

Design Considerations for Stochastic Analog-to-Digital Conversion

Skyler Weaver¹, Daniel Knierim², Roland Wazenried², and Un-Ku Moon¹

¹School of EECS, Oregon State University, 1148 Kelley Engineering Center, Corvallis, OR 97330. USA

²Tektronix, Inc., 14200 S.W. Karl Braun Drive, PO Box 500, Beaverton, OR 97077. USA

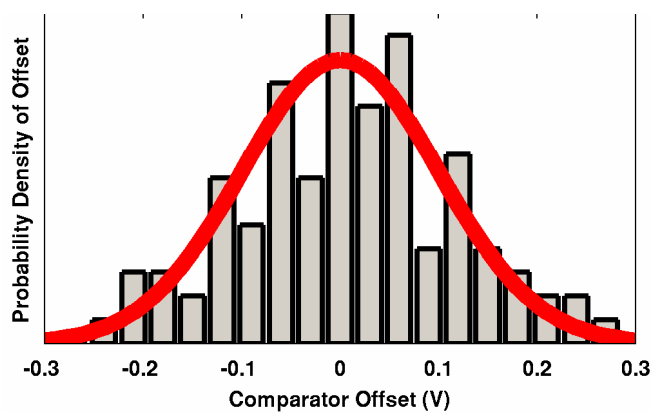
Abstract—Using knowledge of the random nature of device mismatch it is possible to employ stochastic techniques in an ADC. This would allow the use of many smaller and less accurate components, and may make it possible to save power and area while maintaining accuracy. Additional benefits include high scalability and high yield with process variation. Various design considerations are suggested that must be addressed to design a successful stochastic ADC and an example calibration and decoding scheme is described.

I. INTRODUCTION


Two very significant challenges in analog design are dealing with device mismatch and noise, both of which are considerable challenges in the scaling of the analog portion of a mixed-mode system into leading edge digital processes. One promising solution to these limitations is to use the stochastic nature of device mismatch to our advantage as introduced in [1].

Consider some large array of identical comparators that all have some input-referred offset due to device mismatch and process variation. The offset will follow some sort of probability density function (PDF), for example Gaussian, as seen in Fig. 1. The actual distribution will then be a discrete set of the continuous PDF due to a finite number of comparators. If all comparators are then connected in parallel, the digital output of the number of comparators that evaluate high combined with information about the distribution of the comparator offsets may be sufficient to reconstruct the input. For example, if a comparator set was known to have a zero-mean Gaussian offset distribution with standard deviation (σ) of 100-mV and half of the comparators evaluate high while the other half evaluate low, then the input is close to the reference. Likewise, if eighty-four percent evaluate high, then the input is close to 100-mV above the reference.

The scope of this paper is to suggest some possible applications of this technique and discuss some of the issues that must be addressed in order for a full stochastic analog-to-digital converter (ADC) to be successful.



(a)



(b)

Figure 1. a) Probability density function (PDF) of a Gaussian distribution, continuous and finite (with $N=100$). b) Demonstrating that $1-\sigma$ more than half of the comparators evaluating high corresponds with input being $1-\sigma$ above the reference voltage (with $\sigma = 100\text{mV}$).

II. APPLICATIONS

One application of this method would be to consider comparators in the second stage of a 2-step subbranging flash ADC. Depending on the speed, noise, and other specifications, these analog comparators will generally have high power requirements and dominate a large area footprint

This research was funded by Tektronix, Inc. and in part by DARPA under contract F33615-02-1-1179.

in order to achieve the desired level of accuracy. Offset cancellation techniques are usually required such as output-offset-storage (OOS) as extensively seen in [2]. Generating accurate voltage references for the subranging flash stage can also be challenging. It may be possible to replace these large comparators with many small, less accurate comparators. Since conventional offset cancellation will not be required, there can be an overall area and power savings. Another benefit of this method is that the reference ladder could be replaced with a single reference, since the random input-referred offsets due to device mismatch act as random reference voltages with some distribution that is related to the PDF of device mismatch.

Another plausible application would be to use a stochastic ADC as the comparator of a successive approximation register ADC (SAR). The comparator could be quickly calibrated by shorting the inputs to the stochastic ADC and counting the number of comparators that evaluate high. During normal operation the sum of the comparator outputs would be measured against the calibrated value to determine the appropriate bit for the SAR. The added performance would come at the final comparison when the input signal is very small. Instead of resolving one LSB, the knowledge of the actual number of comparators that evaluate high could be used to resolve multiple bits. In fact, this concept of resolving extra LSBs on the final comparison could be applied to other ADC architectures, such as pipelined ADCs.

One additional benefit of a stochastic comparator is its robustness and superior yield in the presence of process defects. Compare the performance of a comparator consuming a given amount of power and area, and a stochastic comparator consisting of n -comparators each consuming $1/n$ power and area. The overall area and power cost is the same for both cases. The stochastic comparator should have superior yield compared to a single, larger comparator since a stochastic comparator could theoretically still operate even if a large number of its comparators were completely useless due to process defects whereas a single comparator may not have any functionality if a large portion of its area were compromised.

III. DESIGN CONSIDERATIONS

A. Intentional Offset

The offset distribution for a given set of comparators is dependent on the comparator architecture, device sizes, process parameters, and many other factors. Since many comparators are required, it may be advantageous to limit the comparators to a relatively simple architecture with near minimum sized devices so that power and area are reasonable. If this is the case, then the variance of the offset distribution is largely a function of the process and may give an undesirable spreading of the comparator offsets. Device mismatch is usually modeled as a Gaussian distribution [3], yet a Gaussian distribution is actually not the most desirable for two reasons.

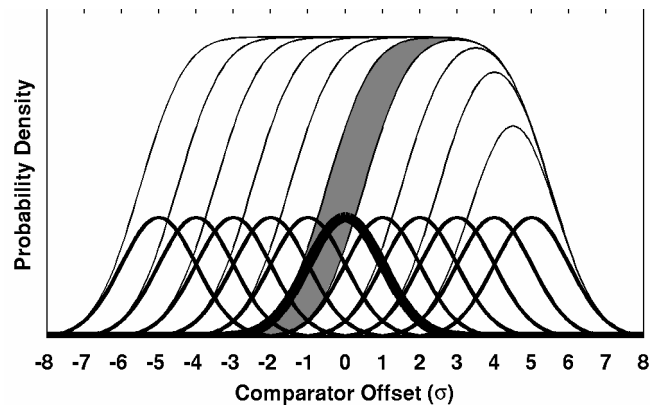


Figure 2. The combined PDF of many Gaussians spaced $1\text{-}\sigma$ apart. The darker Gaussian PDF corresponds to the shaded area of the combined PDF.

First, since a Gaussian distribution is not uniform, there will be more comparators with trip-points close to the mean of the distribution than with trip-points near $1\text{-}\sigma$ away from the mean. There will be even less comparators around the $2\text{-}\sigma$ point. From this one can infer that higher accuracy is attainable around the mean than at points further away from the mean. Consider that for a given ADC accuracy a certain average comparator offset spacing is required. If σ is equal to 100-mV and the specified accuracy is required at 300-mV from the mean, then the total number of comparators needed to satisfy this condition will be very large; however, the comparator offset spacing around the mean will then be equivalent to a much higher accuracy than required. This results in wasted area and power by having unnecessary and redundant comparators near the mean.

Second, any non-uniform distribution implies that the average spacing of comparator offsets is also not uniform. Since an ideal ADC should have a linear transfer function, more post-processing will be required to generate a linear transfer characteristic.

To remedy this, there are generally two types of intentional offset that can be applied to a comparator array in order to adjust the distribution. These two offset types can be thought of as fine and coarse offsets.

Fine offset describes applying a different intentional offset to each comparator according to some function in an attempt to control the overall distribution. If a large array of comparators were designed to have built-in linearly increasing intentional offsets, their actual offsets would be a function of the intentional offset and the random unintentional offset. The effective PDF of this comparator array will be the convolution of the intentional offset distribution and the PDF of the random offset. This could be implemented as an intentional mismatch in the comparator input pairs, or in a latch-based comparator as a mismatched output capacitance.

Coarse offset describes applying a single intentional offset to an entire group of comparators. This is in essence

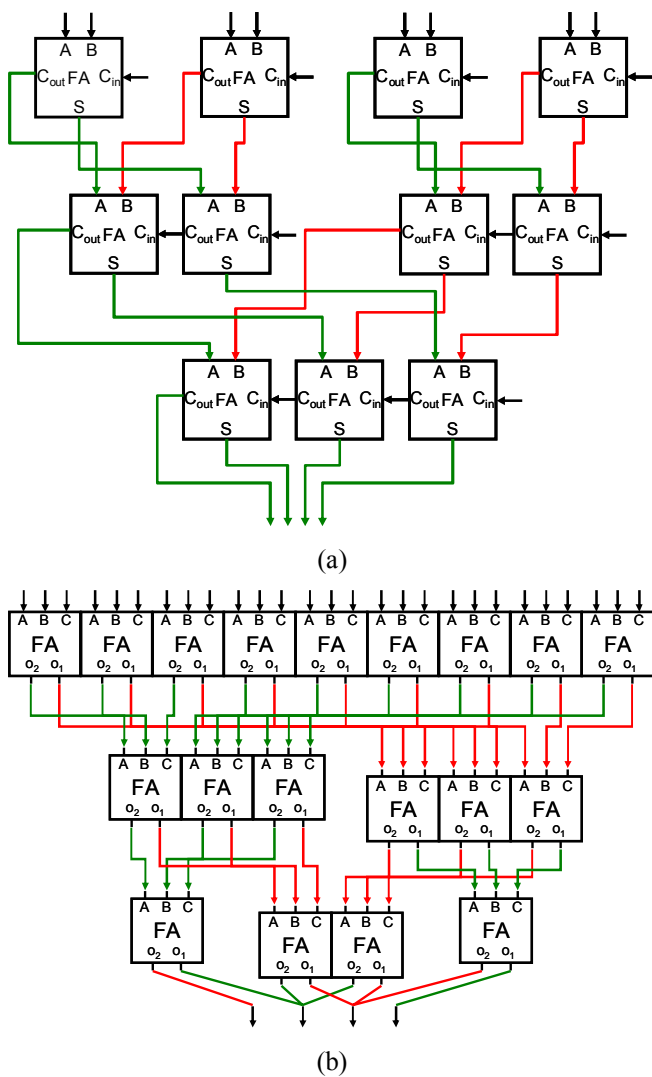


Figure 3. a) A 15-input, 4-bit output ripple carry binary counter based on ones adder. b) A 27-input Wallace tree based ones-adder where o_1 and o_2 are bit-weighted equal to the input and twice the input, respectively (i.e. $o_1 = S$ and $o_2 = C_{out}$).

the structure of a typical FLASH with each single comparator replaced with a stochastic comparator. These offsets could be implemented as different reference voltages or any of the intentional mismatches already described. Spacing a number of identical comparator groups with Gaussian PDFs $1-\sigma$ apart, as seen in Fig. 2, will result in an effective PDF with a linear region at the center of the distribution, since the areas of each PDF will be maintained under the combined PDF.

B. Counting

For a standard flash ADC, $2^n - 1$ comparators are required to resolve n -bits. In a stochastic version a much larger number of comparators are required to achieve the same accuracy since by definition each comparator is less

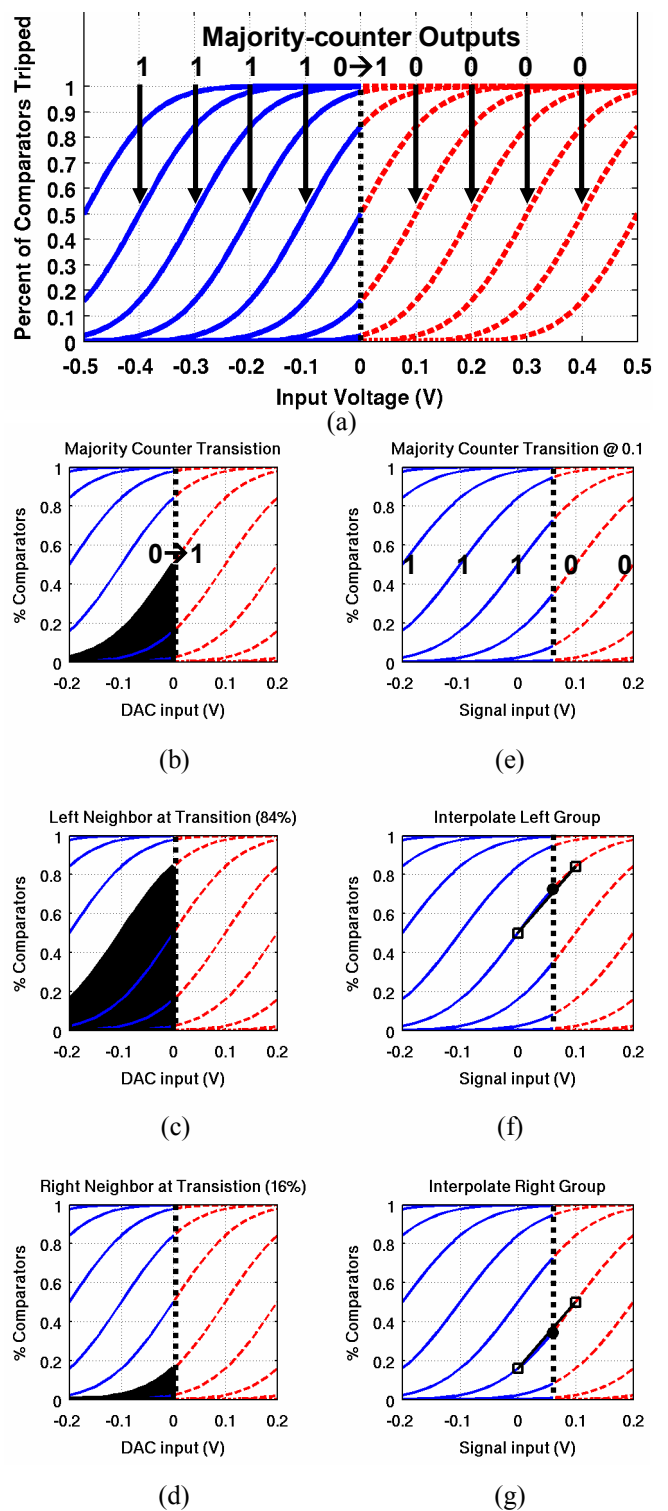


Figure 4. a) CDFs of $1-\sigma$ spaced comparator groups ($\sigma = 100\text{-mV}$) with zero input, group majority counts shown. b) On majority counter transition, DAC value is stored as mean of group. c,d) Full counts of neighboring groups stored for interpolation. e) Group majority counts for $V_{in} = 60\text{-mV}$. f,g) Input is resolved by interpolation.

accurate. Given that the number of comparators is large there needs to be a fast way to sum all of the comparator outputs into a binary weighted word.

One solution is a hierarchy of binary addition, as seen in Fig. 3a, where all comparator outputs are taken as 1-bit numbers and added using full adders. The result is many 2-bit numbers that are then added together to result in 3-bit numbers. This process continues until the output has enough bits to represent the number of comparator outputs. Another method of adding many single bit-valued inputs is by implementation of a Wallace-tree based ones-adder as suggested in [4] as a thermometer-to-binary converter. Both of these methods consume a similar amount of power and area, so either might be chosen by layout considerations.

It is important to note that it may not be necessary to count every comparator output if the comparators have coarse intentional offsets as shown in Fig. 2. If this is the case, then the cumulative density functions (CDF) of the comparator offsets will appear as in Fig. 4a. Note that in this configuration some comparator groups will have almost all high outputs, while others will have almost all low outputs. The most information is contained in the group or groups where some comparators evaluate high and some evaluate low. Time spent counting could be saved if only the groups where nearly half evaluated in either direction are counted. This could be facilitated by implementation of a majority-counter that merely needs to determine if more than half of comparator outputs are high or low; in essence, treating each comparator group as a single comparator. This will result in a thermometer code from the output of the majority-counters. The groups where there is a transition in the thermometer from high to low are the only groups that need to be counted, thus saving time.

C. Calibration and Decoding

Although the trip-points of all comparators should follow some PDF, the actual offset value of each comparator instance will not be identical from process run to process run, or even chip to chip from the same wafer. Therefore, it is necessary to implement a start-up calibration scheme that will be used to characterize the comparator offsets.

One possible calibration and decoding scheme is to consider the case where there are comparators in coarse-offset groups with majority count and full count information available for each group. Calibration could be achieved by connecting a slow, accurate DAC to the input. The input voltage will ramp until the first majority-counter flips from zero to one. On this transition, the value of the DAC is stored as the median of this particular comparator group (Fig. 4b). As shown in Fig. 4c,d, the full counts of neighboring groups

are also stored. The DAC continues to ramp until all majority-counters output high.

With the calibration process complete, the input signal can now be applied. First, the majority-counters create a thermometer code which indicates the point of interest at the transition from ones to zeros (Fig. 4e). The full counts of the two groups about the thermometer code transition are used to resolve the input. Since the input will lie between the median of one group and the median of its neighboring group, the input is known to lie within the range of two DAC values obtained from calibration. Further resolution can be achieved by interpolating between two known points of each group, as shown in Fig. 4f,g. Interpolation of the group below the majority count transition and interpolation of the group above the majority count transition will result in two answers. The final output would be the average of these two results. Linear interpolation may be sufficient; however, Platt scaling described in [5] may give a better fit, as it maps a sigmoid function to a straight line. Here we have assumed that the majority count thermometer code will be monotonically increasing. If monotonicity can not be assured, then another mechanism (as opposed to the majority count transition) will be required to choose the groups to be used for interpolation.

IV. CONCLUSION

Taking advantage of the random nature of device mismatch by employing stochastic techniques in an ADC may make it possible to save power and area while maintaining accuracy. Additional benefits include high scalability and high yield with process variation. Various design considerations were suggested that must be addressed to design a successful stochastic ADC, and an example calibration and decoding scheme was described.

REFERENCES

- [1] J. L. Ceballos, I. Galton, G. C. Temes, "Stochastic Analog-to-Digital Conversion," *Circuits and Systems*, 48th Midwest Symposium on, pp. 855-858, 2005.
- [2] D. J. Huber, R. J. Chandler, and A. A. Abidi, "A 10b 160MS/s 84mW 1V Subranging ADC in 90nm CMOS," *ISSCC Dig. Tech. Papers*, pp. 454-455, February, 2007.
- [3] S. C. Wong, K. H. Pan, and D. J. Ma, "A CMOS Mismatch Model and Scaling Effects," *IEEE Electron Device Letters*, vol. 18, no. 6, pp. 261-263, June 1997.
- [4] E. Säll and M. Vesterbacka, "Comparison of Two Thermometer-to-Binary Decoders for High-Performance Flash ADCs," *Proceedings of IEEE Norchip Conference*, Oulu, Finland, Nov. 21-22, 2005.
- [5] J. Platt, "Probabilistic Outputs for Support Vector Machines and Comparisons to Regularized Likelihood Methods," *Advances in Large Margin Classifiers*, MIT Press, 1999.