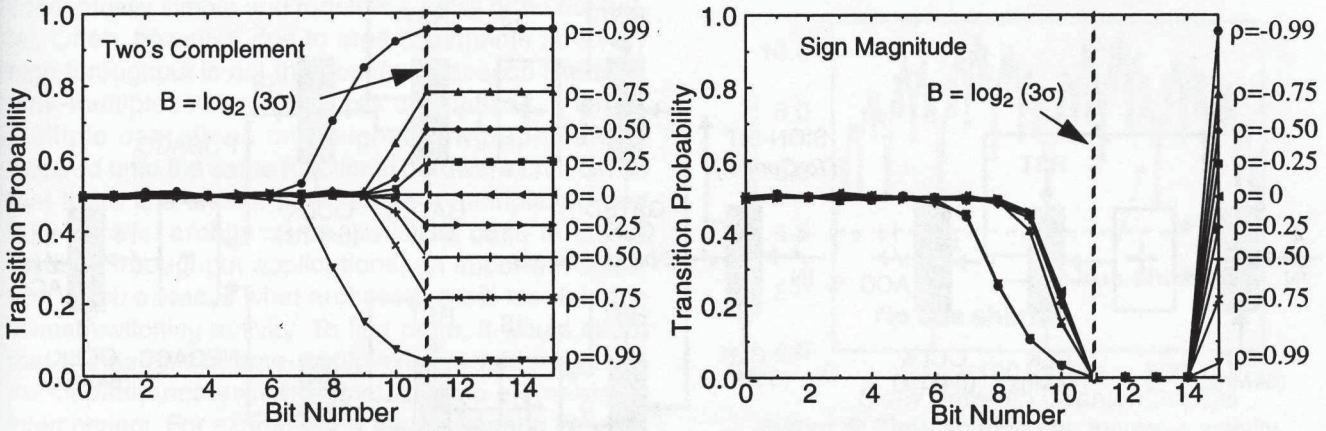


Figure 2: Reducing switching activity by optimizing number representation.

-1 to 0 will result in all of the bits toggling). This is especially significant when the dynamic range of the signals being processed is much smaller than the peak value for the bit-width used since a lot of the msb bits will perform sign-extension. Even when a signal spans the entire bit-width, scaling operations can reduce the dynamic range.

An approach to minimizing the switching activity in the msbs is to use a sign-magnitude representation, in which only one bit is allocated for the sign and the rest to the magnitude [1]. In this case, only one of the msb bits will toggle when the signal switches sign, as opposed to the two's complement representation where due to sign-extension several of the bits will switch. To illustrate this, consider Gaussian data applied to a 16-bit data-bus, which has a mean of 0 and a $3\sigma = 2^{11}$ (the upper breakpoint which represents the dynamic range of the signal lies at $\log_2(3\sigma) = 11$ bits in this case). Figure 2a shows the transition probabilities (both the 0->1 and 1->0) vs. bit position number for two's complement representation as a function of the first order correlation coefficient. A ρ=0 indicates that the data is uncorrelated and therefore the transition probability for the msb's is 1/2 (i.e., there is a 50% chance the output is going to transition between positive and negative). A large positive correlation coefficient implies that the signal changes very slowly and therefore switches sign very infrequently. Similarly, a negative correlation coefficient with a large magnitude implies that the signal changes sign frequently. In two's complement representation, bits 11 to 15 are used for sign-extension. If a sign-magnitude representation is used (Figure 2b), the higher order bits will have a low transition activity and bit 15 (the sign bit) will have the same activity as the two's complement case.

The reduction in the number of transitions for sign-magnitude over two's complement is a function of both the dynamic range of the signal and the signal correlation coefficient. Figure 3 shows the ratio of the number of transitions required in the two's complement representation to that required in the sign-magnitude representation as a function of the signal correlation and normalized dynamic range. The normalized dynamic range of a signal is defined as $3 * \sigma$ / Peak Amplitude and the number of transitions is the sum of the transitions for all the bits per clock cycle (i.e., number of transitions = $\sum \alpha_i$, $0 \leq i \leq 15$). The biggest advantage for using sign-magnitude is when the dynamic range is small and the signal has a large negative correlation. Sign-magnitude is clearly useful for driving large busses, where the overhead for converting back to two's complement is quite insignificant compared to the reduction in capacitance switched in the large busses. The scheme presented here is only one example of coding data for reducing power. There are several other schemes which can be used such as log representation, gray-coding [2], conditional inversion coding [3], etc.

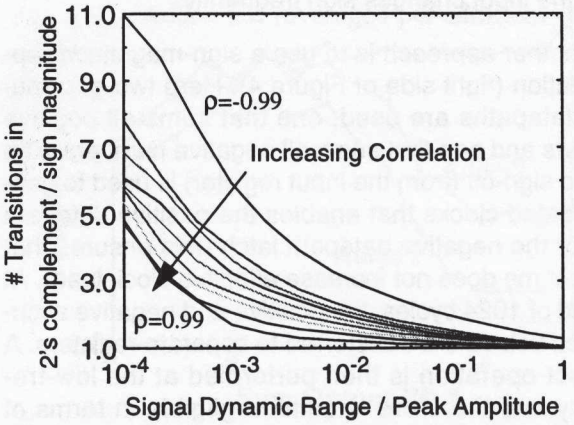


Figure 3: Activity reduction using sign-magnitude.

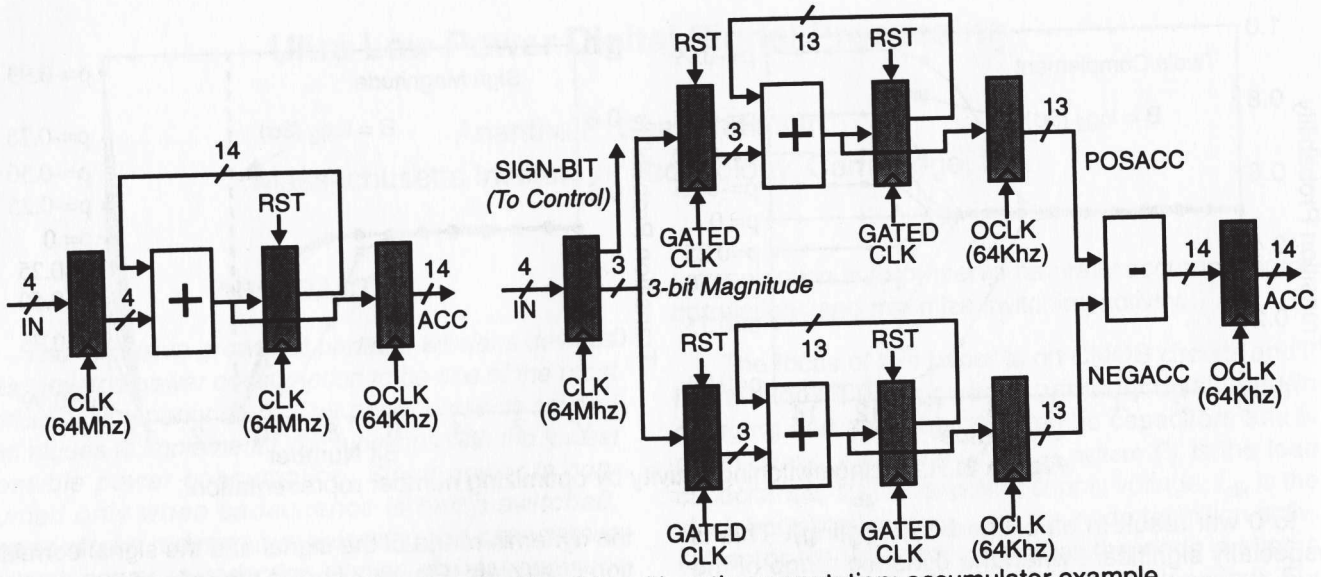


Figure 4: Number Representation trade-off in arithmetic computation: accumulator example.

While sign-magnitude has some advantages in terms of the number of transitions on busses, addition/subtraction operations are difficult to implement. However, in several applications, that require multiple additions inside a sample period (e.g., Multiply Accumulate Units or DSP filters), an architecture optimization can be used that trades silicon area for lower switching activity without altering the throughput. To illustrate this consider a CDMA radio correlator example in which the correlation length is 1024; the samples, whose values ranges from -7 to +7, are accumulated at 64Mhz.

On the left side of Figure 4, is a conventional architecture for an accumulator that uses two's complement representation. The accumulated result is transferred to an accumulator register at 64Khz. In this architecture the msb-bit, bit 3, is tied to bits 4-13 of the adder input for sign-extension and therefore the adder has high switching activity in the msb-bits (due to glitching) when the input changes sign frequently.

Another approach is to use a sign-magnitude representation (right side of Figure 4). Here two accumulator datapaths are used, one that sums all positive numbers and one that sums all negative numbers. The latched sign-bit (from the input register) is used to generate gated clocks that enables the positive datapath latch or the negative datapath latch (this ensures that this scheme does not increase effective clock load). At the end of 1024 cycles, the positive and negative accumulated values are transferred to separate registers. A subtract operation is then performed at the low-frequency and therefore is quite negligible in terms of capacitance overhead. In fact, the higher order bits (3-13) only need an incrementer (as opposed to a full-adder in the two's complement case), reducing the number of gates that are switched in the accumulator. The sign-magnitude implementation, however, requires control circuitry (overhead capacitance) to generate the timing signals for the various latches in the implementation.

By keeping the computation for positive data and negative data separate, the power is not very sensitive to rapid fluctuations in the input data. Table 1 shows the power numbers for various input patterns. Again, the biggest advantage is when the sign toggles frequently. For the case when the input changes very slowly, the sign-magnitude implementation consumes more power (15%) due to the capacitance switched by the control overhead circuitry. For the radio example, data coming in to the accumulator tends to be random.

Table 1. Number representation trade-off for arithmetic.

Input Pattern (1024 cycles)	Two's Complement Power, 3V	Sign Magnitude Power, 3V
Constant (IN= 7)	1.97 mW	2.25 mW
Ramp (-7,-6,...7,-7)	2.13 mW	2.43 mW
Random	3.42 mW	2.51 mW
Min -> Max -> Min (-7, +7,-7,+7,...)	5.28 mW	2.46 mW

3. Time-sharing of Resources

One strategy for implementing signal processing algorithms is the direct mapping approach, where there is a one to one correspondence between the operations on the signal flow graph and operators in the final implementation. Such an architecture style is

conceptually simple and requires a small or no controller. Often, however, due to area constraints or if very high-throughput is not the goal (e.g., speech filtering), time-multiplexed architectures are utilized in which multiple operations on a signal flowgraph can be mapped onto the same functional hardware unit. Given that there is a choice between time-multiplexed and fully-parallel architectures, as is the case in low to medium throughput applications, an important question which arises is what architecture will result in the lowest switching activity. To first order, it would seem that the degree of time-multiplexing would not affect the capacitance switched by the logic elements or interconnect. For example, if a data flowgraph has five additions to be performed inside the sample period, it would seem that there is no difference between an implementation in which one physical adder performs all five additions or an implementation in which there are five adders each performing one addition per sample period. It would seem that the capacitance switched should only be proportional to the number of times the additions were performed but this is not the case since resource sharing can destroy signal correlations and increase switching activity. To demonstrate this trade-off, consider two examples of resource sharing: sharing busses and sharing execution units.

First consider two implementations of two 8-bit counters clocked every cycle one with bus sharing and one without. In the parallel implementation, the two counters are driving two separate busses running at frequency f while the time-multiplexed implementation has a single shared bus running at twice the frequency. For the parallel case, the activity is easy to compute. The number of transitions for each bus = $\Sigma\alpha$ = 1 + 1/2 + 1/4 + ... + 1 / 128 ≈ 2 since the lsb of the counter switches every cycle, the 2nd lsb switches every other cycle, etc. Therefore, the total number of transitions per cycle is approximately 4 for the parallel implementation.

Figure 5 shows the plot of the average number of transitions for the time-shared case as a function of the skew between the counter outputs. The figure also shows the number of transitions for the non-time-shared implementation, which is independent of the counter skew. As seen from this figure, except for one value of the counter skew (when the skew is = 0), the parallel implementation has the lower switching activity. This example illustrates that time-sharing can modify the signal characteristics and cause an increase in the switching activity.

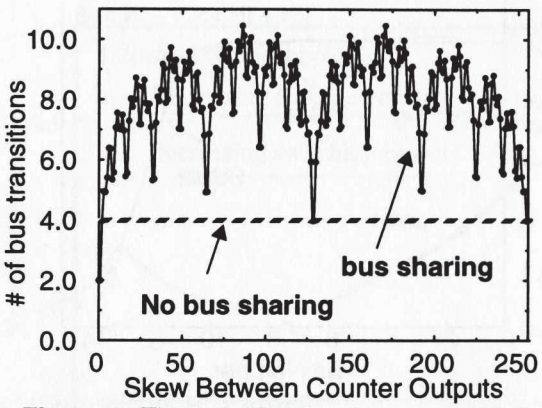


Figure 5: Time-sharing can increase activity.

The second example is a simple second order FIR filter which will demonstrate that time-multiplexing of execution units can increase the switched capacitance. The FIR filter is described as follows:

$$Y = a_0 \bullet X + a_1 \bullet X@1 + a_2 \bullet X@2 = A + B + C \quad \textbf{(EQ 2)}$$

where @ represents the delay operator and $A = a_0 \bullet X$, $B = a_1 \bullet X@1$ and $C = a_2 \bullet X@2$. Also let $O1 = A + B$. The value of the coefficients are: a_0 = 0.15625, a_1 = 0.015625, and a_2 = -0.046875.

One possible implementation might be to have two physical adders, one performing $A + B = O1$ and the other performing $O1 + C$. An alternate implementation might be to have a single time-multiplexed adder per-

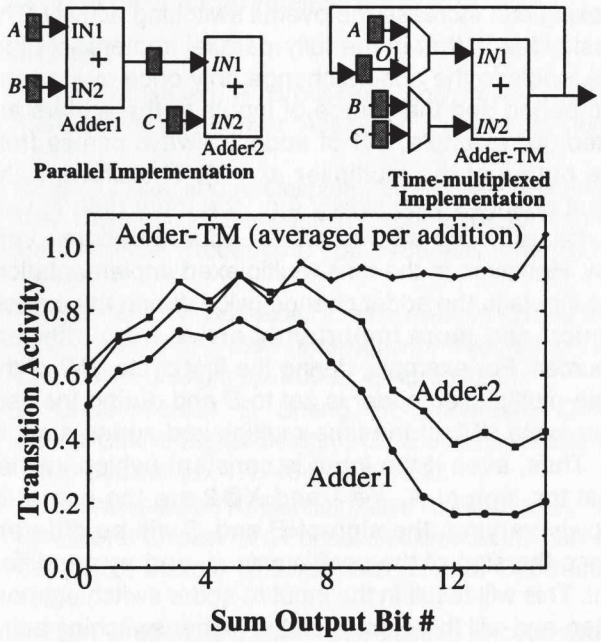


Figure 6: Activity trade-off for time-multiplexed hardware: adder example.

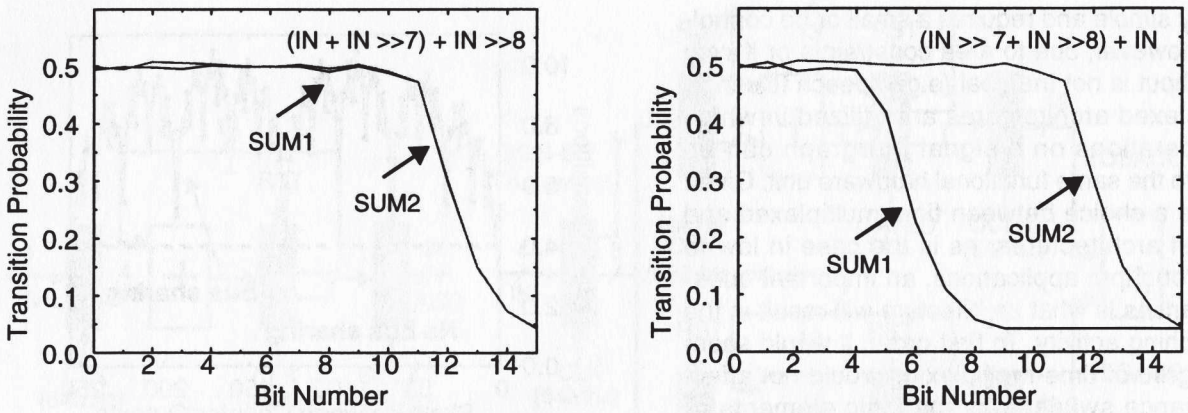


Figure 7: Reducing switching activity by re-ordering inputs.

forming both additions. So, in cycle 1, $A + B$ is performed, and in cycle 2, $O1 + C$ is performed. The topologies for the two cases are shown in Figure 6. In the first implementation, the adders can be chained and therefore the effective critical path can be reduced (chaining two ripple carry adders of N bits has a delay = $(N + 1)$ • Delay of one bit); it is clear that this will result in extra glitching activity but since the objective here is to isolate and illustrate the activity modification only due to time-multiplexing, the first implementation is assumed to be registered. The transition activities (obtained using IRSIM assuming speech input data) for the two adders for the parallel implementation and the average transition activity per addition for the time-multiplexed implementation are shown in Figure 6.

The results once again indicate that time-multiplexing can increase the overall switching activity. The basic idea is that in the fully-parallel implementation, the inputs to the adders change only once every sample period and the source of inputs to the adders are fixed (for example, *IN*1 of adder1 always comes from the output of the multiplier $a_0 * X$). Therefore, if the input changes very slowly (i.e., the input data is very correlated), the activity on the adders become very low. However, in the time multiplexed implementation, the inputs to the adder change twice during the sample period and more importantly arrive from different sources. For example, during the first cycle, *IN*2 of the time-multiplexed adder is set to B and during the second cycle, *IN*2 of the time-multiplexed adder is set to C. Thus, even if the input is constant (which implies that the sign of X, $X@1$ and $X@2$ are the same) or slowly varying, the sign of B and C will be different since the sign of the coefficients a_1 and a_2 are different. This will result in the input to adder switching more often and will therefore result in higher switching activity. That is, even if the input to the filter does not change, the adder can still be switching. For this example, the time multiplexed adder switched 50% extra capacitance per addition compared to the parallel case (even without including the input changes to adders or the multiplexor overhead).

4. Optimizing Ordering of Operation

The switching activity can be reduced by optimizing the ordering of operations in a design. To illustrate this, consider the problem of multiplying a signal with a constant coefficient and assume that the multiplication is decomposed into IN + IN >>7 + IN>>8. Figure 7 shows the results of two alternate implementations for the two required additions. In the first implementation, IN and IN >>7 are added first and the sum (SUM1) is then added to IN >>8. In this case, the SUM1 transition characteristics is very similar to that of the input IN since the amplitude of IN >>7 is much smaller than IN and the 2 inputs have identical sign-bits. Similarly SUM2 is very similar to SUM1 since SUM1 is much larger in magnitude than IN >>8. In the second implementation, the two small number IN>>7 and IN>>8 are added first and the output SUM1 is then added to IN. In this case, the output of the first adder has a small amplitude and therefore lower switching activity. The second implementation switches 30% less capacitance than the first implementation for this data.

5. Voltage Scaling

Since the dominant component of power consumption for a properly designed CMOS circuit is proportional to the square of the supply voltage, operating circuits at the lowest voltage is the key to minimizing the energy consumed per operation. However, the individual circuit elements run slower at lower supply voltages and this must be compensated for through appropriate architectural design. One important class of applications are those which have no advantage in exceeding a bounded computation rate, as found in

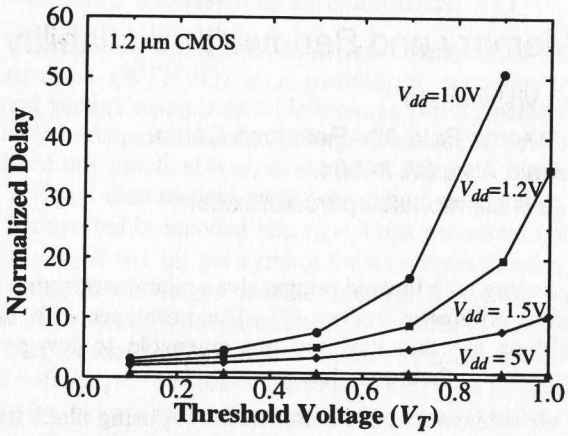


Figure 8: Effect of threshold reduction on the delay for various supply voltages.

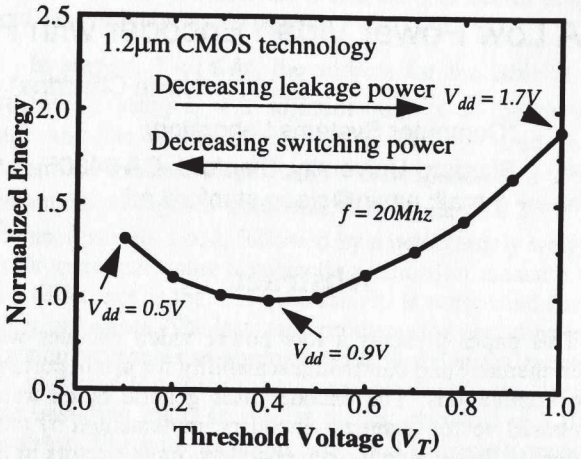


Figure 9: Compromise between dynamic and leakage power dissipation through V_t variation.

real-time signal processing. For such applications, we have previously presented an architecture driven voltage scaling approach in which parallel and pipelined architectures can be used to compensate for increased delays at reduced voltages [4][5]. As the amount of parallelism increases, the amount of overhead circuitry also increases and results in an optimum level of parallelism for low-power. The optimum voltage corresponding to this level of parallelism was found to be close to 1.5V ($\approx 2V_T$) using a standard technology with high threshold devices (V_{TN}=0.7V and V_{TP}=-0.9V).

Reducing the threshold voltage of the device allows the supply voltage to be scaled down (and therefore lower switching power) beyond the architectural optimum without loss in performance. For example, a circuit running at a supply voltage of 1.5V with V_T=1V will have approximately the same performance as the circuit running at a supply voltage of 0.9V and a V_T=0.5V. Figure 8 shows a plot of normalized delay vs. threshold voltage for various supply voltages.

Since significant power improvements can be gained through the use of low-threshold MOS devices, the question of how low the thresholds can be reduced must be addressed. Figure 9 shows a plot of energy vs. threshold voltage for a fixed throughput for a 16-bit datapath ripple carry adder (which essentially represents the power to perform the operation). Here, the power supply voltage is allowed to vary to keep the throughput fixed. For a fixed throughput (e.g., that is obtained at a 20Mhz clock rate), the supply voltage and therefore the switching component of power can be reduced while reducing the threshold voltage. However, at some point, the threshold voltage and supply reduction is offset by an increase in the leakage currents, resulting in an optimal threshold voltage for a given level of logic complexity. That is, the optimum threshold voltage must compromise between improvement of current drive at low supply voltage operation and control of the sub-threshold leakage. Circuits have been designed with power supply voltages as low as 200mV using reduced threshold devices [6].

Acknowledgments

The work presented in this paper is a result of a collaboration with Prof. Brodersen and graduate students at U.C. Berkeley. This research was performed as a part of the InfoPad wireless terminal design.

References

[1] A. Chandrakasan, R. Brodersen, *Low Power Digital CMOS Design*, Kluwer Academic Publishers, July 1995.

[2] C. Su, C Tsui, and A. Despain, "Low-power Architecture Design and Compilation Techniques for High-Performance Processors," pp. 489-498, *Compcon 1994*.

[3] M. Stan and W. Burleson, "Limited-weight Codes for Low-power I/O," *1994 International Workshop on Low-power Design*, pp. 209-214, April 1994.

[4] A. Chandrakasan, S. Sheng, R. Brodersen, "Low-Power CMOS Digital Design", *IEEE Journal of Solid-State Circuits*, pp. 473-484, April 1992.

[5] A. Chandrakasan, A. Burstein, and R. Brodersen, "A Low-power Chipset for a Portable Multimedia I/O Terminal," *IEEE Journal of Solid-State Circuits*, December 1994.

[6] J. Burr, J. Shott, "A 200mV Self-Testing Encoder/Decoder using Stanford Ultra-low Power CMOS", *IEEE ISSCC*, pp. 84-85, 1994.