

CMOS Integration of Capacitive, Optical, and Electrical Interconnects

Jon Lexau¹, Xuezhe Zheng², Jon Bergey³, Ashok V. Krishnamoorthy²,
Ron Ho¹, Robert Drost¹, Jack Cunningham²

¹Sun Microsystems, 16 Network Circle, Menlo Park, CA; ²Sun Microsystems, 9515 Towne Centre Dr, San Diego, CA

³Forza Silicon, 776 E. Green St. Suite 200, Pasadena, CA
jon.lexau@sun.com

Abstract

We present a 90nm test chip integrating proximity communication, optics using external lasers and photodiodes, and CML electronics on a single CMOS chip which can route data at multi-Gb/s rates through any combination of its three interconnect interfaces. A robust and flexible unlocked datapath allows independent timing and margin characterization of each path.

Motivation

Proximity communication provides a high-bandwidth, high-density channel between two chips. Test results show far lower per-pin power, latency, and area costs when compared to traditional solder balls [2][3]. However, because this technique relies on capacitive coupling between two chips placed face-to-face, it only works if the chips are in very close proximity. In contrast, optical networks provide a proven communication technology for larger distances, ranging from backplanes to wide-area-networks [4].

Systems that integrate many multi-chip packages together can benefit from both communication technologies: they can use proximity communication within each package for very high bandwidth and low power data transfers, and they can communicate between packages using optical networks. In addition, such a system would likely require high-speed electrical channels for system I/O, testing, and configuration.

To explore this diverse three-way marriage of optics, proximity communication, and high-speed electrical I/O, we built a CMOS test chip in the ST Microelectronics 90nm process integrating all three. In this paper, we discuss interface requirements, floorplanning, and test results from this chip.

Chip Overview and Interface Details

Proximity communication requires face-to-face chip overlap, optics requires external lasers and photodetectors which can be either wirebonded or flip-chip bonded to our chip, and the high-speed electrical interface also requires wirebonding or flip-chip connections. Figure 1 is a sketch of these varied interfaces, showing how the pieces might fit together.

Figure 2 shows the three chip interfaces and their possible interconnections. For full testability, any port can be routed to any other, as well as back to itself, through a series of multiplexers and demultiplexers controlled via scan chains. Each port interface contains four independent channels. Using four channels enables testing for possible cross-talk between channels, and matches commonly available laser (VCSEL) and




Figure 1. Interface Physical Constraints

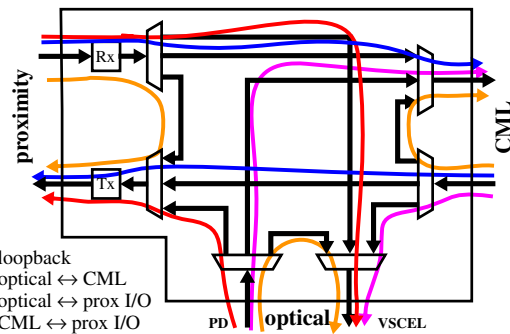


Figure 2. High-Level Datapath

photodetector (PD) arrays

To simplify the test chip, we created a flow-through datapath for all three interfaces. Therefore, none of the ports are clocked, and data simply flows in and out without synchronization. Cross-channel skew is minimal with only these four channels, but future versions with wider datapaths will be explicitly clocked.

A. Electrical Interface

The electrical interface uses differential current-mode logic (CML) for high speed and for compatibility with existing optical circuits. On-chip poly resistors provide pin termination. The CML interface allowed for easier system testing, because its input can be directly AC-coupled to a pattern generator, without needing to set an appropriate common-mode voltage. The CML output buffer can drive a 50Ω transmission line on the test board, and is tunable using three digital waveshaping signals. Two selections modify the output voltage swing and edge rate, and a squelch control completely disables all of the output channels, if desired.

B. Optical Interface

This integration employs commercially-available discrete lasers and photodetectors flip-chip bonded or wirebonded to our chip. Future prototypes will integrate waveguides and photodetectors using SOI-based Silicon photonics to reduce cost, complexity, and power.

The optical transmit channel takes an electrical signal and converts it to a current to drive the external VCSEL. The transmitter consists of a level-shift block and a modulator core. The level shifter performs several waveshaping functions: edge-rate control, output current peaking, and eye crossing control, all set using off-chip adjustable current sources. The modulator core provides a constant laser bias current and modulates a signal current on top of it, again using external current sources. This flexibility allows us to experiment with a wide range of VCSELs and bonding options and also to fine-tune the link performance.

The optical receive channel consists of a transimpedance amplifier (TIA) followed by a limiting amplifier. The TIA takes the very small, single-ended PD current and converts it to a differential voltage. The 5-stage limiting amplifier provides full-swing CMOS levels at its output, and feeds its output back to its input through an error-correcting amplifier in order to zero out voltage offsets.

C. Proximity Interface

Conceptually, the proximity transmit and receive plates, made of rectangles in top-layer metal, form parallel plate capacitors that can send data wirelessly across the gap between face-to-face chips as shown in Figure 1. We employ a two-part solution to manage the critical problem of chip misalignment [1].

First, each transmit and receive pair includes a set of “where” blocks which can precisely detect the relative positions of two chips using electronic means. The where block resembles a large array of small transmit and receive plates - for this chip, 20x20 arrays of 25μm plates. By putting certain patterns on the transmit plates, sweeping a corresponding pattern on the receive plates, and then measuring the strength of the detected signal, we can determine any misalignment in each axis. Second, given some fixed chip misalignment, we can electrically steer the transmit data by a small amount in any direction to correct for that misalignment. We accomplish this by dividing each transmit pad into a 4x4 array of 16 micro-pads and using a scan chain to shift the data up to half a transmit plate in any direction. Two additional microplates are required around the entire array perimeter, but this extra cost is easily amortized over a larger array.

For this chip, the transmit and receive plates are arrayed on 125μm centers. The receive plates are slightly smaller than the transmit micro-pad arrays so that a receiver will be less likely to “see” noise from an adjacent transmit bit under a small chip misalignment. The layout of a typical proximity experiment is shown in Figure 3. The receive array is on the left side and transmit on the right. The small squares on the transmit side indicate the microplates. The four differential channels are arranged in a 2x4 array of plates with a diagonal pattern to reduce first-order cross-coupling between bits. In use, one chip will be flipped left-to-right and then placed over the other, as shown in Figure 1, causing the proper bit positions to align.




Figure 3. Typical Proximity Experiment Layout

We included five separate proximity experiments on this chip in order to test different receiver and transmitter circuit designs, as well as to provide single-chip testing capability. In addition to three chip-to-chip proximity communication experiments, there are two separate on-chip experiments where transmitters and receivers use on-chip parallel plate capacitors to emulate face-to-face chips. Some bits are even wired directly between transmitter and receiver. These experiments allow us to do initial test and debug of the circuits before placing and aligning two chips. In fact, these tests can even be performed with PGA packaged parts.

Chip Floorplan

The 4mm x 6mm chip floorplan is shown in Figure 4. Conventional chip I/O is around the right half of the die. All wirebonded connections must be kept far from the proximity communication pads, because two chips arranged face-to-face will block any nearby wirebond sites. The CML and optical circuits sit on the right half of the die. The “data steering” block contains most of the mux/demux circuits shown in Figure 2. The proximity communication channels are on the left half of the die. The three off-chip experiments (B,C,D) stack in a column around a common line of symmetry. The on-chip proximity experiments (A,E) are located closer to the center of the chip because they need no chip overlap to function. The four large squares at the top and bottom left are areas for experimental MEMS processing which will enable precise mechanical chip

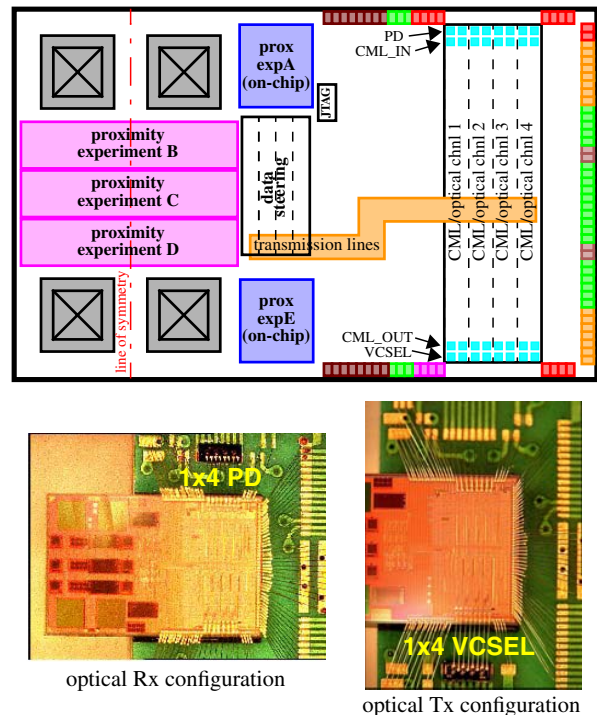


Figure 4. Chip Floorplan and Die Photos




Figure 5. Optical Receiver Sensitivity and Eye Diagram

alignment. Figure 4 also shows photographs of our chip mounted on its test board in optical transmit and receive configurations. Note how the proximity communication circuits lie off the board, like a “diving board” outrigger. This eases chip-to-chip overlap in a package or test setup.

Testing

Using only one test board, we can test the CML-to-CML and CML-to-optical paths. The CML loopback connection works cleanly at data rates exceeding 5Gb/s. For the optical output path, a lens focuses the VCSEL output onto the end of a fiber which connects to a stand-alone high-speed (10Gb/s) photodetector, whose output is then measured electrically. Alternately, the optical signal can be connected directly to a plug-in module on the BERT. The optical transmit path works at 5Gb/s with an extinction ratio of 6.7dB. We test the optical receiver by configuring one test board for CML-to-optical to generate data for a second board which does optical-to-CML. Figure 5 shows the receiver sensitivity and eye diagram for this complete link at 5Gb/s. No apparent noise floor is observed to a BER of 10^{-14} . The receiver sensitivity is -11.5dB at a BER of 10^{-12} .

Proximity communication required the most complicated test setup, where two boards must be positioned face to face. We fix one board, and mount the other on a precision 6-axis positioner, which is able to move the board in steps as small as $0.1\mu\text{m}$. The fixed board is mounted right side up, with a VCSEL array bonded next to our die. An optical fiber routes the data to a third test board which has a PD input. The CML output of this final board is connected to a data communications analyzer for performance measurement. A portion of this test setup is shown in Figure 6.

Figure 7 shows the timing margin (as a percent of the bit period) for a range of data rates and test configurations. From top to bottom, the curves are: (1) our custom pattern generator local loopback, (2) CML electrical loopback, (3) the three-board setup described above, and (4) a two-board setup with CML in, proximity communication to a second board, and then




Figure 6. Three Board Proximity/Optical/CML Test Setup




Figure 7. Proximity Timing Margin vs Data Rate




Figure 8. Proximity Timing Margin vs Gap at 1.85Gb/s

back to CML. Note how the optical path cleans up some of the degradation introduced by accumulated jitter in the non-clocked proximity communication link.

For a proximity communication channel, we are also concerned about the link behavior relative to the gap between the chips. Using the six-axis positioner, we can step this spacing and measure the eye opening. Figure 8 shows the results of such a test. The data rate is 1.85Gb/s for all points, and the two curves represent cases (3) and (4) from the previous test. Again, adding the optical path improves the timing margin.

Conclusions

Scaling of computer systems performance will continue to drive integration of diverse I/O technologies. In this experiment we combined capacitively-coupled I/O, VCSEL-based optical I/O, and traditional electrical I/O on a single CMOS platform, and verified full interoperability between all three interfaces at speeds exceeding 2.5Gbps. This chip enables future prototypes to more fully characterize metrics such as power or latency, to further explore die packaging, and to push more aggressive interconnects such as silicon-based photonics.

Acknowledgements

The authors gratefully acknowledge the assistance of D. Beckman, J. Gainsley, A. Chow, T. Ono, F. Liu, D. van Blerkom, R. Bosnyak, A. Chakraborty, and P. Bell-Bosnyak in the design and testing of our prototype chip. This work was supported in part by DARPA contract NBCH30390002.

References

- [1] R. Drost *et al.*, “Electronic alignment for proximity communication,” *IEEE ISSCC*, 2004.
- [2] R. Drost, *et al.*, “Proximity communication,” *IEEE JSSC*, 9/2004.
- [3] R. Hopkins, *et al.*, “Circuit techniques to enable 430 Gb/s/mm/mm proximity communication,” *IEEE ISSCC*, 2007.
- [4] A. V. Krishnamoorthy, “Photonics-to-electronics integration for optical interconnects in the early 21st century,” *Optoelectronics Letters*, 5/2006.