Circuits and Techniques for High-Resolution Measurement of On-Chip Power Supply Noise

Elad Alon, *Student Member, IEEE*, Vladimir Stojanović, *Member, IEEE*, and Mark A. Horowitz, *Fellow, IEEE*

***Abstract*—This paper presents a technique for characterizing the statistical properties and spectrum of power supply noise using only two on-chip low-throughput samplers. The samplers utilize a voltage-controlled oscillator to perform high-resolution analog-to-digital conversion with minimal hardware. The measurement system is implemented in a 0.13-μm process along with a high-speed link transceiver. Measured results from this chip validate the accuracy of the measurement system and elucidate several aspects of power supply noise, including its cyclostationary nature.**

***Index Terms*—Analog–digital conversion, power delivery validation, random noise, spectral measurement, supply noise measurement.**

I. Introduction

As CMOS technology has scaled, supply voltages have dropped while chip power consumption has remained constant or even increased, causing chip currents to increase. Even at constant chip power, in order to maintain a fixed percentage noise budget on the power supply, the combination of lower voltage and higher current forces the required supply grid impedance to drop with the scaling factor squared. As shown in Fig. 1, in today's 1-V 100-A high-performance microprocessors, the required impedance for a 10% noise budget is roughly 1 mΩ. Achieving this low impedance across a broad range of frequencies can be extremely challenging, and therefore supply noise has become an issue even for digital logic circuits.

CAD tools are using increasingly sophisticated models of the distribution network in order to give chip designers the ability to find potential problems in simulation [1]. However, circuits to measure supply noise and validate these models have not been as fully developed.

In general, it is very difficult to measure a full time-domain supply voltage waveform during the entire course of chip operation because either an extremely high speed analog interface or analog–digital (A/D) converter would be required to avoid filtering or aliasing. Therefore, measurement circuits have been designed to observe specific properties of supply noise that can be extracted without high-speed interfaces or converters.

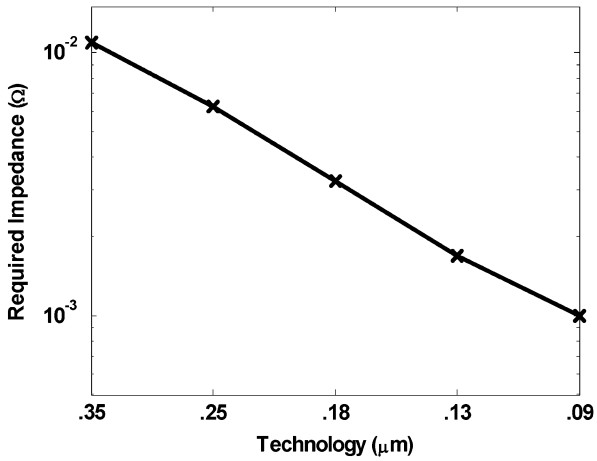

Fig. 1. 10% noise budget required impedance for a 100-W chip across technologies.

Muhtaroglu *et al.* proposed a circuit that is effective in detecting overshoot and undershoot events over a window of time [2]. By performing repeated measurements at different threshold levels and tracking the percentage of time that the supply crosses each threshold, this type of circuit can provide information about the distribution of supply noise. Under controlled testing conditions, chips can be operated in a repetitive manner to facilitate the use of sub-sampling "on-chip oscilloscopes" such as those used by Takamiya *et al.* to measure the supply voltage [3].

Unfortunately, sub-sampled waveforms measured in a repetitive mode may not capture supply noise behavior during normal chip operation. Samplers that do not rely on averaging for accuracy could be used to collect data about the distribution of the supply noise at each point in time (e.g., an oscilloscope with "infinite persistence" turned on). While this measurement can provide a large amount of information about the behavior of the supply noise, the dynamics of the noise would remain uncharacterized. This is a limitation because both the distribution and frequency spectrum of the noise must be known in order to characterize the effect of supply noise on circuits—especially on sensitive analog or mixed signal circuits.

We extend the previous measurement techniques by treating supply noise as a random process, and use its statistical properties—in particular, its autocorrelation [4]—to measure the noise spectrum. Autocorrelation can be measured using only two samplers with precise sampling instants, but low sampling rates. We further extend this technique to measure the dynamics of a

Manuscript received August 31, 2004; revised November 29, 2004. This work was supported in part by C2S2, the MARCO Focus Center for Circuit and System Solutions, under MARCO Contract 2003-CT-888.

E. Alon and M. A. Horowitz are with the Department of Electrical Engineering, Stanford University, CA 94305 USA, and also with Rambus, Inc., Los Altos, CA 94022 USA (e-mail: elad@stanford.edu).

V. Stojanović was with Rambus, Inc., Los Altos, CA 94022 USA and also with the Department of Electrical Engineering, Stanford University, CA 94305 USA. He is now with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: vlada@mit.edu).

Digital Object Identifier 10.1109/JSSC.2004.842853

repetitively time-varying, or cyclostationary [5], noise process by recording the sampling instants in addition to the distance between the samples. This additional information captures the periodic structure of the noise, and better characterizes its behavior. Since only one of the two samplers required to measure autocorrelation is needed to collect the distribution of the noise, this technique can measure all of the supply noise properties necessary to calculate the behavior of circuits affected by this noise.

In order to achieve the sub-mV level resolution necessary to collect supply noise statistics with relatively simple circuitry, we used a voltage-controlled oscillator (VCO) to perform A/D conversion on the noise samples. Results from a 0.13-μm high-speed link transceiver chip on which the samplers were implemented characterize the measurement system's accuracy and present both deterministic and cyclostationary properties of supply noise.

II. Random Supply Noise and Autocorrelation

While supply noise may actually be deterministic in nature, the number of state variables that would need to be tracked to precisely calculate it is enormous. Therefore, instead of finding its exact behavior, supply noise is usually modeled as a random process, and can be characterized by its autocorrelation or its frequency spectrum. This spectrum is one of the properties necessary to calculate the effects of the noise on sensitive circuits. In addition, the shape of the frequency spectrum can give insights into the source of the noise.

For a time-invariant (also known as stationary) random process, the autocorrelation (R) is defined as

$$R(\tau) = E[v_{\text{noise}}(t+\tau/2)v_{\text{noise}}(t-\tau/2)] \qquad (1)$$

where τ represents the separation of the two noise samples in time and $E[\,]$ returns the expected value of the random process. The Fourier transform of the autocorrelation gives the power spectral density (PSD) of the noise process. For example, as shown in Fig. 2(a), a white, zero mean noise process has a spectral density that is flat and infinitely wide—correspondingly, its autocorrelation is an impulse whose magnitude is set by the variance (σ^2) of the noise process. Intuitively, if a noise process contains high frequency components, it will change within a short period of time, and therefore samples of the noise that are far apart in time will not be correlated with each other.

As an additional example of the relationship between autocorrelation and noise spectrum, consider a low-pass filtered white noise process. The filter spreads each noise impulse in time, and thus closely spaced samples of the noise are highly correlated with each other. As shown in Fig. 2(b), this causes the autocorrelation to become broader, matching the shape of the impulse response of the filter.

Thus far, the discussion of autocorrelation has been limited to stationary noise processes that do not vary with time. Most chips run synchronously to one or more clocks, and these clocks can modulate the occurrence of noise events. For example, as shown in Fig. 3, switching events may be more likely to occur at the beginning of the clock cycle than at the middle or end since all of the clock drivers and flip-flops toggle near the rising edge of the clock. Because of this modulation, it is unlikely that the properties of supply noise will be independent of time. However, since this modulation is repetitive, at a particular point in time relative to the cycle, the properties of the noise should be the same. This is known as a cyclostationary process, and can be characterized by measuring the autocorrelation (and hence PSD) of the noise at each time point in the cycle. In other words, the autocorrelation R becomes a function of both the time separation τ and the time t at which the samples were taken.

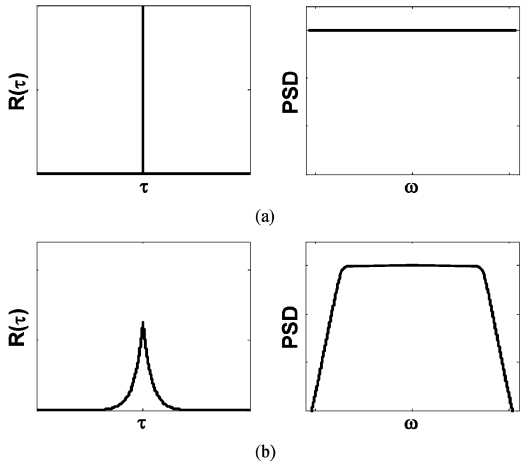

Fig. 2. Autocorrelation and power spectral density for (a) white noise and (b) low-pass filtered white noise.

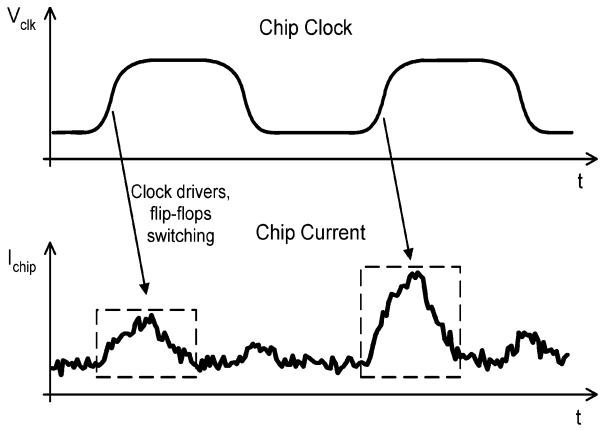

Fig. 3. Example of the clock modulating the chip current consumption waveform.

The most important characteristic of autocorrelation for the measurement system is that it is an average statistical property of the noise, and therefore we do not need to know its exact behavior at all times in order to extract its frequency content. The Nyquist frequency of the measurement is set by the minimum time spacing between the samples—not by the repetition rate of the sampling—which greatly reduces the requirements on the throughput of the sampling circuits. A further benefit of measuring autocorrelation[1] is that any uncorrelated, additive noise that is independent between the samplers (e.g., thermal noise)

[1]In practice, in order to reduce the sensitivity of the measured results to residual offset errors in the samplers, it is often desirable to measure autocovariance, $C(\tau) = E[(v_{\text{noise}}(t+\tau/2) - E[v_{\text{noise}}(t+\tau/2)])(v_{\text{noise}}(t-\tau/2) - E[v_{\text{noise}}(t-\tau/2)])]$.

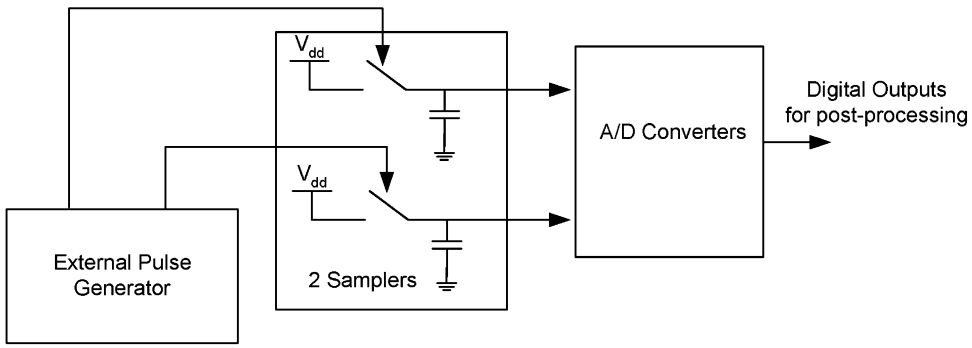

Fig. 4. Supply noise measurement system block diagram.

can be eliminated because the expected value of the product of two uncorrelated noise terms is zero.

III. Measurement Circuits

A block diagram of the noise measurement system is shown in Fig. 4. We used an external pulse generator to control the samplers' timing in order to allow a high degree of flexibility in the range of time spacing between the samplers, the sampling rate, and the ability to lock the samplers to different noise cycles. As we will describe in the calibration section, on-chip hardware is included in the measurement system to allow calibration of any mismatch in the routing of the two sampling signals to the measurement circuit.

To avoid additional noise that would be coupled into the measurement system from communicating analog quantities off the chip, A/D conversion is performed internally. Since all the digital samples are post-processed using measured calibration curves, the linearity and offset requirements of this on-chip converter are not stringent. This allows a simple and compact design that can be integrated onto many different parts of the die.

A. Sampling Switch

The sampling switch is the only circuit in the measurement system that must be on the die in order to measure supply noise relative to on-chip $V_{\rm ss}$—it is also the only circuit whose bandwidth must be high enough to avoid filtering high frequency noise content. The bandwidth of the sampling switch should meet or exceed the bandwidth of the highest bandwidth circuits on the chip so that the measurement will accurately capture the supply noise behavior as it is seen by these circuits. For example, in many digital designs the highest bandwidth circuits are fanout-of-four (FO4) inverters, and therefore the sampling switch would need a bandwidth above that set by the rise time of these inverters. While previous designs such as [3] and [6] used NMOS switches, when sampling $V_{\rm dd}$ it is easier to obtain the required bandwidth using PMOS switches.

In order to minimize coupling between $V_{\rm dd}$ and the sampled node during hold mode, the samplers operate off of their own, slightly higher than nominal supply ($V_{\rm ddQ} = 1.3$ V instead of 1 V). Since any noise on $V_{\rm ddQ}$ (e.g., due to uncoupled ground-bounce) will couple onto the sample node through the sampling switch parasitic capacitances, this supply is heavily bypassed to on-chip $V_{\rm ss}$.

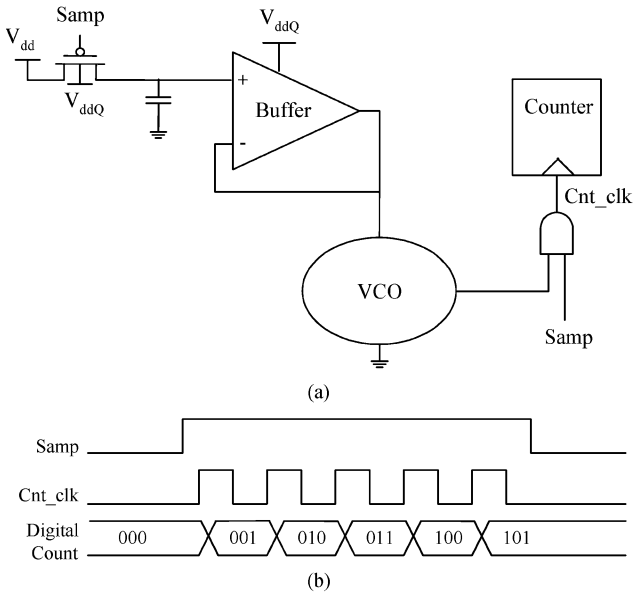

Fig. 5. (a) VCO converter schematic. (b) Example waveforms.

B. VCO Converter

Conversion is achieved by using the sampled voltage to set the frequency of a VCO, and then digitally estimating this frequency by counting the number of clock edges the oscillator outputs over a specified window of time, as shown in Fig. 5.

Despite the fact that the oscillator is essentially free-running, the quantity that is being digitally measured is the average VCO frequency over the conversion window. Therefore, this scheme is insensitive to high frequency jitter (induced either by thermal noise or noise from $V_{\rm ddQ}$) because the counting process averages out cycle-to-cycle variations in the oscillator's period.

Unlike high frequency jitter, the averaging does not attenuate noise that is on the same time scale as the counting window. Fortunately, for noise frequencies too low to be averaged by the counter, on- and off-chip $V_{\rm ss}$ are very well correlated. Therefore, even if the on-chip bypass capacitance does not fully couple $V_{\rm ss}$ to $V_{\rm ddQ}$ at these frequencies, the external regulator and bypass capacitors will. Experimental results presented in the next section show that $V_{\rm ddQ}$ remains well-coupled to $V_{\rm ss}$ even at these frequencies.

The buffer and VCO were implemented with a regulator and pseudo-differential regulated supply CMOS buffers [7]. Since V_{ddQ} must be quiet for the system to perform accurate measurements, the supply rejection requirements of the regulator are not very strict.

The resolution achieved by the VCO converter is set by the minimum change in sampled voltage necessary to measure a difference of one count over the conversion window

$$1\ \mathrm{LSB} = 1/(T_{\mathrm{win}} K_{\mathrm{vco}}) \tag{2}$$

where K_{vco} (in Hz/V) is the VCO gain, and T_{win} is the conversion time. One of the advantages of the VCO-based approach is its flexibility in achievable resolution; as long as leakage in the sampling switch is negligible, the resolution of the converter can be improved by increasing the conversion time.

Since the phase of the oscillator is random with respect to the start of the counting window, the oscillator effectively adds a uniformly distributed noise with a magnitude of 1 LSB to the output of the converter. This noise can be helpful because it creates dither in the output of the converter, allowing the resolution of the converter to increase as the measurements are externally averaged together.

In order to maximize the resolution for a given conversion time, K_{vco} should be as large as possible. Hence, if implemented as a ring, the operating frequency of the VCO should also be high. While ring oscillators can easily achieve 6–8 FO4 cycle times, building a large adder with a latency of less than that would consume too much area and power. Fortunately, the counter's latency does not need to be set by the operating frequency of the VCO—only its throughput must be high enough that no counts are lost. Therefore, we used a simple toggle counter to minimize the required hardware.

If the VCO is implemented as a ring oscillator, the resolution can be further increased by counting the edges out of all phases of the VCO instead of only a single phase.[2] By tracking all phases in the oscillator, the effective VCO cycle time can be reduced to the delay of a single stage of the VCO, which can be 1 FO4 or less. This technique may be desirable in technologies where leakage in the sampling switch limits the resolution of the overall converter.

One of the drawbacks of the VCO-based converter is that the voltage to frequency gain of the oscillator may fluctuate with temperature. Depending upon the VCO implementation, temperature variations can cause offset and gain errors in the measurement. However, with additional circuitry such as a local thermometer—which could be implemented by another VCO converter sampling a constant voltage—the temperature dependence of the oscillator could be measured during calibration and cancelled in post-processing.

IV. Measurement Results

The proposed supply noise measurement circuits were implemented on a 0.13-μm CMOS process along with four high-speed links and a central ASIC that controls the links [8], [9]. The circuits have the capability to measure both the digital supply (V_{dd}) and the supply dedicated to the analog blocks such as the phase-locked loop (PLL) and phase interpolators (V_{ddA}). In addition, noise generators were integrated to allow testing and validation of the measurement system.

[2]This does not actually require multiple counters, only that the phase state of the VCO is recorded at the beginning and end of the conversion window.

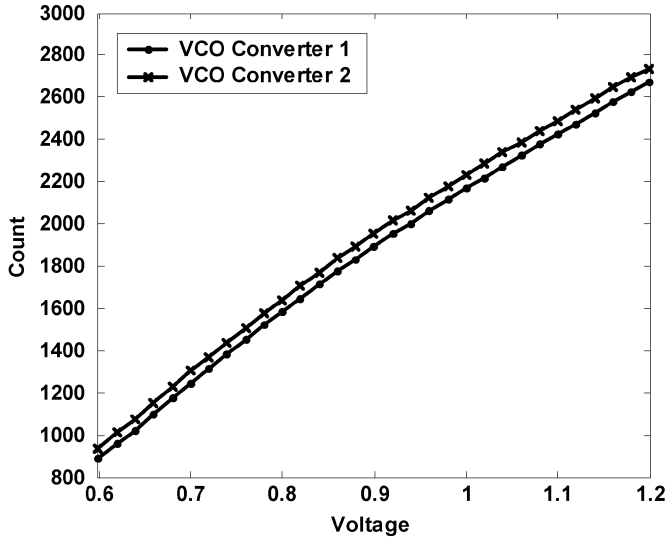

Fig. 6. VCO calibration curves with a 1-μs conversion window.

A. Measurement System Calibration

Calibration is a key element of the measurement system since it enables the realization of a high degree of accuracy with relatively simple (and nonideal) circuits. The first step in the calibration is to characterize the digital code versus input voltage curves of the two VCO-based converters. To perform this calibration, the rest of the chip is placed in a power-down state and the supply voltage is swept across the range of interest using an external power supply while the average digital code at each voltage is recorded. The measured VCO count versus voltage curves of the two samplers with a 1 μsec conversion window are shown in Fig. 6. At a 1-V supply, K_{vco} is roughly 2.6 GHz/V, corresponding to a nominal LSB of 385 μV.

In addition to voltage calibration of the A/D converter, timing calibration is necessary to characterize any skew between the two sampling signals caused by mismatches in routing from the pulse generator to the measurement circuits. To simplify the calibration procedure, two flip-flop phase detectors, shown in Fig. 7(a), are included as part of the measurement circuit.

To measure the timing skew, the external pulse generator is used to sweep the timing offset between the two sampling signals and the outputs of the phase detectors are collected multiple times at each setting. The result of this measurement, shown in Fig. 7(b), is the percentage of time that each phase detector outputs a 1 (i.e., for phase detector 1, Samp_1 arrives after Samp_2) at each timing offset. If the setup times of the two flip-flops are identical, the timing skew can be calculated by finding the center point between the two 50% probability points.

Clearly, in using this timing calibration method any mismatch in the setup times of the two phase detectors will result in an error in the estimated skew. Since the setup time of the flip-flops is generally on the order of the minimum step-size in τ, this

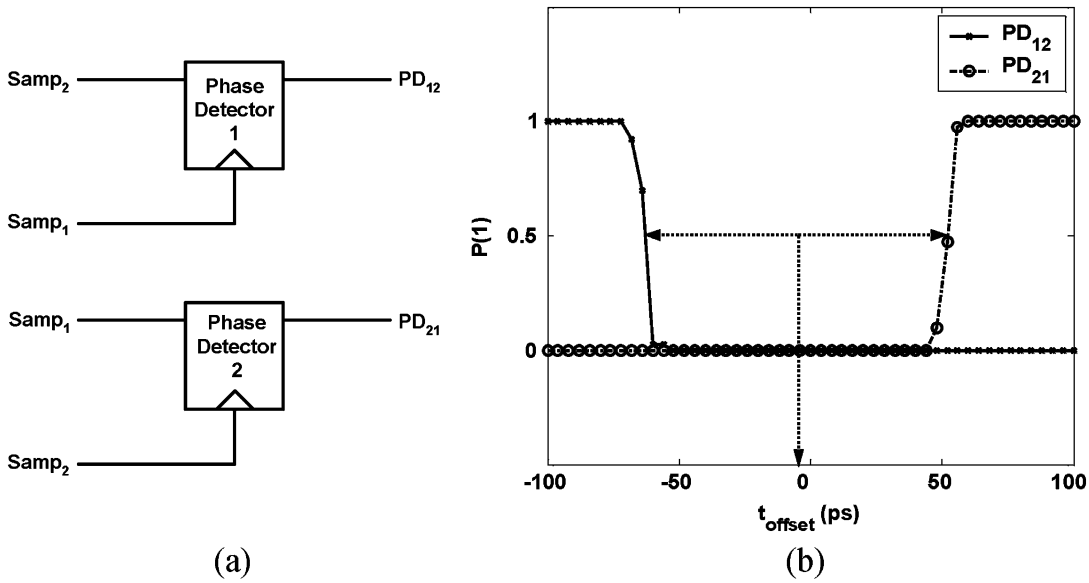

Fig. 7. (a) On-chip flip-flop phase detectors for timing calibration. (b) Measured probability of each phase detector outputting a 1 versus timing offset between Samp_1 and Samp_2. Negative t_{offset} indicates Samp_1 is skewed to arrive after Samp_2.

timing error will not cause a large error in the measured autocorrelation. In addition, this calibration error can be removed by making use of the fact that the autocorrelation will be maximized when the two sampling signals are exactly aligned.[3]

B. Measurement System Validation and Characterization

Since noise on V_{ddQ} directly affects the sampled node, the main accuracy concern for this measurement system is the degree to which V_{ddQ} remains coupled to on-chip V_{ss}. In order to characterize this effect, we used the noise generator to inject square wave current onto the supply grid while the rest of the chip was inactive.

The injected current causes the chip's power supplies to collapse toward each other. As shown in Fig. 8(a), if V_{ddQ} is not perfectly coupled to on-chip V_{ss},[4] it will effectively move in the same direction as V_{dd} because of the change in V_{ss}. Since the VCO integrates V_{ddQ} noise over the conversion window, if the conversion window is the same width as the injected pulse, this uncoupled ground-bounce would cause the measured waveform to be somewhat triangular instead of square. The height of this triangle relative to the amplitude of the square provides a worst-case bound on the relative error of the measurements due to shifts in V_{ss} causing V_{ddQ} noise. With a 500-ns conversion window, the measured 1-MHz and 4-MHz waveforms of Fig. 8(b) exhibit only negligible differences in magnitude and shape, showing that uncoupled ground-bounce on V_{ddQ} is minimal for this chip.

Despite the fact that these generated waveforms are deterministic, their spectral densities can be characterized using autocorrelation to validate the measurement procedure. A 32-MHz square wave was created on V_{dd} and the measured waveform and PSD are shown in Fig. 9. The PSD exhibits the odd harmonics associated with the square-wave waveform.

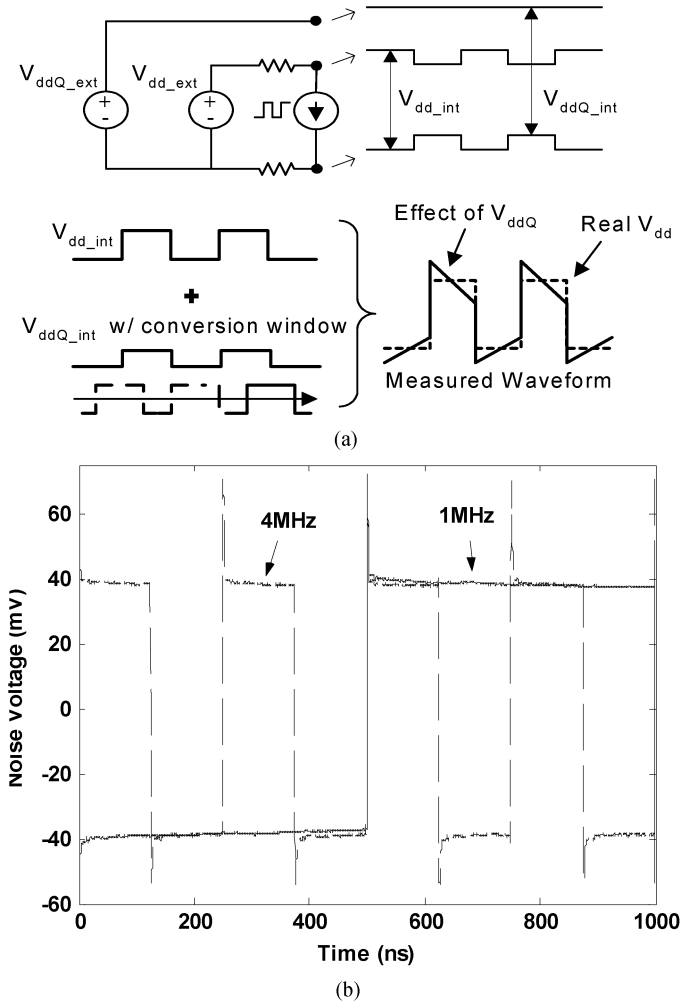

Fig. 8. (a) Potential effect of uncoupled V_{ss} noise through V_{ddQ} on measured waveform. (b) Measured noise injected on V_{ddA} (mean subtracted) at 1 MHz and 4 MHz with a 500-ns conversion window.

Finally, to complete the characterization of the samplers, the noise floor of the system was measured in two different ways.

[3]Because of the variance associated with estimating autocorrelation using only a finite number of samples, this method will also result in calibration errors; however these errors can be minimized with a large number of samples.

[4]That is, if on-chip V_{ddQ} and V_{ss} do not shift by exactly the same amount relative to a fixed external reference.

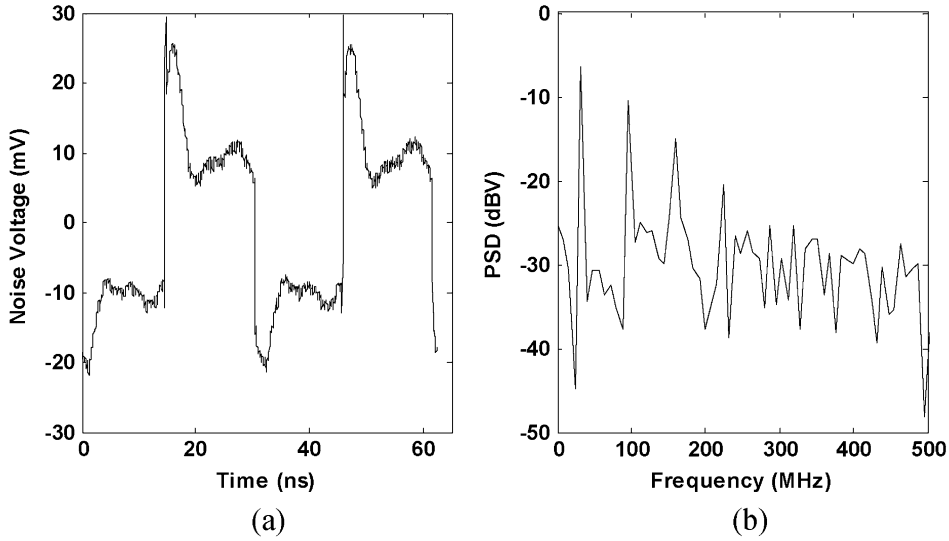

Fig. 9. Measured (a) waveform and (b) PSD of injected 32-MHz square wave on $V_{\rm dd}$. The PSD has been scaled by the Nyquist frequency (10 GHz) to calculate the power and shown with the unit (dBV) instead of the usual (dB V^2/Hz). For PSDs in dBV, the average level in dB corresponds roughly to the noise σ.

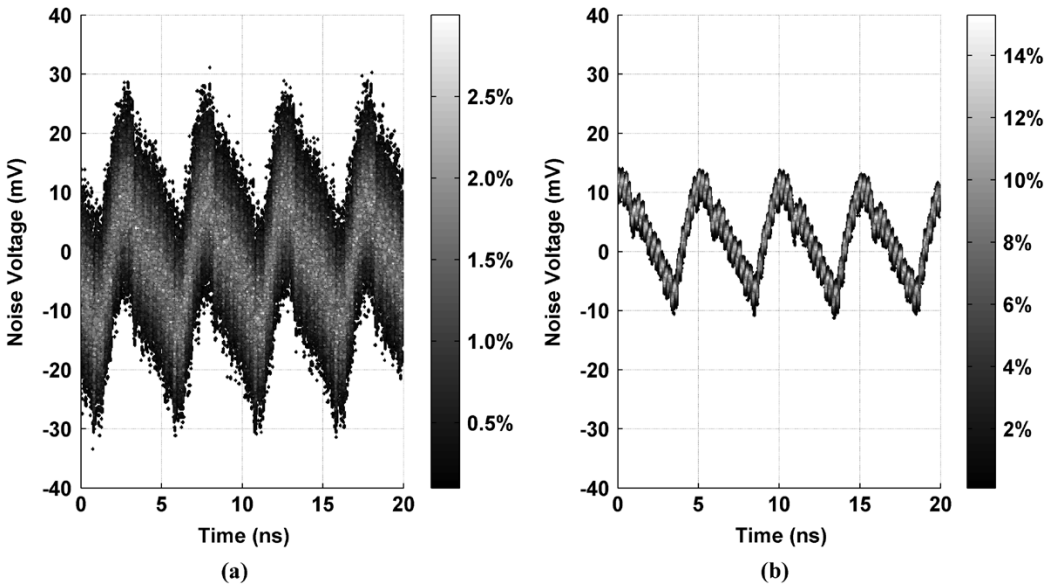

Fig. 10. Measured noise distributions on (a) $V_{\rm dd}$ (b) $V_{\rm ddA}$.

First, using a single sampler with the chip powered down, 100 K samples were taken and the voltage distribution around the mean was recorded; the measured σ was 800 μV. Second, both samplers were used to measure the noise floor in the frequency domain. The measured noise floor of the PSD (overlaid in Fig. 12) was 300 μV. The reduction from the single sampler is due to the averaging of the noise sources that are uncorrelated between the two samplers.

C. Measured Supply Noise

With the characterization of the samplers complete, the chip was activated with the links running at 4 Gb/s in serial loop-back on 2^{31} PRBS data. Before proceeding to measure autocorrelation, we used a single sampler as a sub-sampled oscilloscope to collect the distribution of the supply noise at each point in time (Fig. 10). This measurement can provide a large amount of information about the behavior of the power supplies. For example, the noise on $V_{\rm dd}$ exhibits much larger random variations (and correspondingly a larger peak-to-peak variation) than the noise on $V_{\rm ddA}$. In fact, the large majority of the noise on $V_{\rm ddA}$ is due to a deterministic, repetitive noise waveform; this waveform can be extracted by taking the mean of the supply noise voltages at each point in time (Fig. 11).

To characterize the dynamics of both this deterministic waveform and the random supply noise, we measured the stationarized, or time-averaged PSD with the chip operating under the same conditions (Fig. 12). As can be seen by the spikes in the spectral densities of Fig. 12, the most dominant component of the repetitive waveforms appears at 200 MHz. Since the clock frequency of the ASIC core is 200 MHz for this data rate, and the core has its own supply voltage but a shared $V_{\rm ss}$, this noise

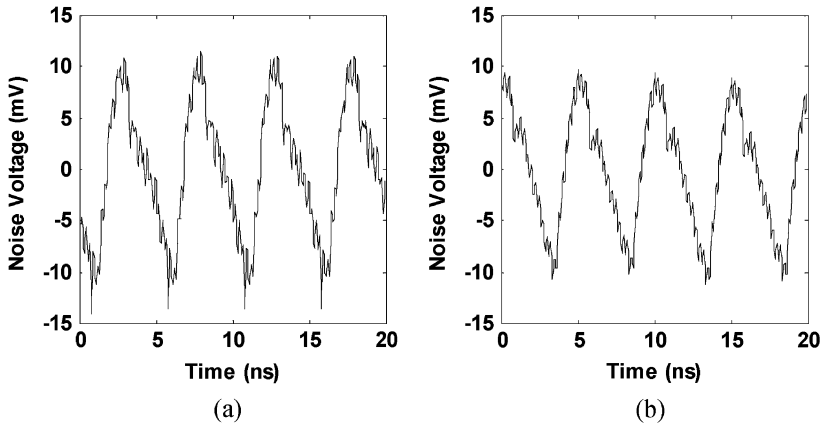

Fig. 11. Measured deterministic noise on (a) $V_{\rm dd}$ and (b) $V_{\rm ddA}$.

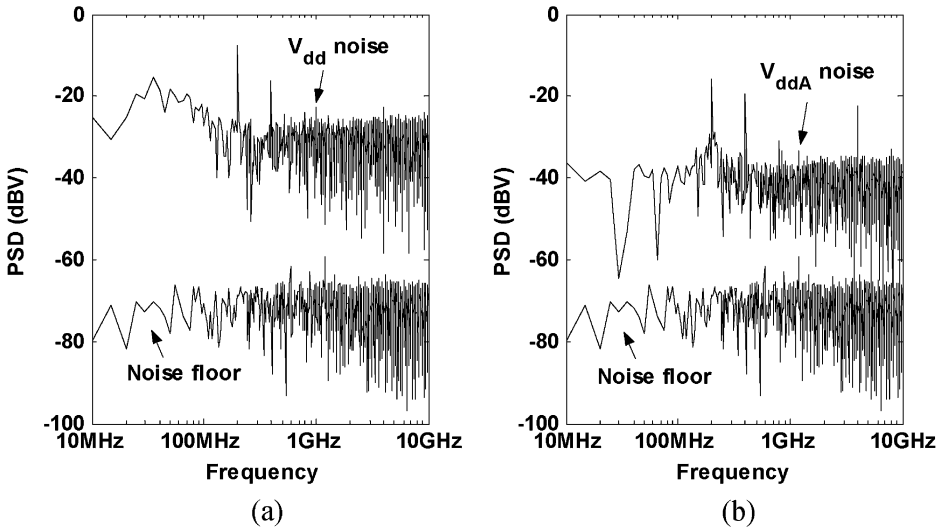

Fig. 12. Measured PSD with the noise floor for (a) $V_{\rm dd}$ and (b) $V_{\rm ddA}$.

is likely due to ground or substrate noise that is not fully decoupled to the supplies, explaining why the waveforms on $V_{\rm dd}$ and $V_{\rm ddA}$ look very similar. An additional noise component exists at 400 MHz, which is the frequency of the reference clock that is fed to the PLLs in the links and drives some of the digital logic. Finally, the noise waveform also has a component at 4 GHz, likely due to transitions in the clock buffers that drive both polarities of the 2-GHz main clock and modulation of the tail currents in the differential pairs used in the transmitter, VCO, clock buffers, and phase interpolators.

Both the distribution measurements from Fig. 10 and the spectral densities from Fig. 12 indicate that a large portion of the supply noise is due to these deterministic variations. Except for the frequency components of the deterministic signal, the PSD of $V_{\rm ddA}$ is nearly flat, implying that the noise is essentially white in nature. The PSD of $V_{\rm dd}$ is also relatively flat; however it exhibits some low frequency peaking in the 20–100-MHz range. Since the voltage waveform shown in Fig. 9(a) created by injecting a square wave current onto $V_{\rm dd}$ exhibits damped sinusoidal ripple, this peaking is most likely due to the transfer function of the supply distribution network. In contrast, the square wave current response of $V_{\rm ddA}$ exhibits only overshoot on the transitions and no sinusoidal ripple, matching the relatively flat PSD. The difference in the impedance of the two distribution networks is intentional; $V_{\rm ddA}$ has a more resistive distribution network because it supplies less power than $V_{\rm dd}$ and because the additional series resistance is beneficial in isolating $V_{\rm ddA}$ from external noise.

For circuits that respond to noise equally at all points in time, the stationary PSD gives a complete description of the supply noise dynamics. However, the random noise is actually cyclostationary in nature. In the time-averaged spectral densities of Fig. 12, the noise appears white because its time-varying behavior has been averaged by the measurement procedure, which for the stationary measurement samples the supply voltage uniformly throughout the noise cycle.

For cyclostationary noise processes, both the distribution and the dynamics of the noise can vary within the cycle. The cycle at which the characteristics of the noise repeat is not necessarily the clock cycle of the chip, especially if the chip has multiple clocks or may exhibit somewhat repetitive modes of operation (e.g., a code loop running on a processor). Since the noise cycle may not be known *a priori*, the most straightforward method to detect the noise cycle is a guess-

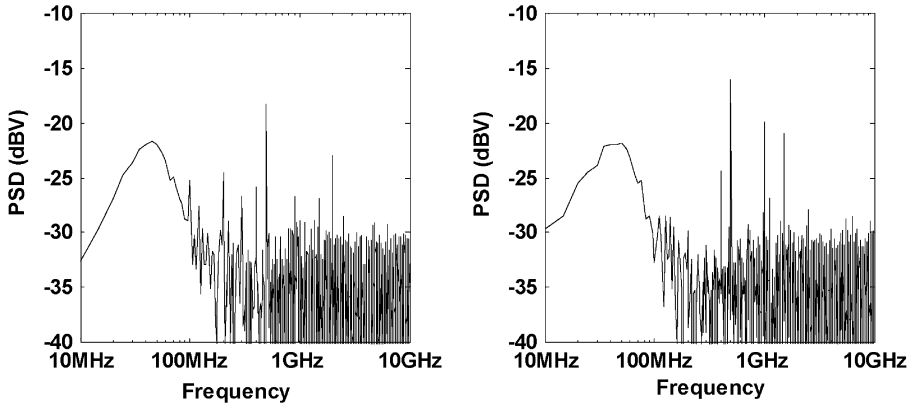

Fig. 13. Measured PSDs of $V_{\rm dd}$ noise at two different times in the noise period with the links operating at 2 Gb/s.

and-check approach. In many chips all of the clocks and/or events are harmonically related to each other, and therefore the measurement can be taken assuming that the noise cycle is some integer multiple of the slowest clock cycle in the system. If the integer multiple that was chosen is large enough that the assumed noise cycle encompasses two or more true noise cycles, the measured characteristics of the supply noise will repeat within the assumed cycle.

On this chip the cyclostationary behavior was not very pronounced, and therefore, instead of presenting the cyclostationary spectral densities at each time in the cycle, we measured the PSD of the random noise at two different times in order to observe examples of noise variation. Additionally, in order to make the cyclostationary behavior more apparent, we decreased the operating frequency of the links to 2 Gb/s. At this reduced rate, the majority of the logic should complete before the beginning of the next clock cycle, causing a period of relative calm on the supply.

Although the total power in the PSD shown on the right of Fig. 13 is only slightly higher than that of the PSD on the left, their dynamics are clearly different—the PSD on the right has a strong component at 1 GHz that is not present in the other PSD. Since a large portion of the transceiver runs off of a 1-GHz clock at this data rate, these two spectral densities correspond to times of relatively low and relatively high switching activity in the link circuits.

V. Conclusion

We have presented a measurement system for characterization of power supply noise using two low-rate samplers to enable autocorrelation (and hence noise spectrum) measurement. The system made use of a VCO-based A/D converter and calibration procedure in order to achieve high-resolution measurements with relatively simple circuitry.

The system was used to measure supply noise on a high-speed link chip, and a deterministic signal of $\sim$20 mV peak-to-peak was detected on both the digital and analog supplies. In addition to this deterministic signal, the nondeterministic portion of the noise was shown to exhibit cyclostationary behavior, and example time-indexed spectra showing the cyclically varying properties of this noise were presented.

With the integration of these supply noise measurement systems, designers can use verified noise distributions and spectral densities to characterize both the supply grid and the impact of supply noise on their circuits. Furthermore, the autocorrelation based measurement technique can be extended to measure other types of noise, such as clock jitter. As the bandwidths of on-chip signals rise and testing with external equipment becomes more challenging, sub-sampling and autocorrelation based on-chip measurement techniques will become increasingly useful in the characterization of on-chip signals and noise behavior.

Acknowledgment

The authors would like to acknowledge the help and support of V. Abramzon and B. Nezamfar of Stanford University, and A. Ho, F. Chen, C. Werner, J. Zerbe, S. Pamarti, K. Chang, F. Assaderaghi, B. Daly, Y. Frans, B. Garlepp, J. Hsu, I. Ghibanescu, and R. Kollipara of Rambus, Inc. E. Alon and V. Stojanović would like to thank M. Lee and I. Stojanović for their constant help and support.

References

[1] S. R. Nassif and J. N. Kozhaya, "Fast power grid simulation," in *Proc. IEEE Design Automation Conf.*, Jun. 2000, pp. 156–161.

[2] A. Muhtaroglu, G. Taylor, and T. Rahal-Arabi, "On-die droop detector for analog sensing of power supply noise," *IEEE J. Solid-State Circuits*, vol. 39, no. 4, pp. 651–660, Apr. 2004.

[3] M. Takamiya, M. Mizuno, and K. Nakamura, "An on-chip 100 GHz-sampling rate 8-channel sampling oscilloscope with embedded sampling clock generator," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2002, p. 182.

[4] A. Leon-Garcia, *Probability and Random Processes for Electrical Engineering*, 2nd ed. Reading, MA: Addison-Wesley, Jul. 1993.

[5] B. J. Skinner, F. M. Ingels, and J. P. Donohoe, "Stationary and cyclostationary random process models," in *Proc. IEEE Southeastcon'94*, Apr. 1994, pp. 450–454.

[6] R. Ho *et al.*, "Applications of on-chip samplers for test and measurement of integrated circuits," in *IEEE Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 1998, pp. 138–139.

[7] K.-Y. K. Chang *et al.*, "A 0.4–4 Gb/s CMOS quad transceiver cell using on-chip regulated dual-loop PLLs," *IEEE J. Solid-State Circuits*, vol. 38, no. 5, pp. 747–754, May 2003.

[8] V. Stojanović *et al.*, "Adaptive equalization and data recovery in a dual-mode (PAM2/4) serial link transceiver," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2004, pp. 348–351.

[9] A. Ho *et al.*, "Common-mode backchannel signaling system for differential high-speed links," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2004, pp. 352–355.

Elad Alon (S'02) was born in Haifa, Israel, in 1979. He received the B.S. and M.S. degrees in electrical engineering from Stanford University, Stanford, CA, in 2001 and 2002, respectively, and is currently working toward the Ph.D. degree at Stanford University.

He was a visiting researcher at Hewlett Packard Company, Ft. Collins, CO, in the summer of 2003, where he implemented supply noise measurement circuits on an Itanium processor. He has also been a visiting researcher at Rambus Inc., Los Altos, CA, since 2003, where he has worked on supply noise measurement and high-speed signaling circuits. His research interests include efficient on-chip power supply regulation and distribution, noise measurement techniques, high-speed interface design, and applications of optimization and nonlinear control to high-speed mixed-signal circuits.

Vladimir Stojanović (S'96–M'05) was born in Kragujevac, Serbia, Yugoslavia. He received the Dipl. Ing. degree in electrical engineering from the University of Belgrade, Serbia, in 1998, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 2000 and 2005, respectively.

He was also with Rambus, Inc., Los Altos, CA, from 2001 to 2004 and a Visiting Scholar with the Advanced Computer Systems Engineering Laboratory, Department of Electrical and Computer Engineering, University of California, Davis, in 1997-1998. He is currently an Assistant Professor in the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA. His current research interests include design, modeling, and optimization of integrated systems, from standard VLSI blocks to CMOS-based electrical and optical interfaces. He is also interested in design and implementation of digital communication techniques in high-speed interfaces and high-speed mixed-signal IC design.

Mark A. Horowitz (S'77–M'78–SM'95–F'00) received the B.S. and M.S. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1978, and the Ph.D. degree from Stanford University, Stanford, CA, in 1984.

He is the Yahoo Founder's Professor of Electrical Engineering and Computer Science at Stanford University. His research area is in digital system design, and he has led a number of processor designs including MIPS-X, one of the first processors to include an on-chip instruction cache, TORCH, a statically scheduled, superscalar processor that supported speculative execution, and FLASH, a flexible DSM machine. He has also worked in a number of other chip design areas including high-speed and low-power memory design, high-bandwidth interfaces, and fast floating point. In 1990 he took leave from Stanford to help start Rambus Inc., Los Altos, CA, a company designing high-bandwidth chip interface technology. His current research includes multiprocessor design, low power circuits, memory design, and high-speed links.

Dr. Horowitz received the Presidential Young Investigator Award and an IBM Faculty development award in 1985. In 1993, he received the Best Paper Award at the IEEE International Solid-State Circuits Conference.