

On-chip oscilloscopes for noninvasive time-domain measurement of waveforms

K. L. Shepard and Y. Zheng
Columbia Integrated Systems Lab
Department of Electrical Engineering, Columbia University
New York, NY 10027
{ shepard,yu } @cisl.columbia.edu

Abstract

High-speed digital design is becoming increasingly analog. In particular, interconnect response at high frequencies can be non-monotonic with "porch steps" and ringing. Crosstalk (both capacitive and inductive) can result in glitches on wires that can produce functional failures in receiving circuits. Most of these important effects are not addressed with traditional ATPG and BIST techniques, which are limited to the binary abstraction. In this work, we explore the feasibility of integrating primitive sampling oscilloscopes on-chip to provide waveforms on selective critical nets for test and diagnosis. The oscilloscopes rely on subsampling techniques to achieve sub-10 psec timing accuracy. High speed samplers are combined with DLLs and a simple 8-bit ADC to convert the waveforms into digital data that can be incorporated as part of the chip scan chain. We will describe the design and measurement of a chip we have fabricated to incorporate these oscilloscopes with a high-frequency interconnect structure in a TSMC 0.25 μm process.

1 Introduction

There is strong recent interest in the ability to noninvasively measure waveforms in the time-domain in integrated circuits. In digital design, this interest stems from the inability of traditional digital test methodologies (e. g., ATPG and BIST) to address the more analog issues of high-speed design such as crosstalk noise and complex nonmonotonic waveforms resulting from the inductive response of high-speed interconnect. E-beam probing and picoprobeing are the only alternatives commonly available for measuring analog waveforms; these techniques are expensive, difficult due to the need to have top-level metal available for probing, and frequently invasive. Moreover, the advent of systems-on-a-chip design is driving the need for testing analog blocks embedded within largely digital integrated circuits. Test bus standards are being developed to allow limited analog access to internal nodes. An alternate approach to mixed-signal test is to excite the analog block with an on-chip waveform generator (with an on-chip D/A converter) and capture the response in the time-domain with some type of on-chip oscilloscope[1].

MOS transistors have f_T 's beyond several gigahertz, making it possible to generate very-high bandwidth signals on-chip. One can imagine two approaches to carrying the information in a high-bandwidth signal off-chip. With a fixed (but presumably known) latency, one could buffer the digital data off-chip, but all of the analog information would be lost. To preserve the shape of the waveform, one could use an amplifier with unity-gain feedback to buffer the signal off-chip, but practical bandwidth limitations

of the amplifier would limit the signal frequencies that could be sensed to hundreds of megahertz. Fundamentally, the challenge of on-chip measurement circuits is that the circuits performing the measurement are in the same technology as the circuits being measured and, therefore, cannot be made intrinsically "faster."

The key to being able to measure fast waveforms is subsampling. This approach is used in digital sampling oscilloscopes and has been employed in several contexts previously for on-chip measurement circuits[2, 3, 4, 5, 6, 7]. The approach can be understood from both a time-domain and frequency-domain perspective. From a time-domain point-of-view, imagine that we have two clocks, one of period T (and frequency $f_0 = 1/2\pi T$), which we call the *trigger clock* and the other of period $T + \Delta t$ (and frequency $f_s = 1/2\pi(T + \Delta t)$), which we call the *sample clock*. We assume that the waveform that we wish to measure is triggered by the leading edge of the signal clock, as shown in Figure 1, and as such is repeated once each T seconds. If we assume that the sample clock samples the data on its leading edge (and that the sample-and-hold circuit holds the sampled value), then a new time point is sampled each time the waveform is repeated. The output of the sample-and-hold circuit is, therefore, a "spread-out" version of the waveform we wish to measure (as shown in Figure 1), allowing the analog-to-digital converter or other circuits processing the data to be very low bandwidth.¹ From a frequency-domain perspective, the waveform to be sampled has frequency content above f_0 as shown in Figure 2. The sampling process is tantamount to mixing with the frequency f_s . The result of this mixing process is a downshifted spectrum above the beat frequency $f_0 - f_s$. Shifted to lower frequency, the signal is easier to measure with "slower" circuits.

Previous work has considered employing on-chip samplers and on-chip samplers with A/D conversion[3, 4, 5, 6, 7]. These approaches relied on external clocks to generate all timing edges. Reference [2] integrates only the comparator on-chip and relies on the comparator switch point to determine the sample time; this time must be calibrated through an off-chip delay path. In this work, we combine high-bandwidth samplers and on-chip A/D conversion with a digital-to-time converter to produce the first fully-integrated digital oscilloscope on chip. In Section 2, we review possible sampling circuits and consider the sampling circuit used in our design. Section 3 considers the unique features of our digital-to-time converter which allows sub-10-psec timing resolution. Section 4 presents the overall testchip design. Preliminary measurement results are presented in Section 5. Section 6 concludes.

¹In fact, the time scale is magnified by a factor of $T/\Delta t$.




Figure 1: Subsampling in the time domain.

2 Samplers

A critical circuit to the subsampling technique is the sample-and-hold. This is the only circuit component of the on-chip oscilloscope that must have a high bandwidth since it must be able to quickly capture the voltage at the sample clock edge. Most of these samplers are based on a master-slave configuration that is similar to a master-slave flip-flop. One possible sampler circuit is shown in Figure 3[3]. The “master” consists of a nFET pass transistor $M1$ feeding a pFET source-follower unity-gain amplifier. The “slave” is a full pass transistor feeding a second pFET source-follower. The pFET source follower stages provide several advantages. In addition to nearly unity gain, the use of pFET transistors limit the effect of substrate noise. Also the output range of the buffer matches nicely the input range of an nFET differential pair in the preamplifier stage of a comparator. For fast sample clock transitions, the bandwidth of the sampler is dominated by the time constant of transistor $M1$ charging the capacitance of node $N1$, which can easily be made greater than $5GHz$ in $0.25\mu m$ technology[8]. Transistor $M2$ (at half the width of $M1$) is present to help cancel clock-feedthrough and charge-injection noise associated with $M1$. The main limitation of this sampler is that the source-follower buffers cut off at input voltage greater than $V_{DD} - |V_{Tp}|$ and, therefore, one cannot sample full-rail signals.²




Figure 3: Master-slave sample-and-hold with pFET source-follower buffers on both master and slave.

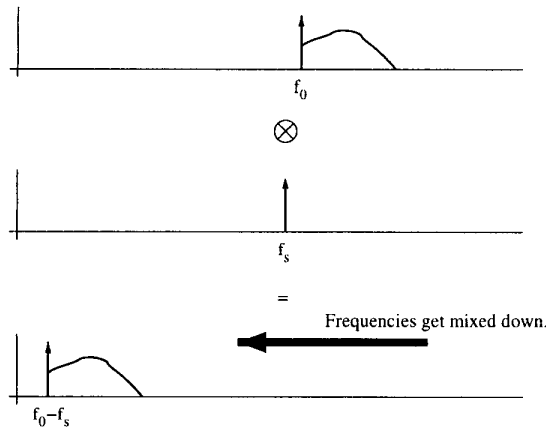


Figure 2: Subsampling in the frequency domain.

The range limitation of the sampler of Figure 3 can be avoided if the buffer is removed from the master, as shown in Figure 4. This is a variation of the sampler used in Reference [6] and is the sampler used in our testchip. Each of the switches is implemented as a full pass transistor.³ In this case, charge-sharing between the implicit capacitances C_1 and C_2 divides down the input voltage to be below the $V_{DD} - |V_{Tp}|$ cutoff of the unity-gain buffer. A variation of this charge-sharing master-slave approach as shown in Figure 5 is used in Reference [7]. In this case, the sampled voltage is converted to a current and amplified by a current mirror to be driven off-chip for measurement.

An important part of the use of the sampler of Figure 4 in our design is that it is calibrated with a separate calibrate input. This calibrate input is driven by an off-chip reference, which can calibrate the entire measurement path to digital output, eliminating errors due to analog mismatch, nonlinearities, and offset in both the sampler and the analog-to-digital converter.

²Of course, the use of a single nFET switch for the master limits the maximum value of the sampled input to $V_{DD} - V_{Tn}$, but the buffer limitation would remain even if a complementary pass-transistor switch were used.

³Note that the first source follower is a “dummy” to match the capacitive loading of the master to the slave. Capacitors C_1 and C_2 are implicitly created by the device loading.




Figure 4: Charge-sharing, master-slave sample-and-hold.




Figure 5: Charge-sharing, master-slave sample-and-hold that buffers to an off-chip current output.

3 Digital-to-time converter

One of the limitations of the subsampling approach illustrated in Figure 1 is that one must generate two tightly controlled clocks off-chip, with the resolution limited by the jitter with which these clocks can be generated. Instead, we wish to generate the trigger and sample edges on chip, with the interval between them determined by digital control, a digital-to-time converter. The circuits required are similar to those employed in time-to-digital converters[9]. The simplest way to build a digital-to-time converter is with a delay-locked loop (DLL) as shown in Figure 6. In this case, the N -stage voltage-controlled delay line (VCDL) is locked to one period (T_{clk}) of the reference clock. This gives each buffer stage of the VCDL a precise delay of T_{clk}/N .⁴ By multiplexing out the outputs of the buffer stages, one could create sample and trigger edges separated by multiples of T_{clk}/N .




Figure 6: Delay-locked loop.

The use of a single DLL, however, limits the time resolution to a gate delay. One technique for overcoming this would be to introduce a circuit to interpolate between the delay stages[10]. Instead, we decided on an approach using two DLLs as shown in Figure 7. In this case, one DLL has a VCDL with N stages and the other has a VCDL with M stages, both locked to the same reference clock. In this case, the delay of each buffer in the first VCDL is locked to a delay of T_{clk}/N and the delay of each buffer in the second VCDL is locked to a delay of T_{clk}/M . By choosing the sample clock from one DLL and the trigger clock from the other, one can achieve multiples of a timing resolution of $T_{clk}/N - T_{clk}/M$, which can be a fraction of a gate delay.

⁴This assumes that the buffer stages are perfectly matched.




Figure 7: Two delay-locked loops used to achieve a timing resolution of $T_{clk}/N - T_{clk}/M$.

The DLLs used in this design will be embedded in a hostile digital environment. As such, they must be as immune as possible to jitter caused by substrate and power supply noise. To accomplish this, the VCDL is constructed with differential buffers as shown in Figure 8 with “symmetric” loads defined by a diode-connected pFET (with a diode-like characteristic) in parallel with a biased pFET (with a triode-like characteristic)[11]. The opposite curvatures of the two characteristics combine to produce a nearly linear load, limiting the conversion of common-mode supply noise into differential jitter. In addition, the buffers are self-biased by a half-replica of the differential pair, locking the lower limit of the output swing to the control voltage V_{ctrl} [12]. There are stability issues associated with this control loop. The loading at the output of the differential amplifier must be sufficient to produce dominant-pole compensation and an overall phase margin of at least 35° , but the loading cannot be so large as to reduce the closed-loop bandwidth excessively and limit dynamic power-supply noise rejection. A loading of about ten buffer stages per bias generator appears to be the appropriate compromise.




Figure 8: Two delay-locked loops used to achieve a timing resolution of $T_{clk}/N - T_{clk}/M$.

The digital-to-time converter on our testchip combines two DLLs, one with $N=32$ and the other with $M=34$ with a 200 MHz reference clock. The buffer stages are carefully matched and the outputs of the buffers are multiplexed to produce time separations between the trigger and sample clock in steps of 9.2 psec ($T_{clk}(1/N - 1/M)$) up to 5 nsec (T_{clk}) as shown in Figure 9.

4 Test chip design

The overall design of the testchip is shown in Figure 10. The control logic steps the digital-to-time converter through increments of 9.2 psec, from a user-specified start time. The 8-bit analog-to-digital converter (ADC) uses a successive-approximation (SA) algorithm and a two-capacitor serial DAC[13] as shown in Figure 11. The capacitors in the serial DAC are implemented using metal-insulator-metal (MIM) capacitors between *metal4* and a special metal layer, giving a capacitance of $1 fF/\mu m^2$. The comparator design is shown in Figure 12[14]. In track mode, the comparator has a gain of approximately 12.5 with the gain around the




Figure 9: Digital-to-time converter combining two DLLs locked to the same reference clock.

positive feedback loop shunted to be less than one (ensuring stability). In latch mode, the regenerative action is enabled, producing nearly full-rail output. This track-and-latch architecture gives good comparator resolution without the need for a multistage amplifier. The overall SA ADC design, though slow, is fairly area-efficient, consuming only 0.015mm^2 .

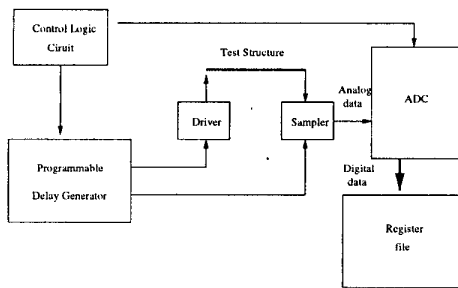


Figure 10: Overall testchip design.

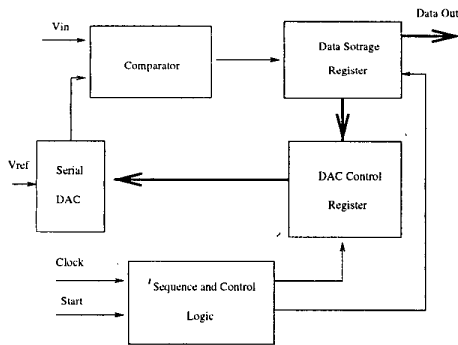


Figure 11: ADC design.

The testchip was designed in the TSMC $0.25\mu\text{m}$ 5M1P process. This is a 2.5-V process with transistor saturation currents at maximum overdrive of about $600\mu\text{A}/\mu\text{m}$ for the nFET and $300\mu\text{A}/\mu\text{m}$ for the pFET. There are five levels of AlCu interconnect. The first four levels have sheet resistivities of 0.076 Ohms/square. *metal5* has a sheet resistivity of 0.044 Ohms/square. A die photo of the fabricated testchip is shown in Figure 13. Sixteen samplers, with the circuit schematic of Figure 4, are positioned to measure various waveforms within a snaking




Figure 12: Track-and-latch comparator.

4-mm-long, 16-bit bus structure. The drivers of the bus are designed to switch with one of three strengths or hold the net high or low. The receiver loads are also variable with MOS switches determining variable amounts of MOS capacitance that can be added to the far-end. The configuration of the testsite is determined by a set of scan-only flip-flop which set the driver and receiver configurations and enable one of the samplers. All of the samples are stored in a 2048-bit register file, which can be scanned out after measurement completion.




Figure 13: Die photo of test chip.

5 Results

The testchip layout was extracted (resistances, capacitors, and inductors) using the Assura RLCX extraction engine[15]. The bus structure was explicitly designed to accentuate inductance effects. Figure 14 shows the waveforms predicted from extraction for bit 7 switching alone (solid curves) and in the presence of simultaneous same-direction switching of all of the bits of the bus (dotted curve) for the strongest driver setting. In the case of simultaneous switching, the response is clearly inductive, with a “porch step” at the near-end and ringing at the far end.

The *actual measured* near-end and far-end responses in the presence of simultaneous switching for the same driver conditions are shown in Figures 15 and 16 (circles at the data points). The solid curve is the simulation result with *no inductance*. Because

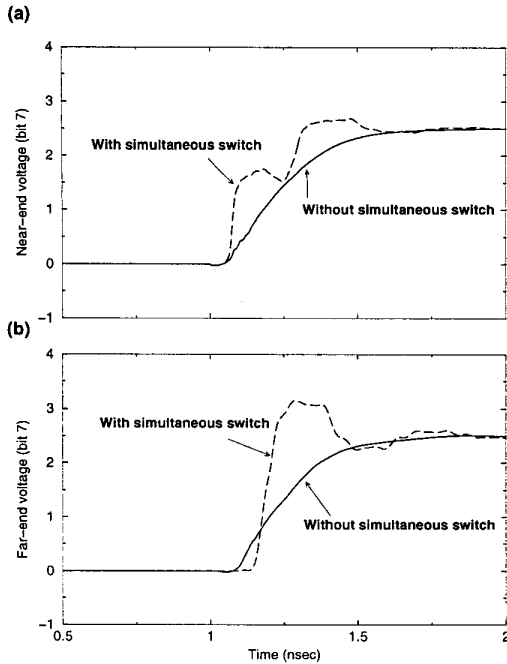


Figure 14: Simulation results: (a) near-end and (b) far-end response of a switching bit 7 of the 16-bit parallel bus with and without simultaneous switching of the other 15 bits.

we also have a sampler positioned on the trigger clock, the (x -axis) delays represent true measured values. Even though a little ringing appears to exist in the far-end waveform, the overall waveform appears to be strictly RC. Crosstalk noise due to the switching of all the other bits while bit 7 is quiet is shown in Figure 17. The solid-curve is the strictly RC simulation. We suspect that we have overestimated the impact of inductance in this test structure because of the neglect of current returning in the substrate.⁵

The timing resolution of the digital-to-time converter is limited by the jitter of the DLLs as well as by error in the matching of the buffer stages of the VCDL. We measured the rms jitter of the digital-to-time converter sample and trigger clocks by measuring the phase-noise spectral density around 200 MHz with a spectrum analyzer⁶ and integrating it according to[16]:

$$t_{rms}^2 = 8T_{clk}^2 \int_0^\infty S_\phi(f) \sin^2(\pi f T_{clk}) df$$

We find $t_{rms} \cong 18psec$. Subtracting (in an rms way) the measured jitter of the reference clock of approximately $17.5psec$ rms, we estimate a jitter in the difference of the trigger and sample edges of approximately $6psec$ rms.

There were several bugs in the chip, which we are correcting in another fab release. A short circuit in one of the pads resulted in large substrate currents. This resulted in leakage currents at the samplers, producing most of the “noise” in the measurement results presented. There was also a small timing bug in the controls to the digital-to-time generator which corrupted some of the data

⁵The substrate for this chip is epitaxial, with a lightly-doped epitaxial layer approximately $7 \mu m$ thick on top of a heavily-doped substrate.

⁶We did not have access to a high-bandwidth oscilloscope to measure the peak-to-peak jitter directly in the time domain.




Figure 15: Circles represent the measured near-end data on bit 7 with all the bits of the bus switching simultaneously. The solid curve is the RC-only simulation.




Figure 16: Circles represent the measured far-end data on bit 7 with all the bits of the bus switching simultaneously. The solid curve is the RC-only simulation.

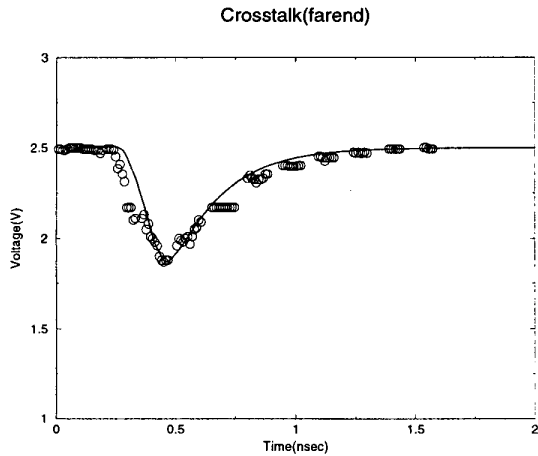


Figure 17: Circles represent the measured crosstalk on bit 7 (far-end) due to the simultaneous switching of all of the other bits of the bus. The solid curve is the RC-only simulation.

point (these were omitted in the measurement data presented, resulted in the evident “gaps” in the data). We are also extending Assura RLCX to incorporate an extracted model of the substrate to enable better correlation between our extraction and measurement results.

6 Conclusions and applications

In this paper, we have described the first, self-contained, full on-chip sampling oscilloscopes for the measurement of high-speed analog waveforms in digital and mixed-signal integrated circuits. The chip employs subsampling techniques enabled by an on-chip digital-to-time converter with sub-10-psec resolution. 8-bit digital data from an area-efficient successive-approximation analog-to-digital converter is stored in a scannable register file.

To employ this technique within the design-for-testability (DFT) methodology of a digital integrated circuit, samplers would have to be positioned near each critical net “tap” point. The samplers themselves are very small, consuming only $100\mu\text{m}^2$. The digital-to-time converter and the ADC can be shared across all of the samplers and can be positioned anywhere on the chip. The digital-to-time converter utilized in our testchip is far larger than it needs to be because we were generating edges from a fairly slow 200 MHz clock, which we could easily bring from off-chip. The DLLs could be much smaller with a high-frequency (PLL-derived) system clock.

References

- [1] G. W. Roberts. Improving the testability of mixed-signal integrated circuits. In *International Custom Integrated Circuits Conference*, pages 214–221, 1997.
- [2] K. Soumyanath, S. Borkar, C. Zhou, and B. A. Bloechel. Accurate on-chip interconnect evaluation: A time-domain technique. *IEEE Journal of Solid-State Circuits*, 34(5):623–631, 1999.
- [3] P. Larsson and C. Svensson. Measuring high-bandwidth signals in CMOS circuits. *Electronics Letters*, 29(20):1761–1762, September 1993.

- [4] K. Lofstrom. Early capture for boundary scan timing measurement. In *Proceedings of the IEEE International Test Conference*, pages 417–422, October 1996.
- [5] M. A. Burns. Understampling digitizer with a sampling circuit positioned on an integrated circuits. U. S. Patent 5578935, 1996.
- [6] A. Hajjar and G. W. Roberts. A high speed and area efficient on-chip analog waveform extractor. In *Proceedings of the IEEE International Test Conference*, pages 688–697, October 1998.
- [7] R. Ho, B. Amrutur, K. Mai, B. Wilburn, T. Mori, and M. Horowitz. Applications of on-chip samplers for test and measurement of integrated circuits. In *Symposium on VLSI Circuits Digest of Technical Papers*, pages 138–139, 1998.
- [8] H. O. Johansson and C. Svensson. Time resolution of NMOS sampling switches used on low-swing signals. *IEEE Journal of Solid-State Circuits*, 33(2):237–245, 1998.
- [9] J. Christiansen. An integrated high resolution CMOS timing generator based on an array of delay locked loops. *IEEE Journal Solid-State Circuits*, 31:952–957, July 1996.
- [10] C.-K. K. Yang and M. A. Horowitz. A $0.8\mu\text{m}$ CMOS 2.5 Gb/s oversampling receiver and transmitter for serial links. *IEEE Journal Solid-State Circuits*, 31:2015–2023, December 1996.
- [11] J. G. Maneatis and M. A. Horowitz. Precise delay generation using coupled oscillators. *IEEE Journal of Solid-State Circuits*, 28(12):1273–1282, 1993.
- [12] J. G. Maneatis. Low-jitter process-independent DLL and PLL based on self-biased techniques. *IEEE Journal of Solid-State Circuits*, 31(11):1723–1732, 1996.
- [13] R. E. Suarez, P. R. Gray, and D. A. Hodges. All-MOS charge redistribution analog-to-digital conversion techniques – Part II. *IEEE Journal Solid-State Circuits*, SC-10(6):379–385, December 1975.
- [14] B. Song, H. S. Lee, and M. Tompsett. A 10-b 15-MHz CMOS recycling two-step A/D converter. *IEEE Journal Solid-State Circuits*, 25(6):1328–1338, December 1990.
- [15] K. L. Shepard, D. Sitaram, and Y. Zheng. Full-chip, three-dimensional, shapes-based RLC extraction. In *Proceedings of the International Conference on Computer-Aided Design*, 2000.
- [16] A. Hajimiri, S. Limotyrakis, and T. H. Lee. Jitter and phase noise in ring oscillators. *IEEE Journal Solid-State Circuits*, 34(6):790–804, June 1999.