Integrated Systems

Academic and Research Staff
Professor Vladimir Stojanović

Collaborators
Prof. Lizhong Zheng, Prof. Joel Dawson

Graduate Students
Nataša Blitvić, Fred Chen, Byungsub Kim, Maxine Lee, Sanquan Song, Ranko Sredojević,

Administrative Staff
Janice Balzer

Introduction

Integrated Systems Group is focused on several key design aspects of modern integrated systems. The group is focused on building cutting edge, energy-efficient integrated systems through vertical optimization encompassing communications and signal processing algorithms and architectures, and digital and mixed-signal circuits. The main research topics include modeling of noise and dynamics in circuits and systems, application of convex optimization to digital communications, analog and VLSI circuits.

Most integrated systems today are limited by tight power and/or throughput constraints. In this environment, the usual system hierarchy becomes inefficient and novel vertical system design approaches are needed to create overall most efficient integrated systems. We believe that this can be achieved by blurring of the strict hierarchy boundaries and modification of standard communication techniques and circuits to perform the operations in a most energy-efficient way.
Some of the example systems that we are interested in are high-speed electrical and optical interfaces, UWB transceivers and sensors, on-chip signaling, clock generation and distribution for system-on-a-chip.

1. Channel-and-Circuits-Aware, Energy-Efficient Coding for High-Speed Links

Sponsors
MARCO Interconnect Focus Center

Project Staff
Nataša Blitvić, Maxine Lee, Professor Vladimir Stojanović and Professor Lizhong Zheng

In order to achieve high throughput, while satisfying energy and density constraints, both the data rates and the energy efficiency of high-speed chip-to-chip interconnects need to increase. In this project we aim to extend the link system design to incorporate energy-efficient channel coding techniques. Using novel energy-efficient coding techniques for non-Gaussian noise and residual interference, we will both increase the achievable data rates and the energy-efficiency of links by drastically off-loading the low-BER target burden and hence decreasing the complexity of the equalization/modulation level (Figure 1). Presently, both a statistical simulator and an experimental setup are being developed with the purpose of streamlining the code design process. The statistical simulator will be the first link simulator to include channel coding and the effects of data correlation. The current focus is on the modeling of the residual inter-symbol interference (ISI), but the approach will be extended to deal with crosstalk, timing jitter and other circuit-related effects. Our recent developments have addressed the difficulty in computing ISI probability distributions for realistic channel lengths, in presence of data correlation in the form of a single parity bit. This approach is presently being extended to linear block codes.

The resulting simulator will provide the capability to model data correlation both as a plug-in for the existing analytical statistical link simulators, or as the basis of time-domain behavioral link simulation software. In order to mitigate the inadequacies of analytical system models (Figure 2), limited by system complexity and link-specific noise sources, we consider advanced statistical methods based on modifications of the standard Monte-Carlo technique. The generality of the Monte Carlo technique will allow us to accurately encompass the system's complexity in a behavioral time-domain framework, without resorting to overly-restrictive simplifications (like linearity) necessary in the fully analytical approach. Furthermore, the sample-size reduction techniques, such as importance sampling, coupled with conditioning through our interference calculation methods, will allow us to efficiently simulate very low target BERs, not reachable by standard Monte Carlo simulation. The promise of this approach lies in the large deviation theory and the theory of asymptotically efficient estimators.

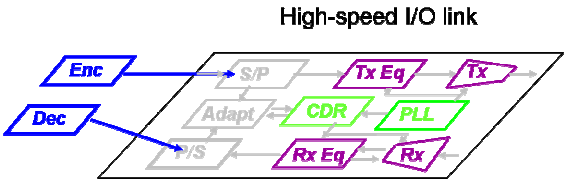

Figure 1: Model of the high-speed link where the serializer and deserializer blocks are replaced by encoder/decoder. By relaxing the target BER, channel coding will have the benefit of lowering the energy associated with timing and equalization

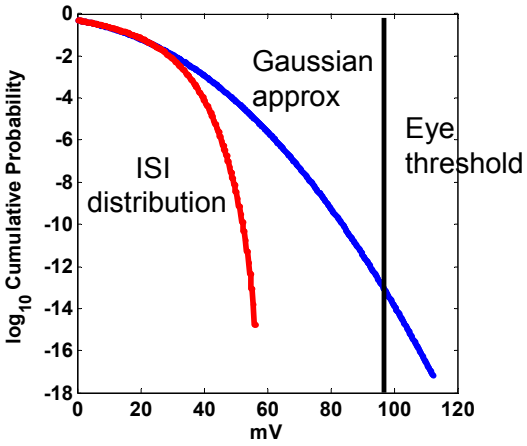

Figure 2: Error incurred in modeling the link noise by additive white Gaussian noise (AWGN). As shown, the simplification can be adequate at low BER, but becomes largely inaccurate by the time we reach the target BER range (~ 10^{-15})

References

[1] V. Stojanovic and M. Horowitz "Modeling and analysis of high-speed links," *Custom Integrated Circuits Conference, 2003. Proceedings of the IEEE 2003*, pp. 589-594, 2003.

2. Convex Optimization of Integrated Communication Systems

Sponsors
Lincoln Laboratories Advanced Concepts Committee

Project Staff
Ranko Sredojević, Tania Khanna, Professor Vladimir Stojanović and Professor Joel Dawson

In system design, allocation of circuit resources like power and noise budgets is a problem with an often unclear solution and results in long negotiations between both circuit and system designers. It is difficult to know the optimal distribution or even the feasible set of distributions of resources. This uncertainty results in an iterative approach with frequent re-design of circuit blocks for different distribution schemes. Insight into the trade-offs between resources within

each circuit block can aid in finding the optimal distribution and eliminate the need for re-design ultimately speeding up the design cycle.

Thus far, work done in analog circuit optimization has applied convex optimization techniques, specifically geometric programming (GP), in order to formulate and solve for optimality. GP is convenient because there is a specific formulation that can efficiently be solved [1]. We wish to follow the style of past circuit optimization attempts [2]-[4], but reformulate them in our more general hierarchical approach to optimize a fully integrated system.

With a hierarchical optimization, GP is used to formulate and optimize each circuit block in a given system. We stress that formulation is not trivial and requires circuit design experience for correctness. From this optimization, trade-off curves describing the performance specifications are created. Trade-off curves are then related in the system formulation to optimization performance criteria for each block, which can then retrieve the optimal system design. This optimization flow can be seen in figures 1 and 2. We also anticipate that a hierarchical approach will allow for an interchange of block topologies, which has not been allowed in the past.

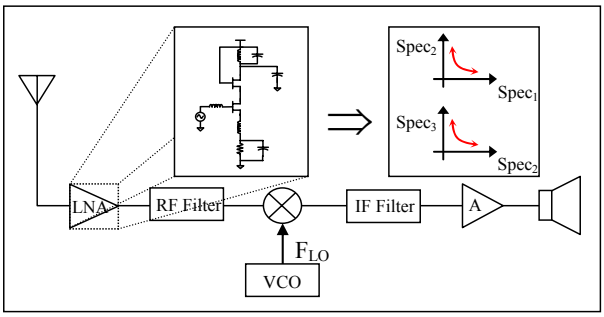

Figure 1: Circuit performance trade-off curve generation

Figure 2: System optimization for performance specifications. Circuit blocks are abstracted into trade-off curves, and performance specifications are optimized.

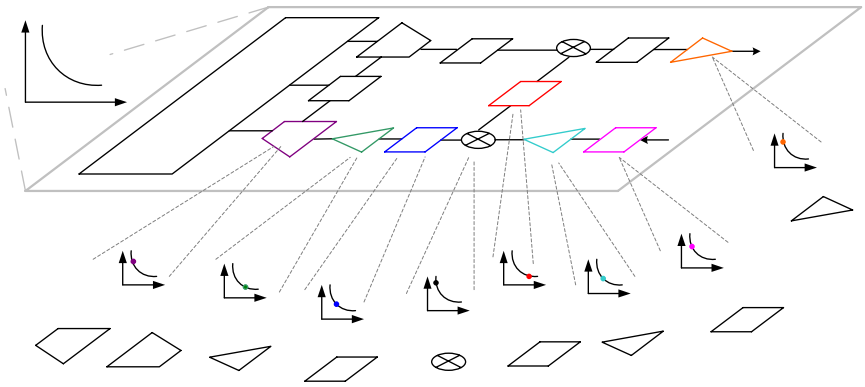

References

[1] S. Boyd, L. Vandenberghe, *Convex Optimization*. New York: Cambridge University Press, 2004.
[2] M.M. Hershenson, "Efficient Description of the Design Space of Analog Circuits," *Design Automation Conference,* 2003.
[3] M.M. Hershenson, "Design of Pipeline Analog-To-Digital Converters via Geometric Programming," ICCAD, 2002.
[4] Y. Xu, L. Pileggi, S. Boyd, "ORACLE: Optimization with Recourse of Analog Circuits including Layout Extraction," *Design Automation Conference,* June 2004.

3. Efficiency of High Speed On-Chip Interconnects: Trade-Off and Optimization

Sponsors
NEC Research Fund

Project Staff
Byungsub Kim and Professor Vladimir Stojanović

Signaling over global on-chip wires has been an increasingly difficult problem for the last several generations of VLSI technologies. As the technology scales, global wires scale poorly, causing a large increase in module-to-module communication. Traditionally, a repeater insertion [1] is used to overcome the latency problem but the power consumption of the signaling increases due to high speed requirement for the repeater. To address the limited latency and energy-efficiency of the repeater chains, alternative techniques such as RF-modulation [2] and pulse width modulation [3] have been suggested.

These past studies, however, have not considered the interconnect as a part of a dense on-chip network which must be optimized for the area-normalized metric such as cross-sectional throughput and power density instead of singe link metric such as throughput and power consumption. Given the global constraints such as power and total die area, designer must jointly optimize interconnect circuits and wires to find the best trade-off between energy dissipated in circuits and wires. In this project, we aim to establish a framework for analysis and comparison of various interconnect methods under a set of performance and cost metrics including bandwidth, latency, chip area, and power consumption.

Figure 1 shows power density versus data rate density of optimized repeater-inserted interconnect of predicted bulk 32nm CMOS process model for a given target delay-to-symbol period ratio, Nd=1, 2, 4. Figure 2 shows power density versus data rate density of optimized pulse width modulation interconnect of the same 32nm CMOS model. The latency of this point-to-point link is one bit time at highest data rate (equivalent to Nd=1 repeater case). The trade-off curves are calculated when all practical design parameters (such as driver size, wire width and space) are optimized to meet given performance specifications. The two figures show that the pulse width modulation is a more energy-efficient signaling method than repeater for comparable data rate density. Our analytical method also provides the information of best interconnect design for given performance specifications.

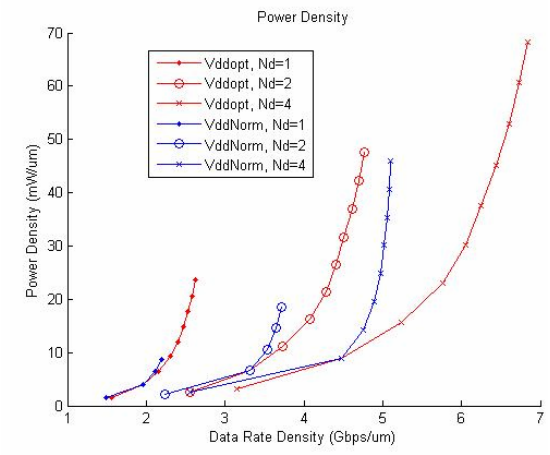

Figure 1: Power density (mW/um) versus data rate density (Gbps/um) of repeater-inserted interconnect for given delay to symbol period ratios (Nd=Td/Ts=1,2,4).

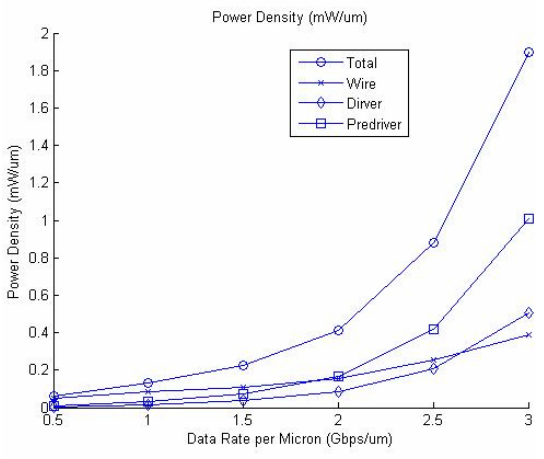

Figure 2: Power Density (mW/um) versus data rate density (Gbps/um) of PWP with one delay to symbol period ratio (Nd ~1).

References

[1] R. Ho, K. Mai, M. Horowitz, “Efficient On-Chip Global Interconnect,” Digest of Technical Papers, IEEE Int’l. Symposium on VLSI Circuits, pp. 271-4, June 2003.

[2] R.T. Chang, C.P. Yue, S.S. Wong, “Near Speed-of-light On-chip Electrical Interconnect,” Digest of Technical Papers, IEEE Int’l. Symposium on VLSI Circuits, pp. 18-21, June 2002.

[3] D. Schinkel, E. Mensink, E. Klumpernik, E. van Tuijl, B. Nauta, “A 3Gb/s Transceiver for RC-limited On-Chip Interconnects,” in Proceedings of IEEE International Solid State Circuits Conference, February 2005.

Publications

Meeting Papers, Presented

C. Werner, C. Hoyer, A. Ho, M. Jeeradit, F. Chen, B. Garlepp, W. Stonecypher, S. Li, A. Bansal, A. Agarwal, E. Alon, V. Stojanović and J. Zerbe "Modeling, simulation, and design of a multi-mode 2-10 Gb/sec fully adaptive serial link system," Custom Integrated Circuits Conference, Proceedings of the IEEE 2005, pp. 704-711, 2005.

H. Hatamkhani, F. Lambrecht, V. Stojanović and C.-K. Yang “Power-Centric Design of High-Speed I/Os,” *Design Automation Conference*, 2006.

S. D. Vamvakos, V. Stojanović, J. L. Zerbe, C. W. Werner, D. Draper and B. Nikolić “PLL On-Chip Jitter Measurement: Analysis and Design,” Digest of Technical Papers, IEEE Int’l. Symposium on VLSI Circuits, June 2006.

A. Amirkhany, A. Abbasfar, V. Stojanović and M. A. Horowitz “Analog Multi-Tone Signaling for High-Speed Backplane Electrical Links,” accepted for publication at *IEEE Global Telecommunications Conference, 2006. GLOBECOM '06.*