


DIGITAL INTEGRATED CIRCUITS

A DESIGN PERSPECTIVE

SECOND EDITION



JAN M. RABAEY
ANANTHA CHANDRAKASAN
BORIVOJE NIKOLIĆ

PRENTICE HALL ELECTRONICS AND VLSI SERIES
CHARLES G. SODINI, SERIES EDITOR

Library of Congress Cataloging-in-Publication Data on file.

Vice President and Editorial Director, ECS: *Marcia J. Horton*

Publisher: *Tom Robbins*

Editorial Assistant: *Eric Van Ostenbridge*

Vice President and Director of Production and Manufacturing, ESM: *David W. Riccardi*

Executive Managing Editor: *Vince O'Brien*

Managing Editor: *David A. George*

Production Editor: *Daniel Sandin*

Director of Creative Services: *Paul Belfanti*

Creative Director: *Carole Anson*

Art and Cover Director: *Jayne Conte*

Art Editor: *Greg Dulles*

Manufacturing Manager: *Trudy Piscioti*

Manufacturing Buyer: *Lisa McDowell*

Marketing Manager: *Holly Stark*

About the Cover: Detail of "Wet Orange," by Joan Mitchell (American, 1925–1992). Oil on canvas, 112 × 245 in. (284.5 × 622.3 cm). Carnegie Museum of Art, Pittsburgh, PA. Gift of Kaufmann's Department Store and the National Endowment for the Arts, 74.11. Photograph by Peter Harholdt, 1995.



© 2003, 1996 by Pearson Education, Inc.
Pearson Education, Inc.
Upper Saddle River, NJ 07458

The author and publisher of this book have used their best efforts in preparing this book. These efforts include the development, research, and testing of the theories and programs to determine their effectiveness. The author and publisher shall not be liable in any event for incidental and consequential damages in connection with, or arising out of, the furnishing, performance, or use of these programs.

All rights reserved. No part of this book may be reproduced, in any form or by any means, without permission in writing from the publisher.

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

ISBN 0-13-597444-5

Pearson Education Ltd., London

Pearson Education Australia Pty, Ltd., Sydney

Pearson Education Singapore, Pte. Ltd.

Pearson Education North Asia Ltd., Hong Kong

Pearson Education Canada Inc., Toronto

Pearson Educación de México, S.A. de C.V.

Pearson Education—Japan, Tokyo

Pearson Education Malaysia, Pte. Ltd.

Pearson Education Inc., Upper Saddle River, New Jersey

- The interconnect network joins the different processor modules to one another, while the input/output circuitry connects to the outside world. For a long time, interconnections were an afterthought in the design process. Unfortunately, the wires composing the interconnect network are less than ideal and present a capacitive, resistive, and inductive load to the driving circuitry. As die sizes grow larger, the length of the interconnect wires also tends to grow, resulting in increasing values for these parasitics. Today, automated or structured design methodologies are being introduced that ease the deployment of these interconnect structures. Examples include *on-chip busses*, *mesh interconnect* structures, and even complete *networks on a chip*. Some components of the interconnect network typically are abstracted away on schematic block diagrams, such as the one shown in Figure 8-2, yet are of critical importance to the well-being of the design. These include the power- and clock-distribution networks. Early planning of these “service” networks can go a long way toward ensuring the correct operation of the integrated circuit.

The structure of Figure 8-2 may be repeated many times on a single die. Figure 8-3 shows an example of a *system on a chip*, which combines all the functions needed for the realization of a complete high-definition digital TV set. It combines two processors, memory units, specialized accelerators for functions such as MPEG (de)coding and data filtering, as well as a range of




Figure 8-3 The “Nexperia” system on a chip [Philips99]. This single chip combines a general-purpose microprocessor core, a VLIW (very large instruction word) signal processor, a memory system, an MPEG coprocessor, multiple accelerator units, and input/output peripherals, as well as two system busses.

peripheral units. Other applications such as wireless transceivers or hard-disk read/write units may even include some sizable analog modules.

Choosing an effective implementation approach strongly depends upon the function of the modules under consideration. For example, memory units tend to be very regular and structured. A module compiler that stacks cells in an arraylike fashion is thus the preferred implementation approach. Controllers, on the other hand, tend to be unstructured, and other implementation approaches are desirable. The choice of the implementation strategy can have a tremendous effect on the quality of the final product. The challenge for the designer is to pick the style that meets the product specifications and constraints. What works well for one design may well be a disaster for another one.

Example 8.1 Trading Off Energy Efficiency and Flexibility

A design that embraces flexibility (or programmability) is very attractive from an application perspective. It allows for “late binding,” in which the application can still be changed after the chip has gone to fabrication. Flexibility makes it possible to reuse a single design for multiple applications, or to upgrade the firmware of a component in the field, reducing the risk for the manufacturer. In contrast, a hard-wired component is totally fixed at manufacturing time and cannot be modified afterwards.

So, why not use flexible or programmable components for every possible design? As always, there is no free lunch. Flexibility comes at a price in both performance and energy efficiency. Providing programmability means adding overhead to implementation. For example, a programmable processor uses stored instructions and an instruction decoder to make a single datapath perform multiple functions. Most designers are not aware of the large cost of flexibility. The impact is illustrated in Figure 8-4, which compares the *energy*




Figure 8-4 Trading off flexibility versus energy efficiency (in MOPS/mW or millions of operations per mJ of energy) for different implementation styles. The numbers were collected for a 0.25 μm CMOS process [Rabaey00].

priate software tools (such as compilers and debuggers for embedded processors), prediction models, and test benches. The latter are quite important because they represent the only means for the end user to verify that the module delivers the promised functionality and performance.

The design of a system on a chip is rapidly becoming an exercise in reuse at different levels of granularity. At the lowest level, we have the standard cell library; at a level higher, we have the functional modules such as multipliers, datapaths and memories; next, we have the embedded processors; and finally, the application-specific megacells. With more and more of the system functionality migrating onto a single die, it is not surprising to see that a typical ASIC consists of a blend of design styles and modules, embedding a number of hard or soft macrocells within a sea of standard cells.

Example 8.6 A Processor for Wireless Communications

Figure 8-15 shows an integrated circuit implementing the protocol stack for a wireless indoor communication system [Silva01]. The majority of the area is occupied by the embedded microprocessor (the Tensilica Xtensa processor [Xtensa01]) and its memory system. This processor allows for a flexible implementation of the higher levels of the protocol stack (Application/Network), and enables changes in the functionality of the chip, even after fabrication. The memory modules are generated using module compilers provided by the process vendor. The processor core itself is automatically generated from a higher level description in Verilog, and uses standard cells for its physical implementation. The advantage of using the “soft-core” approach is that the processor instruction set can be




Figure 8-15 Wireless communications processor—an example of a hybrid ASIC design methodology. The processor combines an embedded microprocessor and its memory system with dedicated hardware accelerators and I/O modules. Observe also the on-chip network module [Silva01].

While both approaches can be tuned to accomplish approximately the same performance levels, the current-mode approach holds a definite edge in terms of (dynamic) power dissipation. Due to its immunity to power-supply noise, it can operate at a much lower noise margin than the voltage-mode network, and thus at a much lower swing as well. The value of the voltage wave propagating over the transmission line can be as low as 100 mV. The main challenge in designing an efficient current-mode circuit is the static power consumption. This is not an issue in ultra high-speed networks, where dynamic power tends to play the major role.

Current-mode CMOS transmission systems have become quite popular in the area of high-speed off-chip interconnections. The ferocious appetite for I/O bandwidth in high performance processors and memories has made this a topic of considerable importance. Readers interested in high-speed interchip signaling are referred to [Dally98] and [Sidiropoulos01].

9.6 Perspective: Networks-on-a-Chip

With chip sizes and clock rates continuing to increase, and feature sizes getting smaller at a constant rate, it is not hard to see that interconnect problems will be with us for quite some time to come. Physics constraints such as the speed of light and thermal noise determine how fast and how reliably communication can be established over “long” distances. As communication bottlenecks have been the major damper on hyper-supercomputer performance, similar constraints are starting to hamper integrated systems (“systems-on-a-chip”). New solutions conceived at the technology and circuit level only help to postpone the problems temporarily. Ultimately, the solution is to address interconnections on a chip as a communication problem and to apply techniques and approaches that have made large-scale communications and networking systems around the world operate reliably and correctly for many years. For example, it is somewhat amazing that the Internet is working correctly given its span and number of connection points. The secret to its success is a well-thought-out protocol stack, which isolates and orthogonalizes the various functionality, performance, and reliability concerns. Hence, rather than considering on-chip interconnections as point-to-point wires, they should be abstracted as communication channels over which we want to transmit data within a “quality-of-service” setting, putting constraints on through-put, latency, and correctness [Sgroi01].

As an example, today’s on-chip interconnect signalling techniques are designed with noise margins large enough to ensure that a bit will always be transmitted correctly. It is probably safe to say that future designs might take a more extreme tack—they may give up integrity entirely for the sake of energy/performance, and thus allow errors to occur on the transmitted signals. These errors can be corrected by other circuitry that is designed to provide error-correcting capabilities and/or retransmit corrupted data.

On a higher level of abstraction, we are already witnessing the emergence of complete “networks” on a chip, as shown in Figure 9-38. Rather than being statically wired from source-to-destination, data is injected as packets into a complete network of wires, switches, and routers, and it is the network that dynamically decides how and when to route these packets through




Figure 9-38 The network-on-a-chip combines multiple processors and their interconnect network on the same die.

its segments [Dally01]. Ultimately, this is the only approach that can work reliably when the discrepancy between device sizes and on-chip distances becomes macroscopic.

9.7 Summary

This chapter has introduced a number of techniques to cope with the impact of interconnect on the performance and reliability of digital integrated circuits. The parasitics introduced by the interconnect have a dual effect on circuit operation: (1) They introduce noise and (2) they increase the propagation delay and power dissipation. In summary,

- *Capacitive cross talk* in dense wire networks affects the reliability of the system and therefore its performance. Careful design using well-behaved regular structures or using advanced design-automation tools is a necessity. Providing the necessary shielding is important for wires such as busses or clock signals.
- Driving large capacitances rapidly in CMOS requires the introduction of a *cascade of buffer stages* that must be carefully sized. More advanced techniques include the lowering of the signal-swing on long wires, and the use of current-mode signalling.
- Resistivity affects the reliability of the circuit by introducing *IR drops*. This is especially important for the *supply network*, where wire sizing is important.
- The extra delay introduced by the *rc* effects can be minimized by *repeater insertion* and by using a better interconnect technology.
- The *inductance of the interconnect* becomes important at higher switching speeds. The chip package is currently one of the most important contributors of inductance. Novel packaging techniques are gaining importance with faster technologies.
- *Ground bounce* introduced by the *L di/dt* voltage drop over the supply wires is one of the most important sources of noise in current integrated circuits. Ground bounce can be