Chapter 8
Energy-Efficient Design of High-Speed Links

Gu-Yeon Wei[1], Mark Horowitz[2], Jaeka Kim[2]

[1]*Harvard University;* [2]*Stanford University*

Abstract: Techniques for reducing power consumption and bandwidth limitations of inter-chip communication have been getting more attention to improve the performance of modern digital systems. This chapter begins with a brief overview of high-speed link design and describes some of the power vs. performance trade-offs associated with various design choices. The chapter then investigates various techniques that a designer may employ to reduce power consumption. Three examples of link designs and link building blocks found in the literature present energy-efficient implementations of these techniques.

Key words: High-speed I/O, serial links, parallel links, phase-locked loop, delay-locked loop, clock data recovery, low-power, energy-efficient, power-supply regulator, voltage scaling, digital, mixed-signal, CMOS.

8.1 INTRODUCTION

Aggressive CMOS technology scaling has enabled explosive growth in the integrated circuits (IC) industry with cheaper and higher-performance chips. However, these advancements have led to chips being limited by the chip-to-chip data communication bandwidth. This limitation has motivated research in the area of high-speed links that interconnect chips [1] [2][10][3][4] and has enabled a significant increase in achievable inter-chip communication bandwidths. Enabling higher I/O speed and more I/O channels improves bandwidth, but this can also increase power consumption, which eats into the overall power budget of the chip. Furthermore, complexity and area become major design constraints when trying to potentially integrate hundreds of links on a single chip. Therefore, there is a need for building energy-efficient high-speed links with low design complexity.

Power in synchronous CMOS digital systems is dominated by dynamic power dissipation, which is governed by the following well-known equation:

$$P_{DYNAMIC} = \alpha C_{SW} V_{dd} V_{SWING} F_{CLK} \quad (8.1)$$

where α is the switching activity, C_{SW} is the total switched capacitance, V_{DD} is the supply voltage, V_{SWING} is the internal swing magnitude of signals (usually equals V_{DD} for most CMOS gates), and F_{CLK} is the frequency of operation. And since power is the rate of change of energy,

$$E_{DYNAMIC} = \alpha C_{SW} V_{DD} \cdot V_{SWING} \quad (8.2)$$

Power consumption in analog circuits is simply set by the static current consumed such that $P_{STATIC} = V_{dd} \cdot I_{STATIC}$. Technology scaling enables lower power and energy in digital systems since the next generation process scales both capacitance and voltage. Transistors also get faster; thus it is possible to run a scaled chip at higher frequencies while still dissipating less power.

Aside from technology scaling, reducing just the supply voltage for a given technology enables significant reduction in digital power and energy consumption since both are proportional to the supply voltage squared. However, voltage reduction comes at the expense of slower gate speeds. So, there is a trade off between performance and energy consumption. Recognizing this relationship between supply voltage and circuit performance, dynamically adjusting the supply voltage to the minimum needed to operate at a desired operating frequency enables one to reduce the energy consumption down to the minimum required. This technique is referred to as adaptive power-supply regulation and requires a mechanism that tracks the worst-case delay path through the digital circuitry with respect to process, temperature, and voltage in order to determine the minimum supply voltage required for proper operation. Although it was first applied to digital systems, adaptive supply regulation can also enable energy-efficient high-speed link design. It is one of several energy-reduction techniques that will be investigated in this chapter.

The design of energy-efficient links relies on optimizing all components of the interface. This optimization requires an analysis of each component comprising the link and making the right power/performance trade-offs. In order to understand these trade-offs, Section 2 presents an overview of link design. Then, Section 3 investigates several approaches used in digital systems that can also be applied to build energy-efficient links. It begins with concepts utilizing parallelism to reduce power consumption. Subsequently, an adaptive supply-regulation technique is introduced that offers a scheme for optimizing energy consumption in the overall link architecture. Section 4 presents implementation details of various test chips