A Serial Link Transceiver Based on 8 GSa/s A/D and D/A Converters in 0.25µm CMOS

William Ellersick[1,3], **Chih-Kong Ken Yang**[2]
Vladimir Stojanovic[1], **Siamak Modjtahedi**[2],
Mark A. Horowitz[1]

[1]Stanford University
[2]University of California at Los Angeles
[3]Analog Devices, Inc.

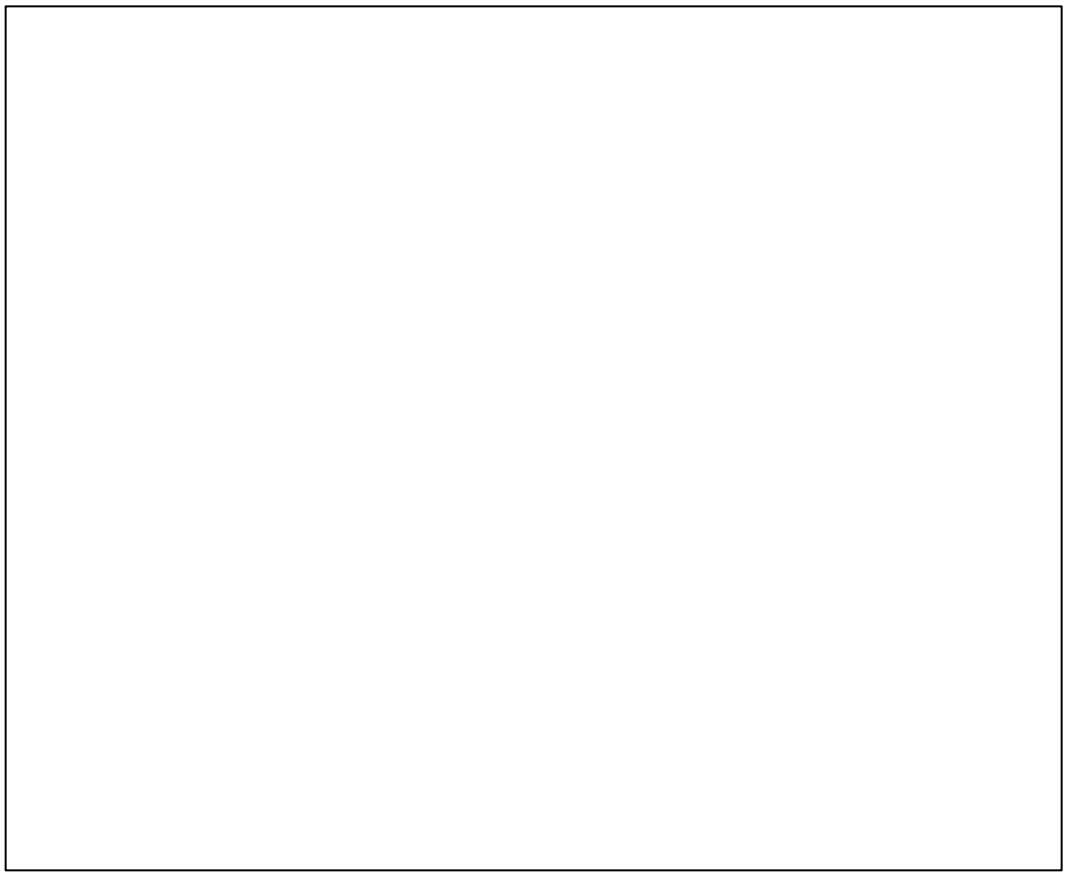

Goal: More Bits on Long Links

- **Wires expensive: use fewer, more complex links**
- **Links between boards and systems**
  - **Communications or computing: tens of Gbit/sec**
  - **Up to 10m cables or 1m backplane traces**
- **CMOS technology for systems-on-a-chip**
  - **Smaller, cheaper, cooler products**
- **More bits/symbol or faster symbol rate**
  - **Improve and equalize parasitic filters**
  - **Reduce other interference sources**

This work explores the design of CMOS transceivers for long links, where wires are expensive, so it makes sense to use fewer, more complex links. These links run between communications or computing systems and their boards, carrying tens of Gbit/sec on cables up to 10m long, or over backplane traces up to 1m in length.

CMOS technology is of the most interest, so that the transceivers for these long links can be integrated with processors and memory, enabling great reductions in system size, cost and power.

There are only two ways to carry more bits on a link: either pack more bits into each symbol, or run at a faster symbol rate. Either way, parasitic filtering and other interference become larger problems at higher bit rates, and must be reduced to increase link data rates.

Equalization of Parasitic Filters

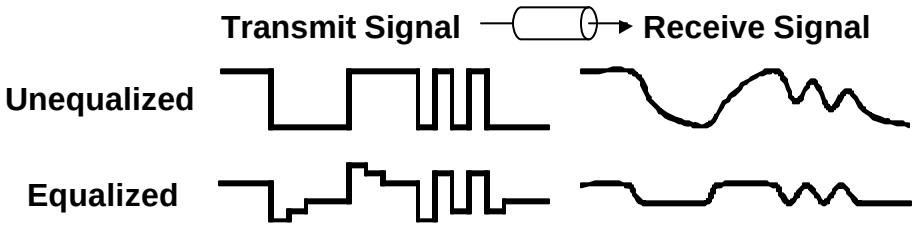

- **Use high bandwidth circuits and RF techniques**
- **Linear equalizer using DAC in transmitter**
  - **Wire and circuit losses are linear, time invariant**
  - **Attenuates low frequencies to match cable loss**
  - **Signal amplitudes decrease with frequency**
- **ADC in receiver for higher performance**
  - **More bits/symbol: multi-level modulation**
  - **Decision Feedback Equalization (DFE)**

Parasitic filters smear signal edges and prevent them from reaching full swing in a single symbol time. This work explores high bandwidth circuits and RF techniques to reduce the parasitic filtering, and shows how to compensate for the remaining circuit and wire with an adjustable high pass filter called a linear equalizer. The transmit equalizer emphasizes the high frequency edges of the signal and then decays to a smaller output level when the signal doesn't change. Because the transmitter is limited in voltage swing, low frequency signal components are attenuated to match the high frequency losses.

An A/D converter in the receiver enables digital communications techniques such as multi-level modulation or Decision Feedback Equalization that can result in larger received signals. But even with these techniques, signal amplitudes decrease at higher bit rates.

Other Interference Reduction

- **Correct linear interference with DAC**
  - **Reflections from connectors and packaging**
  - **Signal and clock coupling**
  - **Correlated supply bounce**
- **Cancel static errors from mismatches**
  - **Timing errors**
  - **DAC nonlinearity**
  - **ADC offset voltage errors**
- **Random time-varying noise**
  - **Thermal noise**
  - **Random phase noise (high bandwidth timing recovery)**

At the same time, other interference sources get worse at higher frequencies. Reflections, coupling and supply bounce are corrected with the same DAC used to implement the transmit equalizer. The signal distortion caused by timing errors increases as edge rates increase. Even DAC nonlinearities and ADC offset voltages can be larger for high bandwidth transceivers, because small transistors are used for their low parasitic capacitances. Still, these static interference sources can be corrected.

What remains are the effects of random, time-varying noise processes that previous results indicate should not limit performance.

Calibrated Link Based on DAC, ADC

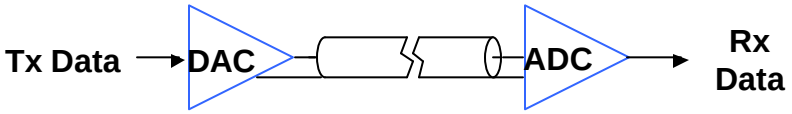

8 GSample/sec: 1 FO4 gate delay in 0.25㎜m CMOS
4 GHz Bandwidth

- **2 bits for 4-level signaling**
- **4 extra bits for equalization, interference correction**
- **2 more bits to explore limits**
- **2 bits for 4-level signaling**
- **2 extra bits for DFE**

Therefore, this work investigates a calibrated transceiver that can measure and correct for parasitic filters and interference. To increase data capacity, symbol rates are as fast as binary transceivers, which have been reported with bit times of 1 FO4 gate delay. This corresponds to 8 Gsamples/sec in 0.25 um CMOS. Nyquist bandwidth of 4 GHz is desired so the circuits don't limit the link performance.

2 bits of resolution are needed to support 4-level signalling. The transmitter has 4 more bits for equalization and interference correction and an extra 2 bits to explore the limits of resolution. The receiver has 2 more bits for signal processing such as Decision Feedback Equalization.

Now the issue with this architecture is really the performance of the ADC and DAC, not the communication techniques, whose performance is well understood.

Organization

- **Introduction**
- ➔**Transceiver design**
  - **DAC circuits**
  - **ADC enhancements**
  - **Inductors to distribute parasitic capacitances**
  - **Clock generation**
- **Experimental results**
- **Conclusions**

Thus, this paper explores the limits of DAC and ADC performance with sample periods of a gate delay. I will present the design of the DAC, and then the enhancements to the ADC over the first implementation, focusing on how to achieve high bandwidth signal paths through circuit design and the use of inductors to distribute parasitic capacitances.

Then, I'll discuss how precision clocks are generated for the data converters. I will present results from the chips we had fabricated, and conclude the talk.

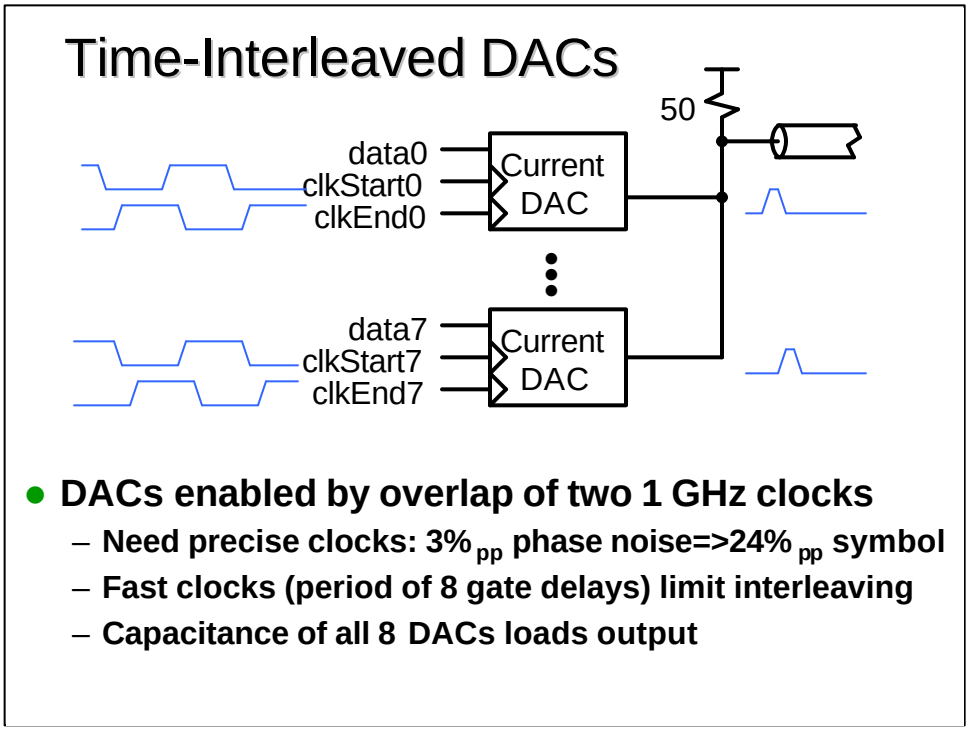

In the transmitter, 8 GSample/sec is achieved by time interleaving eight D/A converters. Each of the interleaved DACs generates a narrow pulse for one symbol time. The output pulses are summed onto the 25 Ohm impedance formed by the 50 ohm onchip termination in parallel with the 50 ohm transmission line.

Because it's difficult to generate narrow clock pulses, each interleaved DAC is enabled by the overlap of two 1 GHz clocks. The clock edges need to be controlled precisely to produce pulses of desired width and position, since with 8-way time interleaving, 3% of a clock cycle of timing noise corresponds to 24% of a symbol time.

The 1 GHz clocks are about as fast as CMOS clocks can be generated or used. This limits the interleaving required to achieve 8 GSa/s, which is important, because the capacitance of all 8 DACs loads the output.

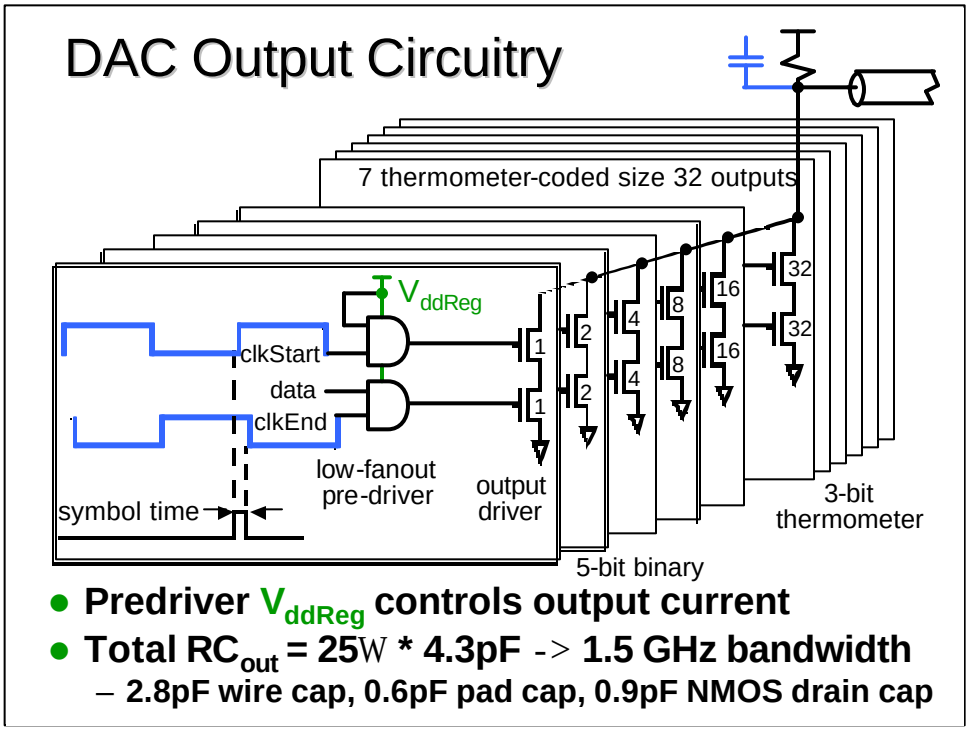

Each of the DACs consists of an array of small current sources, with the least significant bits implemented with 5 binary weighted current sources, and the upper 3 bits with 7 identical current sources. The outputs of the current sources are summed onto the 25 Ohm output impedance.

Each of the current sources consists of 2 stacked NMOS transistors with a grounded source. The output level is controlled by regulating the supply of the predrivers. This maximizes the gate to source voltage at the output, allowing small transistors to be used.

Still, because of the 8 way interleaving, the RC pole at the DAC output limits the bandwidth of the transmitter to 1.5 GHz.

Interleaved Flash ADC Comparator

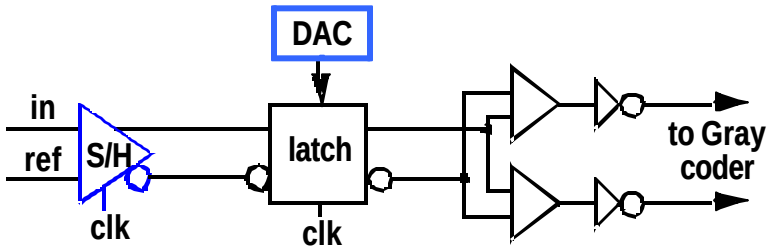

- **Offsets corrected with DAC in 2nd stage latch**
  - **Simple, high bandwidth sample/hold amplifier**
  - **No switches or capacitors: breaks bandwidth vs. offset tradeoff**
- **Enhancements over GAD chip (VLSI-99)**
  - **Improved DAC linearity through better matching of clock coupling**
  - **Clocking error corrected**
  - **S/H gain reduced 30% to 1.3: 7 GHz bandwidth (0.4LSB$_{pp}$ noise)**
- **Total RC$_{in}$ = 25W * 1.9pF -> 3.3 GHz bandwidth**
  - **0.6pF wire cap, 0.6pF pad cap, 0.7pF NMOS gate cap**

The ADC is also time interleaved, and uses a flash architecture. Thus, its performance is largely determined by the design of the comparator, shown here. Offset errors are corrected in the comparator with a DAC that drives a correction signal into the 2nd stage latch. Note that no switches or capacitors are used in the data path. This approach breaks the tradeoff between bandwidth and offset voltage, allowing small transistors to be used for high bandwidth, and correcting for the resultant large transistor mismatch errors. It also allows the use of a simple sample/hold amplifier, with low gain for high bandwidth.

There are several enhancements over the comparator design that was presented at VLSI in '99. The offset correction DAC linearity is improved through better matching of clock coupling. A clocking error was corrected and the sample/hold gain reduced to increase the sampling bandwidth to 7 GHz. The reduced gain also had the effect of increasing the input referred thermal noise to a 6 sigma value of 0.4 LSBpp.

Again, despite the near minimum size devices used, the input capacitance of 1.9pF results in a pole at 3.3 GHz, limiting ADC bandwidth to less than the 4 GHz goal.

ADC Inductor Network

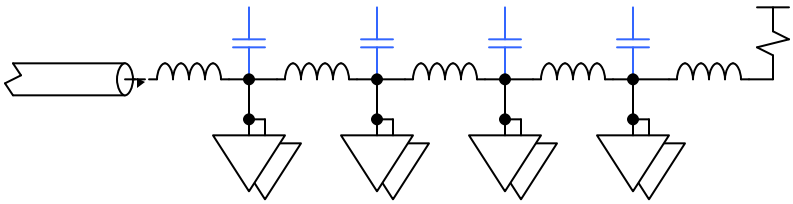

- **Goal: reduce effect of large parasitic capacitance**
  - **Old good idea: distributed amplifiers in oscilloscopes**
- **If sqrt(L/C)=50Ω: lumped transmission line**
  - **Each ADC sees 50Ω to the right, 50Ω to the left**
- **Inputs of pairs of ADCs to separate pads**
  - **Bond wire inductors can optionally be inserted**
- **Bandwidth traded for delay (100ps total on line)**

So we use an old, good idea, used for years to build distributed amplifiers in oscilloscopes. By dividing up the large capacitances, and inserting the right value of inductor,a lumped 50 ohm transmission line is constructed.

Pairs of interleaved ADC inputs are wired to separate pads. This allows bond wire inductors to optionally be inserted between the input capacitances to evaluate performance with and without the inductors.

This distributed ADC technique trades bandwidth for delay - about 100 ps down the entire LC line. A block diagram of the fulll transceiver shows how this delay is managed.

Transceiver Inductors and Clocking

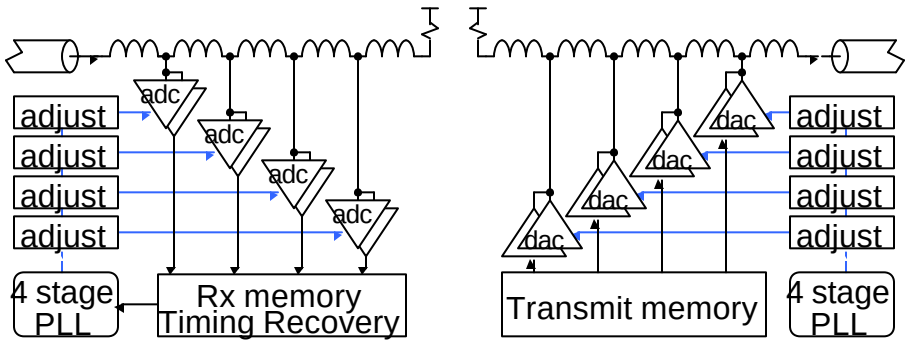

- **Phase adjusters correct LC delay, static errors**
  - **Adjuster: clock mux, 1/16th-symbol clock interpolator**
  - **8 ADC phase adjusters + 1 for timing recovery**
  - **16 DAC phase adjusters (2 clocks for each DAC)**

The receiver is shown on the left. Because each time-interleaved A/D converter has its own clock (shown in blue), the LC delay on the lumped transmission line is compensated for by adjusting the phase of these clocks. The same ability to adjust clock phases is needed to cancel static phase errors, which have dominated timing errors in previous interleaved transceiver chips.

To produce adjustable clocks, a four-stage differential PLL drives 8 clocks with 45 degree spacing into the phase adjusters. Each phase adjusters consists of a clock mux and a digitally controlled clock interpolator with a resolution of 1/16th of a symbol time. One phase adjuster is used for each of the eight time-interleaved ADCs, with an additional phase adjuster in the feedback path of the receive PLL for timing recovery.

A similar inductor network is used to distribute the transmitter output capacitance, with two phase adjusters for each interleaved DAC to correct pulse width and pulse position.

This represents the bulk of the circuitry on the test chip, so let's take a look at the layout.

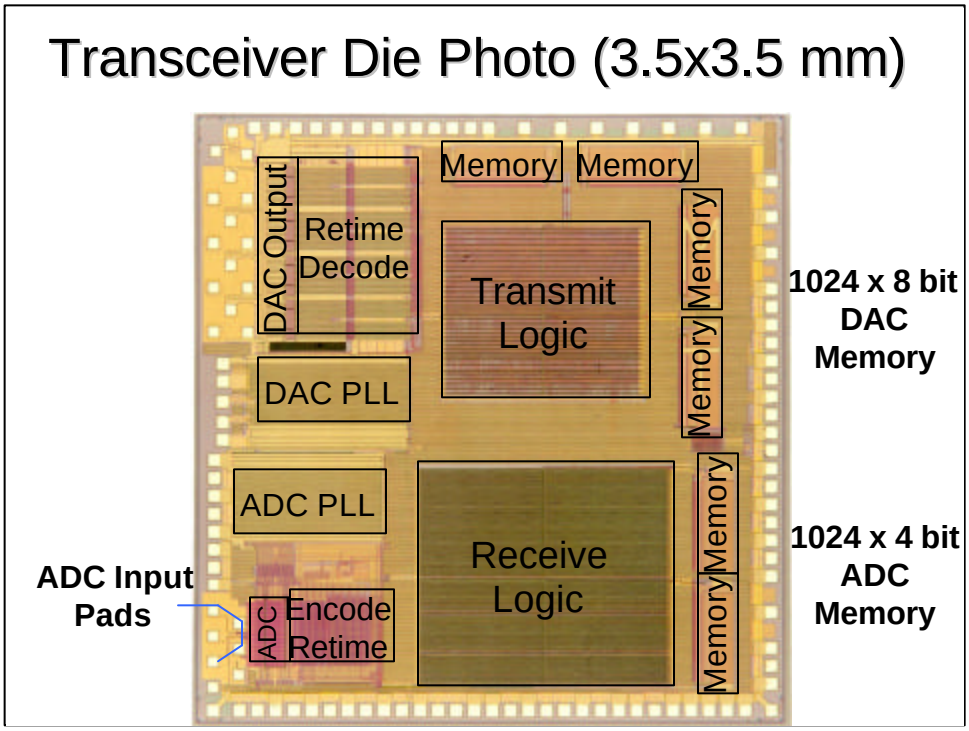

The PLLs can clearly be seen in this die photo, swollen by the large number of phase adjusters and the big devices used to reduce clock buffering. The analog ADC circuitry is quite small, even though it includes 120 comparators and the offset correction circuits and memory. Its output is converted to 4 bit Gray code, retimed to a lower rate clock, and fed into a synthesized logic block, which either stores or compares the data against the contents of a 1024 symbol memory. The synthesized logic is large because of the 128 and 256 bit wide data paths.

The transmit data path is similar but reversed. Data from a 1024 symbol memory is retimed to the high speed DAC clocks and decoded. Again, the analog output circuitry is fairly small, even though most of the design effort and experimental work focused on the analog sections. The die is packaged in a ceramic leadless chip carrier with controlled 50 ohm signal lines.

Measured ADC results are presented first, comparing performance with and without inductors. Note the four input pads in the lower left, each of which is connected to a pair of time-interleaved ADCs. These can either be bonded with short, flat bond wires with no appreciable inductance as shown, or with small inductor coils to distribute the input capacitance. Let's zoom in on these input pads...

ADC Bond Wire Inductor Photo

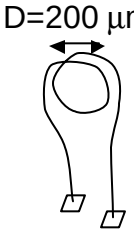

- **2-turn inductors ~1.1 nH**
  - **Not onchip inductors: large R**
- **Goal: inductors in flip-chips**
- **HSPICE: +/-30% variation doesn't affect results**
  - **nH/mm depends on log(D)**
  - **30 gauge wire form: 200 µm**
- **Verify LC match with TDR**

... on a chip bonded with inductors. Each of the inductors consists of a 2-turn coil of bond wire. The input comes in to the first ADC input pad, through a 2-turn bond wire inductor to the second input pad, through more inductors to the third and fourth input pads, and through a final inductor to the onchip termination resistor. Now this is a beautiful piece of bonding work by Pauline Prather, who works in our lab, but it is obviously not mass manufacturable. Bond wire inductors were used instead of onchip spiral inductors to keep series resistance low, with the goal to gather data to help design inductors into a flip chip package.

Precise 1.1nH inductors obviously cannot be bonded, but simulation shows that a 30% variation in inductance is acceptable. Also, because the inductance per unit length depends on the log of physical dimensions, moderate changes in shape and size only result in small changes in the inductance value. To control the diameter of the coils, we looked for the smallest cylinder we could find, which turned out to be a piece of 30 gauge wire, and used it as a form to wind inductors with a 200 um diameter. Coupling between the inductors changes the frequency response, but can be equalized. The LC matching was verified with a Time Domain Reflectometer, showing that the ADC input capacitance has been distributed, so that the frequency response of the ADC should be limited only by the sample/hold amplifier bandwidth of 7 GHz.

ADC Sinewave Loss vs. Frequency

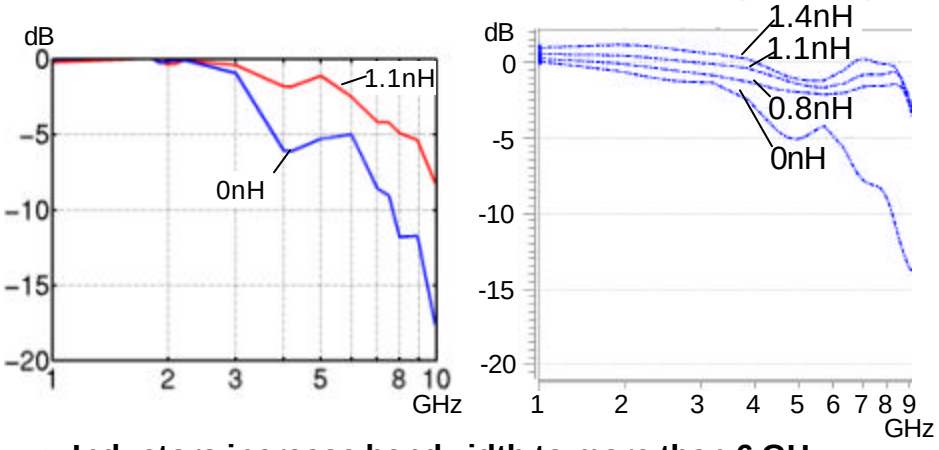

- **Inductors increase bandwidth to more than 6 GHz**
- **Measured TDR matches simulated TDR with L=1.1nH**
- **Bumps are due to connector and package reflections**

And indeed, the ADC has a measured bandwidth of more than 6 GHz with inductors, compared to 3.5 GHz without. The measured data on the left is generated by driving asynchronous sine waves into the ADC and fitting ideal sine waves to the data captured in the onchip memory. The frequency response matches the simulation of a detailed model of the ADC input path, shown on the right, and is flatter with inductors because reflections are reduced. Three values of inductance are simulated to show sensitivity to inductance value.

Now these results looks promising, and show that the ADC has high input bandwidth, but we need to consider a broader measure of ADC performance that includes noise and distortion.

ADC Performance

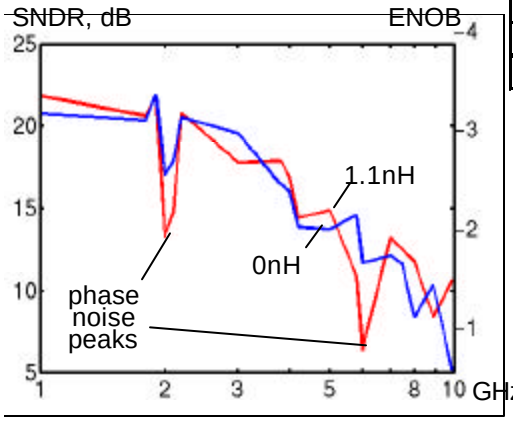

Power	1.1 Watts
Static Error (INL)	0.6LSB$_{pp}$(3 raw)
Static Phase Error	10 ps$_{pp}$ (47 raw)
Random Noise	0.5 LSB$_{pp}$
Timing Variation	15 ps/V

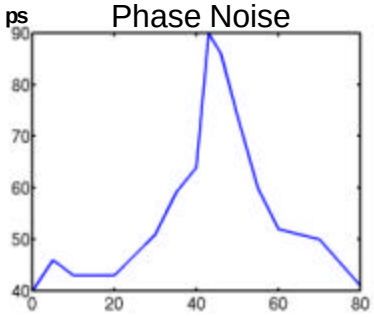

- **SNDR limited by phase noise**
  - **Not improved by inductors: distortion proportional to signal**
  - **Technology file error reduced filter cap: underdamped PLL**
  - **Input couples into PLL, is downsampled: excites noise peaks**

The Signal to Noise and Distortion ratio of the ADC is plotted on the left, showing a resolution bandwidth of only about 3 GHz. This is due to phase noise, which causes distortion that increases with frequency and is proportional to the sampled signal. Inductors do not improve the SNDR, because the larger received signal with inductors is accompanied by an increase in the distortion caused by phase noise.

The phase noise is twice that measured on previous chips, due to an error in the layout technology filethat caused the PLL filter capacitance to be less than half the desired value. Simulations show this results in larger phase noise and an underdamped PLL, and this is confirmed by experimental measurements showing a peak of 90 ps of jitter in response to noise at 42 MHz, even though the PLL was designed for a loop bandwidth of 12 MHz. Also, since coupling from the single-ended ADC input to the PLL reference clock is downsampled by the PLL, phase noise peaks are seen in the SNDR plot near clock harmonics, although there are not enough data points to show all the peaks.

The ADC consumes a little over 1 Watt, dominated by the oversized PLL and phase adjusters. Offset correction using the DACs in each comparator reduces errors from 3 to 0.6 LSBpp. The phase adjusters reduce static timing errors from 47 to 10 ps,pp. Half an LSB of p-p random noise is measured, matching estimates of thermal noise.

DAC Performance

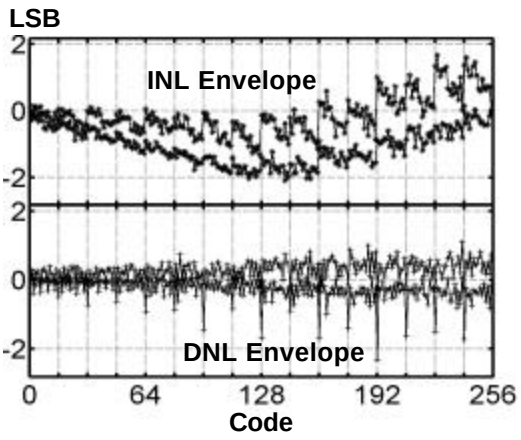

Measurements	LSB:3mV
Power	1.9 Watts
Static Phase Err.	10 ps_{pp}(47 raw)
Clock Coupling	6LSB_{pp}(63 raw)
Voltage Error	1LSB_{pp}(3.7raw)
Bandwidth	3 GHz(1.5 L=0)
Phase noise	45 ps_{pp}

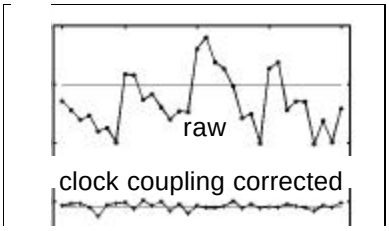

- **Wrong inductors bonded, half ideal value: low BW**
- **Phase noise affects asynchronous SNDR**
  - **Less effect on synchronous transceiver**

The DAC dissipates about twice as much power as the ADC, because it has twice as many clocks. Coupling from all the interleaved clocks is large, but is reduced from 67 to 6 LSBpp by transmitting a correction signal as shown on the right, albeit at the cost of a reduction in usable output swing. Remaining voltage errors are shown in the envelope plots on the left that show the worst case INL and DNL out of the eight time-interleaved DACs, at each output code. The curve in the INL envelope is due to the variation in DAC output impedance with output voltage. The large DNL errors every 16th code fortunately are negative, and thus cause overlapping codes instead of large gaps. Thus, the DAC has about 5 effective bits at low frequencies after calibration. Resolution at higher frequencies is limited by the large phase noise and the 3 GHz output bandwidth, which is low because inductors of half the ideal value were mistakenly bonded.

However, equalizing for the low bandwidth is similar to equalizing a long cable, and phase noise affects synchronous operation less than it affects the asynchronous resolution and SNDR, so let's characterize the transceiver in a link.

Transceiver Equalization, DAC to ADC

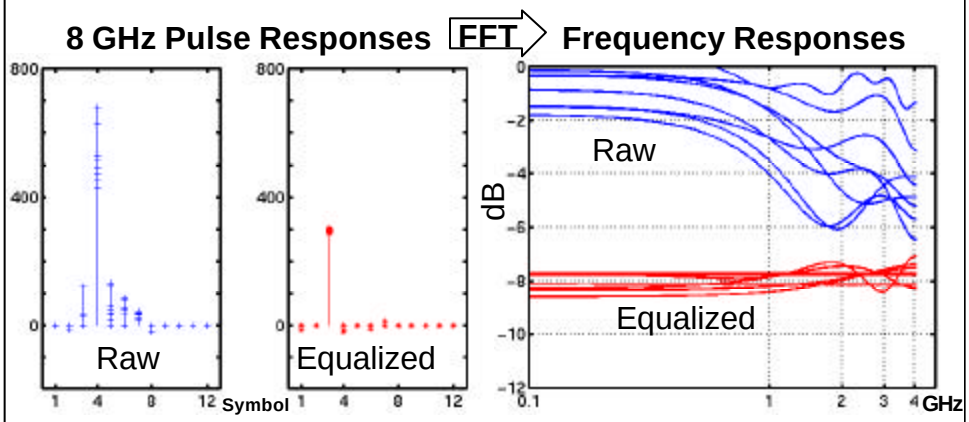

- **8 unique pulse responses: need 8 lookup tables**
- **Short cable used: 2 GHz transceiver bandwidth**
- **Measured with MSB comparators by adding DC**

The plots on the left show the pulse responses of the transceiver before and after equalization. This data was generated by transmitting a pulse from each time-interleaved DAC, and plotting the voltage sampled by the receiver at each symbol time. Parasitic filtering causes the pulses to be spread over several symbol times. This inter-symbol interference varies with the differing positions on the LC lines, and the pulses heights vary because of residual pulse width variations between the time-interleaved DACs. Thus, each transmitter needs a unique equalization lookup table. However, this does not increase complexity, since parallel lookup tables are already needed to run at 8 GSa/s, although our test chip implements the lookup tables in software on a PC. After equalization and gain correction, the pulse heights are reduced nearly the same 300 mV height, and the interference in adjacent symbol times is significantly smaller. Only a short cable was used, as the low transceiver bandwidth of 2 GHz is roughly equivalent to a long cable, and is a good challenge to 8 GSymbol data transmission and equalization. The FFT of the pulse responses show the attenuation of low frequency signal components to match the high frequency parasitic loss.

To avoid quantization errors in equalization, the pulse responses are measured by adding a variable DC voltage to the signal to find the switching point of the most significant bit comparator in each time-interleaved A/D converter.

8 GSymbol/s Transceiver Eyes

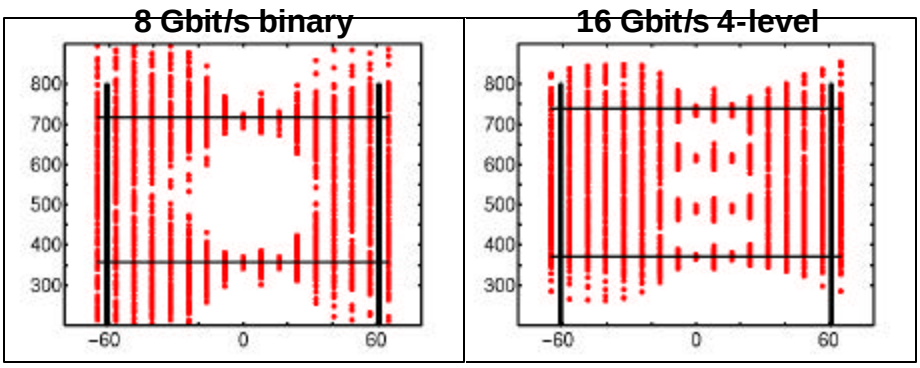

- **Eye openings less than 750 mV swing by 125 ps**
  - **Binary: 300 mV_{pp}, 4-level 100 mV_{pp}; 45 ps_{pp} width**
  - **Random transceiver phase noise: 57 ps_{pp}**
  - **2 GHz bandwidth, large voltage errors corrected**
- **Binary operation also verified at 10^{-10} BER**

A similar technique is used to measure the voltage and phase margins of the link. The schmoo plots shown here are generated by varying the DC voltage and the phase of the receiver clocks, and plotting a dot where bit errors are recorded. The received signal swing is reduced by 8 dB by the 2 GHz bandwidth and large clock coupling correction, and performance is limited by the random transceiver phase noise of 57 ps,pp. These schmoo plots represent a BER of about 10^{-3}, and show that 4-level signalling is possible. Binary operation was verified at a bit error rate of 10^-10, with a smaller eye.

Thus, despite large phase and voltage errors and significant clock coupling and bandwidth limitations, an accurate 8 GSymbol/s transceiver can be built in 0.25 um CMOS.

Conclusion

- **High speed CMOS circuit techniques**
  - **Time-interleave DACs and ADCs for high sample rate**
  - **Inductors to trade bandwidth for delay**
  - **Small transistors for high bandwidth, low power**
- **Calibration to improve accuracy, maintain speed**
  - **Extra DAC bits to correct interference, nonlinearity**
  - **ADC offset cancellation with DAC inside comparator**
  - **Clock interpolators to correct static phase errors**
  - **Correct layout, circuit asymmetries to lower design risk**
- **Data converters at binary transceiver speeds**
  - **Accurate timing is crucial**
  - **Long links: 4-level signaling, DFE**

In conclusion, we've demonstrated techniques for designing high performance transceivers in cheap CMOS technology. High sample rate is achieved by time-interleaving D/A and A/D converters. High bandwidth is achieved by using inductors to trade bandwidth for delay, and by designing with small transistors for their low parasitics and low power consumption.

Calibration techniques are used to improve accuracy while maintaining speed. The DAC in the transmitter corrects for parasitic filtering, interference and simple nonlinearities, and ADC offsets are corrected with a DAC inside each comparator. Timing accuracy is improved with clock interpolators that correct static phase errors. These techniques also correct for layout and circuit asymmetries and thus lower design risk. This was critical for us, because we learned more than we ever expected to on this project. But we were able to use calibration techniques to correct most of the mistakes that we learned from.

This work demonstrates that data converters are possible at binary transceiver speeds, and that accurate timing is crucial to their operation. As circuit speeds increase and transistors become cheaper, data converters will enable the use of digital communications techniques to increase performance on long links.

I would like to acknowledge the assistance of Ken Chang, Azita Emami, Dean Liu and Jackie Wong, and some amazing bonding by Pauline Prather. Finally, this work would not have been possible without the efforts of the students I followed, particularly that of Stefanos Sidiropoulos, who will follow me today.