

Analysis, Design and Implementation of High Speed Electrical Interconnects in Si VLSI

Edward J. Ouellette & Thomas F. Knight, Jr.

Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

<http://www.ai.mit.edu>



The Problem: Interconnect delays have become the major limit in VLSI system speeds today. For frequencies in and above the GHz range on chip, on-chip electrical interconnect bandwidth is limited by the DC resistance, skin effect (AC resistance), capacitance, and dielectric dispersion of wires. [1, 2]. Also, the inductance of long on-chip wires can no longer be ignored when switching frequencies are over several GHz. The speed of transistors in modern ICs increases substantially with every improving generation; however the raw, unbuffered speed of the electrical interconnect wiring is reaching a plateau, limited by material, electrical and geometry constraints. The intrinsic RC time constant of electrical wires stays approximately constant throughout process scaling; the l^2 dependence of the interconnect delay on the length, l , of the connection does not go away for chips with smaller feature sizes. In fact, since today's chips are becoming more complex, the average interconnect distance is actually increasing. While pipelining may seem like a simple solution to this problem, the latency penalty simply cannot be tolerated for some critical signals. As minimum wire widths and pitches decrease in more advanced technologies, wire heights are increased to provide lower parasitic resistance, to increase current capacity and to reduce delays. However, this steep wire aspect ratio translates into greater mutual capacitance coupling between adjacent wires, increasing crosstalk noise. Simultaneous switching noise has also increased due to faster signalling rates and increased complexity of modern ICs. Power supply distribution for both on-chip and in chip packaging has also become a major design issue because of wire limitations. This subsystem must be carefully designed to supply a low-impedance, low-inductance, and low-noise power source to noisy digital circuits, preventing them from interfering with each other and causing bit errors. As chips operate faster, and at lower voltages, noise margins become ever harder to maintain, making the hardware designer's job more challenging.


Motivation: A set of applications exists whereby a class of speed-critical signals (e.g. a pipeline stall signal in a microprocessor) needs to be distributed globally, and the only thing that matters (next to power) is speed. Since electrical interconnection delays have become the significant factor determining overall Si VLSI system speeds, new approaches for distributing high-speed, low-power, low-noise signals on and off chip must be developed [3].

Approach: Most VLSI designers are relying upon process improvements (e.g. low k dielectrics, copper interconnects, SOI technology, feature scaling, SiGe, etc.) and architecture changes only to achieve performance gains needed for the next generation of VLSI technology. However, these improvements will only go so far, and more significant improvements can be obtained by a thorough investigation into the propagation mechanisms of signals on electrical wires, limitations on signal rates, and noise sources. Only then can one design the best circuits and electrical conduits which provide the greatest signalling rates at the lowest power and noise levels.


The traditional approach to sending high-speed signals across large distances on VLSI chips involves a chain of equally-spaced, high-current voltage-mode repeaters in a conventional CMOS design style [4]. New dynamic impedance circuits and low-noise, low-capacitance wire structures are being investigated, which allow signals to propagate down long RC-delay dominated signal lines at speeds much higher than simple repeated lines, and with better signal integrity.

Figure 1 shows simple models for electrical interconnects. Electrical interconnects are divided into RC-dominated ($l < \frac{\lambda}{8}$, where l is the length of the wire, and λ is the wavelength), and transmission line or LC-dominated wires ($l > \frac{\lambda}{8}$). R' , C' and L' are the per-unit-length resistance, capacitance, and inductance, respectively, of the wires. The transmission line shown has intrinsic impedance, Z_0 .

A new differential regenerator circuit and on-chip differential transmission waveguide wire geometries are be-



1. Uniform LCR Line and simplified schematic symbol



2. Uniform RC Line and simplified schematic symbol

Figure 1: Simple Models of Electrical Interconnections

ing combined to give rise to active impedance controlled lines with reduced voltage swings, suitable for high-speed, low-power communication across large chips. The combined effects of small, low-noise signal swings, differential signalling, and active guard-banding results in several orders of magnitude reduction in power consumption and between 10-100 times faster signal propagation compared to simple repeated lines on chip. A differential-transimpedance receiver maintains near-constant voltage on each line segment, forcing the transmission line to operate in current mode effecting a "negative capacitance" and readily rejecting common-mode noise. The driver circuit is small – only large enough to supply the required communication current – in contrast to other voltage-mode drivers, which must worry about swinging large voltage signals on large capacitances in short time periods, wasting valuable power. The circuits also incorporate high power supply rejection, PSRR, tolerating upto 20% VDD RMS noise, and low power supply noise generation, making them ideal for use in noisy digital environments, while operating off a single power supply. 20Gb/s signalling rates per line have easily been simulated in a TSMC 0.13 μ m technology. Actual circuits to be fabricated are expected to yield similar results.

Impact: Cleverly designed wires and low-power circuits on chip and I/O pads off-chip can now carry signals up to 80% of c in a given VLSI technology, without requiring special process modifications.

Future Work: Fab batch of VLSI test chips integrating on-chip test circuits for 20Gb/s signalling, clock recovery, and combining both source and line encoding for higher bitrates for off-chip interconnects.

Research Support: Support for this research was provided in part by Nippon Telegraph and Telephone: Letter of Agreement Date 6/21/1998.

References:

- [1] A. Deutsch, G. Kopcsay, V. Ranieri, J. Catalado, et al. Electrical characteristics of lossy interconnections for high-performance computer applications. *SPIE Proceedings on Microelectronic Interconnects and Packages: Optical and Electrical Technologies*, 1389:161–176, November 1990.
- [2] H. H. Kim. Can optical interconnects improve electrical interconnects of digital processors? Technical Report 11126-910522-05TM, AT&T Bell Labs, June 1991.
- [3] Edward J. Ouellette. *Analysis, Synthesis, and Fabrication of VLSI Si Detector Arrays for Optoelectronic Interconnections*. Massachusetts Institute of Technology SM Thesis. MIT Department of Electrical Engineering and Computer Science, May 1994.
- [4] Jan M. Rabaey. *Digital Integrated Circuits: A Design Perspective*. Prentice Hall Electronics and VLSI Series. Prentice Hall, 1996.