
ENCYCLOPEDIA OF COMPUTER SCIENCE AND TECHNOLOGY

EXECUTIVE EDITORS

Allen Kent James G. Williams

UNIVERSITY OF PITTSBURGH
PITTSBURGH, PENNSYLVANIA

ADMINISTRATIVE EDITOR

Carolyn M. Hall

ARLINGTON, TEXAS

VOLUME 36
SUPPLEMENT 21



MARCEL DEKKER, INC.

NEW YORK • BASEL • HONG KONG

COPYRIGHT © 1997 BY MARCEL DEKKER, INC.
ALL RIGHTS RESERVED

Neither this book nor any part may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, microfilming, and recording, or by any information storage and retrieval system, without permission in writing from the publisher.

MARCEL DEKKER, INC.
270 Madison Avenue, New York, New York 10016

LIBRARY OF CONGRESS CATALOG CARD NUMBER: 74-29436
ISBN: 0-8247-2289-2

Current Printing (last digit)
10 9 8 7 6 5 4 3 2 1

PRINTED IN UNITED STATES OF AMERICA

DESIGN TECHNOLOGIES FOR LOW-POWER VLSI

MOTIVATION

In the past, the major concerns of the very-large-scale integration (VLSI) designer were area, performance, cost, and reliability; power consideration was mostly of only secondary importance. In recent years, however, this has begun to change and, increasingly, power is being given comparable weight to area and speed considerations. Several factors have contributed to this trend. Perhaps the primary driving factor has been the remarkable success and growth of the class of personal computing devices (portable desktops, audio- and video-based multimedia products) and wireless communications systems (personal digital assistants and personal communicators) which demand high-speed computation and complex functionality with low-power consumption.

In these applications, average power consumption is a critical design concern. The projected power budget for a battery-powered, A4 format, portable multimedia terminal, when implemented using off-the-shelf components not optimized for low-power operation, is about 40 W. With advanced nickel-metal-hydride (secondary) battery technologies offering around 65 W hr/kg (1), this terminal would require an unacceptable 6 kg of batteries for 10 hr of operation between recharges. Even with new battery technologies such as rechargeable lithium ion or lithium polymer cells, it is anticipated that the expected battery lifetime will increase to about 90–110 W hr/kg over the next 5 years (1) which still leads to an unacceptable 3.6–4.4 kg of battery cells. In the absence of low-power design techniques, current and future portable devices will suffer from either very short battery life or very heavy battery pack.

There also exists strong pressure for producers of high-end products to reduce their power consumption. Contemporary performance optimized microprocessors dissipate as much as 15–30 W at 100–200 MHz clock rates (2)! In the future, it can be extrapolated that a 10-cm² microprocessor, clocked at 500 MHz (which is a not too aggressive estimate for the next decade) would consume about 300 W. The cost associated with packaging and cooling such devices is prohibitive. Because core power consumption must be dissipated through the packaging, increasingly expensive packaging and cooling strategies are required as chip power consumption increases. Consequently, there is a clear financial advantage to reducing the power consumed in high-performance systems.

In addition to cost, there is the issue of reliability. High-power systems often run hot, and high temperature tends to exacerbate several silicon failure mechanisms. Every 10°C increase in operating temperature roughly doubles a component's failure rate (3). In this context, peak power (maximum possible power dissipation) is a critical design factor, as it determines the thermal and electrical limits of

designs; impacts the system cost, size and weight, dictates specific battery type, component and system packaging, and heat sinks, and aggravates the resistive and inductive voltage-drop problems. It is therefore essential to have the peak power under control.

Another crucial driving factor is that excessive power consumption is becoming the limiting factor in integrating more transistors on a single chip or on a multiple-chip module. Unless power consumption is dramatically reduced, the resulting heat will limit the feasible packing and performance of VLSI circuits and systems.

From the environmental viewpoint, the smaller the power dissipation of electronic systems, the lower the heat pumped into the rooms, the lower the electricity consumed and, hence, the lower the impact on global environment, the less the office noise (e.g., due to elimination of a fan from the desktop), and the less stringent the environment/office power delivery or heat-removal requirements.

The motivations for reducing power consumption differ from application to application. In the class of micropowered battery-operated, portable applications, such as cellular phones and personal digital assistants, the goal is to keep the battery lifetime and weight reasonable and the packaging cost low. Power levels below 1–2 W, for instance, enable the use of inexpensive plastic packages. For high-performance, portable computers, such as laptop and notebook computers, the goal is to reduce the power dissipation of the electronics portion of the system to a point which is about half of the total power dissipation (including that of display and hard disk). Finally, for high-performance, nonbattery operated systems, such as workstations, desk-top computers, and multimedia digital signal processors, the overall goal of power minimization is to reduce system cost (cooling, packaging, and energy bill) while ensuring long-term device reliability. These different requirements impact how power optimization is addressed and how much the designer is willing to sacrifice in cost or performance to obtain lower power dissipation.

The next question is to determine the objective function to minimize during low-power design. The answer varies from one application domain to the next. If extending the battery life is the only concern, then the energy (i.e., the power-delay product) should be minimized. In this case, the battery consumption is minimized even though an operation may take a very long time. On the other hand, if both the battery life and the circuit delay are important, then the energy-delay product must be minimized (4). In this case, one can alternatively minimize the energy/delay ratio (i.e., the power) subject to a delay constraint. In most design scenarios, the circuit delay is set based on system-level considerations, and hence during circuit optimization, one minimizes power under user-specified timing constraints.

SOURCES OF POWER DISSIPATION

Power dissipation in Complementary Metal-Oxide-Silicon (CMOS) circuits is caused by three sources: (1) The leakage current which is primarily determined by the fabrication technology, consists of reverse bias current in the parasitic diodes formed between source and drain diffusions and the bulk region in a MOS transistor as well as the subthreshold current that arises from the inversion charge that exists at the gate voltages below the threshold voltage; (2) the short-circuit (rush-through)

current which is due to the DC path between the supply rails during output transitions; and (3) the charging and discharging of capacitive loads during logic changes.

The diode leakage occurs when a transistor is turned off and another active transistor charges up or down the drain with respect to the first transistor's bulk potential. The resulting current is proportional to the area of the drain diffusion and the leakage current density. The diode leakage is typically 1 pA for a 1- μm minimum feature size. The subthreshold leakage current for long-channel devices increases linearly with the ratio of the channel width over channel length and decreases exponentially with $V_{\text{GS}} - V_t$, where V_{GS} is the gate bias and V_t is the threshold voltage. Several hundred millivolts of "off bias" (say, 300–400 mV) typically reduces the subthreshold current to negligible values. With reduced power supply and device threshold voltages, the subthreshold current will, however, become more pronounced. In addition, at short channel lengths, the subthreshold current also becomes exponentially dependent on drain voltage V_{DS} instead of being independent of V_{DS} , (see Ref. 5 for a recent analysis). The subthreshold current will remain 10^2 – 10^5 times smaller than the "on current," even at submicron device sizes.

The short-circuit (crowbar current) power consumption for an inverter gate is proportional to the gain of the inverter, the cubic power of supply voltage minus device threshold, the input rise/fall time, and the operating frequency (6). The maximum short-circuit current flows when there is no load; this current decreases with the load. If gate sizes are selected so that the input and output rise/fall times are about equal, the short-circuit power consumption will be less than 15% of the dynamic power consumption. If, however, the design for high performance is taken to the extreme where large gates are used to drive relatively small loads, then there will be a stiff penalty in terms of short-circuit power consumption.

The short-circuit and leakage currents in CMOS circuits can be made small with proper circuit and device design techniques. The dominant source of power dissipation is thus the charging and discharging of the node capacitances (also referred to as the dynamic power dissipation) and is given by

$$P = 0.5CV_{\text{dd}}^2E(\text{sw})f_{\text{clk}} \quad [1]$$

where C is the physical capacitance of the circuit, V_{dd} is the supply voltage, $E(\text{sw})$ (referred as the *switching activity*) is the average number of transitions in the circuit per $1/f_{\text{clk}}$ time, and f_{clk} is the clock frequency.

LOW-POWER DESIGN SPACE

The previous section revealed the three degrees of freedom inherent in the low-power design space: voltage, physical capacitance, and data activity. Optimizing for power entails an attempt to reduce one or more of these factors. This section briefly discusses each of these factors, describing their relative importance, as well as the interactions that complicate the power optimization process.

Voltage

Because of its quadratic relationship to power, voltage reduction offers the most effective means of minimizing power consumption. Without requiring any special circuits or technologies, a factor of 2 reduction in supply voltage yields a factor of 4

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.