

Effectiveness of Adaptive Supply Voltage and Body Bias as Post-Silicon Variability Compensation Techniques for Full-Swing and Low-Swing On-Chip Communication Channels

Giacomo Paci[‡], Davide Bertozzi[†], Luca Benini[‡]

[†] ENDIF, University of Ferrara, 44100 Ferrara, Italy.

[‡] DEIS, University of Bologna, 40136 Bologna, Italy.

Abstract

Adaptive body bias (ABB) and adaptive supply voltage (ASV) have been showed to be effective methods for post-silicon tuning of circuit properties to reduce variability. While their properties have been compared on generic combinational circuits or microprocessor circuit sub-blocks, the advent of multi-core systems is bringing a new application domain forefront. Global interconnects are evolving to complex communication channels with drivers and receivers, in an attempt to mitigate the effects of reverse scaling and reduce power. The characterization of the performance spread of these links and the exploration of effective and power-aware compensation techniques for them is becoming a key design issue. This work compares the variability compensation efficiency of ABB vs ASV when put at work in two representative link architectures of today's ICs: a traditional full-swing interconnect and a low-swing signaling scheme for low-power communication. We provide guidelines for the post-silicon variability compensation of these communication channels.

1. Introduction

As technology continues to shrink, process variations can have a significant negative impact on yield due to the wider spread of performance and power consumption. Post-silicon tuning allows the adjustment of device characteristics after a die has been manufactured to compensate for the specific deviations that occurred on that particular die [21, 22]. One of the methods utilizes the transistor body effect to change transistor threshold voltage by applying an adaptive body bias (ABB) to chip devices to modulate performance and power [21, 30]. The other method of performing post-silicon tuning is to adjust the supply voltage (ASV) to trade performance with power, thus achieving a similar effect to ABB in spite of the different physical mechanism, implementation overhead and trade-off curves. The effectiveness of ABB and ASV in reducing variability has been assessed and compared mainly on combinational logic circuits [24], key elements of microprocessor critical paths [22] and ring oscillators [20], sometimes achieving counterintuitive and even conflicting conclusions. According to [24], the difference in effectiveness is so small that choosing one method over the other should mainly be based on implementation overhead. [20] claims that although the frequency and power tuning range of ABB is more limited than that of ASV, its frequency tuning range proves effective for process-dependent performance compensation. In contrast, [22] concludes that using ASV together with ABB is much more effective than using any of them individually and is worth the cost. In essence, the effectiveness of ASV and ABB should not be assessed in gen-

eral, but with reference to the variance of a specific manufacturing process and to the performance and power tuning requirements of the design at hand. With the advent of multi-core integrated systems, the assessment of post-silicon variability compensation techniques cannot be limited to the traditional testbenches of past research any more, such as combinational logic circuits or even microprocessor circuit sub-blocks. In fact, the new architecture trend requires long (global) interconnects for the connection of system-level blocks with each other. Unfortunately, physical properties of these on-chip interconnects are not scaling well with feature sizes, and they are becoming a key limiting factor for performance, reliability and timing closure of the whole system. A common practice is to overcome the effects of interconnect reverse scaling by means of circuit-level techniques, so that on-chip interconnects cannot be viewed as simple on-chip wires any more, but rather as communication channels including complex drivers and receivers [5, 19]. Analyzing the impact of process parameter variations on the performance and reliability of these communication channels and exploring effective means for their compensation is a key design issue. The relative effectiveness of ABB and ASV may greatly depend on the particular communication circuits they are applied to. A traditional design technique for long links consists of inserting equally spaced CMOS repeaters to deal with resistive loss along the wire. However, with the increase in number and density of the wires with each new technology, interconnect area and power are severely impacted [1]. The most effective technique for global interconnects to achieve significant power savings and energy-delay efficiency is to reduce the voltage swing of the signal on the wire [15] and, possibly, to avoid the use of repeater stages, like in [10]. On the other hand, low-swing signaling reduces noise immunity and poses non-trivial circuit design challenges. In many previous works (e.g., [27]) power, area and delay of communication links making use of full-swing vs low-swing signaling have been compared. This work investigates the robustness of the two signaling techniques to process variations, and assesses the effectiveness of ASV and ABB as variability compensation techniques for full-swing and low-swing interconnects. The ultimate objective is not to characterize the implementation cost of these techniques, but to find out which technique is worth the cost for a specific communication channel and under given process parameter variations. However, aware of the need for low cost compensation, this work also investigates the effectiveness of selective compensation of specific communication channel sub-blocks as opposed to channel-wide tuning. All our tests were conducted on an STMicroelectronics 65nm low-power technology. Given the emerging role of networks-on-chip (NoCs) as reference interconnect fabrics for MPSoC platforms [26], our study targets the on-chip wires used for switch-to-switch connectivity at the top level of the NoC architecture hierarchy.

This work is structured as follows. After reviewing previous work (Section 2), the design and characterization of the communication channels under test are illustrated in Section 3. Their inherent robustness to process variations is analyzed in Section 4, while the effectiveness of ASV and ABB for variability compensation is

addressed in Section 5.

2. Related Work

Most research on low-swing interconnects is focused on designing circuit structures with minimal impact on delay, area and power, so the inherent advantages of low-swing signaling are not swamped by transmitter and receiver overhead. An overview of drivers and receivers is illustrated in [3, 15]. [3] makes a comparison with traditional CMOS circuits and is one of the few papers dealing with repeater stages of low-swing interconnects. The use of repeaters is avoided in [17] by means of a swing limiter and an interconnect accelerator at the receiver. Carefully engineered voltage level converters are proposed in [2, 18], while an optimized level restoration scheme based on bootstrapping can be found in [8]. Sense amplifiers are commonly used to detect a small voltage swing in reduced-swing buses [4, 15]. The minimum interconnect swing should be set by the need to overcome noise at the receiver. An adaptive sensing scheme is proposed in [14] to reduce the threshold voltage offset between a driver and a receiver and ensure low-swing reliable operation. An adaptive voltage swing is set at circuit initialization in [13] to drive interconnects based on their delay, thus coping with the increasing interconnect delay spread. To the limit, a self-calibrating interconnect can be designed [6, 12]. Differential current-mode signaling schemes have a distinctive advantage over the single-ended ones in terms of noise immunity and signal integrity [11]. Neighbor-to-neighbor crosstalk can be reduced with twists in the differential interconnect pairs [9]. Differential low-swing interconnects come at the cost of a significant area and power overhead, therefore are not considered in this paper. Current variation models tend to ignore variations in wires [29], however the spread of technology parameters may jeopardize functionality of transmitting and receiving circuits, causing communication performance degradation or even failure. The traditional techniques for post-silicon compensation of variability are adaptive body biasing (ABB) [21, 25] and adaptive supply voltage (ASV) [23]. Comparative studies of ABB vs ASV when put at work for variability compensation in microprocessor sub-circuits or generic combinational logic circuits have not reached a unique conclusion, proving that the choice is tightly design- and technology-dependent. In [20, 22, 24] there is consensus on the fact that ASV has a larger tuning range of circuit properties and the combined use of ASV and ABB further extends this range. However, the measured yield improvements are different depending on the technology and the design at hand, so it is not unambiguous whether hybrid approaches are worth the cost. In many cases, ABB seems to suffice for the required range of post-silicon compensation. Only for core-to-core variations ASV seems the best compensation option [7]. [16] points out the dependence of ABB and ASV efficiency on the device type and operating temperature in 90nm technology, while [24] emphasizes the role of biasing resolution as well.

This work aims at extending the analyses performed so far to the link architectures for on-chip communication. First, the intrinsic robustness of full-swing vs low-swing signaling schemes to process variations will be explored. Second, ABB and ASV will be applied to find out which extent they can restore the nominal performance of sample communication channels affected by process variations and which is the incurred power cost. Our analysis can justify a later investment in the synthesis backend of nanoscale designs to support the most suitable variability compensation technique for a given communication channel and variation scenario.

3. Communication channel design

We at first present the design of the communication channels that will be assessed later on in terms of robustness to process variations and suitability for traditional post-silicon compensation techniques. Without lack of generality, we restrict our analysis to an intermediate layer wire with a length of 2mm, which is already the typical length of a switch-to-switch link in a regular network-on-chip architecture [26]. Inserting repeaters to reduce delay of

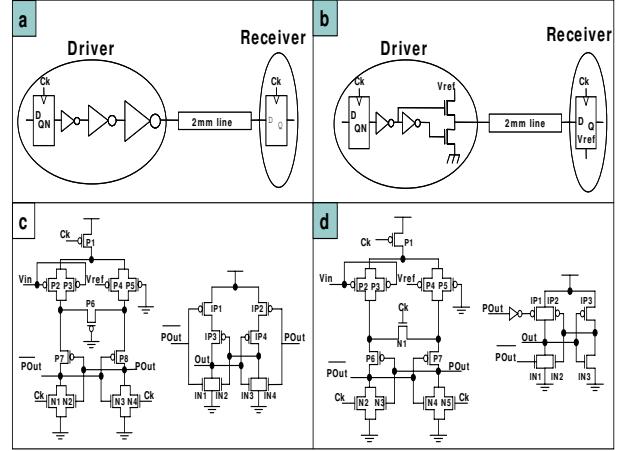


Figure 1. a) CMOS full-swing interconnect. b) Low-swing interconnect. c) PDIFF low-swing receiver from [15]. d) Optimized PDIFF low-swing receiver.

a wire is effective only when the wire is at least twice as long as the critical length of the technology and of the specific routing layer. In our target 65nm technology, a 2mm wire falls below this threshold and the choice is therefore for an unrepeatable interconnect. Even for longer links, solid network-on-chip implementation works like [10] suggests the use of unrepeatable wires for the point-to-point communication links between switches, unlike other scenarios where high-fanout nets are required. To the limit, link pipelining can be used to break long timing paths. Following these indications, this paper assumes the use of unrepeatable wires for network-on-chip communication. We model the on-chip wire by a π distributed RC model. In this work, we assume that capacitive cross-talk has been tackled by means of physical-level techniques such as shielding or proper wire spacing. We leave the analysis of the interaction between cross-talk and variability compensation for future work. The reference link architecture uses a 1V full-swing signaling (Fig.1.a). The driver consists of a library flip-flop and a chain of buffers sized based on exponential horn methodology for minimum delay. The receiver is another library flip-flop. The alternative communication scheme is the low-swing pseudo-differential interconnect architecture reported in Fig.1.b. The basic circuit is taken from [15]. The driver is an NMOS-only push-pull driver which allows the use of very low power supplies and a quadratic energy reduction as a function of the voltage reference/swing V_{ref} . The receiver is still clocked but requires the voltage reference as an additional input. The original receiver circuit proposed in [15] is the clocked sense amplifier followed by a static latch illustrated in Fig.1.c. This pseudo-differential scheme uses single wire per bit while still retaining most advantages of differential amplifiers such as low input offset and good sensitivity. The major reliability degradation may come from the local device mismatch between the double input transistor pairs and from the variation between distant references of the driver and the receiver. In contrast, receiver operation is largely insensitive to V_{dd} supply noise, as opposed to other schemes. This was the basic motivation for picking up this scheme from [15]. However, we applied some improvements to this receiver, ending up with the circuit in Fig.1.d. First, PMOS transistor P6 in Fig.1.c has the task of equalizing the connected nodes, however it remains active even after the initialization, thus slowing down node transients. Moreover, it is not very conductive when the connected nodes reach an initialization value approaching its voltage threshold. In Fig.1.d it has been replaced by an NMOS transistor driven by the clock, thus achieving a better equalization and a faster node transition. Second, although the NOR static latch in Fig.1.c appears to be symmetric, it features uneven 0-to-1 and 1-to-0 switching times. Balancing rise and fall times makes the circuit actually asymmetric. The solution in Fig.1.d allows an easier balancing of these times while

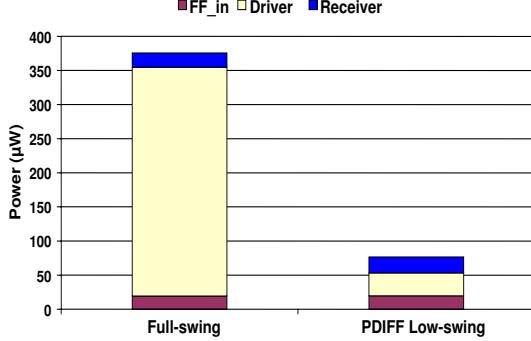


Figure 2. Power breakdown at 1.68 GHz, i.e. the maximum performance achievable by full-swing signaling.

keeping the cross-coupled inverter pair fully symmetric: the outputs of the pseudo-differential receiver in fact directly drive the transistors (dis-)charging the flip-flop output capacitance, while the cross-coupled inverter pair keeps the sampled values. Output capacitance for the differential signal was tuned to be the same for POUT and POUTN signals. As a side effect, the flip-flop in Fig.1.d turns out to scale better from a performance viewpoint and enables higher operating frequencies for a comparable area than that of Fig.1.c. The voltage swing was chosen to be 200mV. Transistor sizing for the low-swing communication channel was done to keep the same (maximum) performance of the full-swing interconnect (1.68 GHz): driver sizing was used to achieve the same link propagation delay, while receiver and static latch sizing was used to enforce the same clock propagation time, so that the next logic stage fed by the communication channel is impacted in the same way. In particular, the library constraints for such propagation time were enforced.

3.1 Characterization of communication channels

We now explore power and area incurred by the full-swing vs PDIFF low-swing signaling schemes in order to provide the same performance. Power results with 100% input switching activity are reported in Fig.2. Our low-swing channel consumes almost 5x less power than the full-swing one, confirming the power efficiency of this solution. Most of the power savings obviously come from the driver and from its reduced reference voltage. The input flip-flop is the same, and so is the power. Moreover, the PDIFF receiver almost equals the power of the library flip-flop in the full-swing scheme, which was chosen with the minimum driving strength. Low-swing signaling also achieves 28.5% lower leakage power. Most of the savings come again from the driver, but also the PDIFF receiver has a lower leakage than the library flip-flop, due to the power gating PMOS transistor in pre-charge mode and to the minimum area NMOS transistors that are switched off in evaluation mode. As regards area, the low-swing channel has a negligible 1% increase in area. The low-swing receiver has a slightly larger area than the library flip-flop, which is counterbalanced by the lower area footprint of the low-swing driver. Please observe that the PDIFF receiver consumes the same total power of the library flip-flop with more area, and this is due to the fact that some of its internal nodes switch with a lower swing. Finally, by modeling and simulating wire lengths larger than 2mm, we got almost the same quadratic delay increase for the full-swing and the low-swing interconnects, since the time constant stays the same. Given a target frequency for a network-on-chip design, the NoC must ensure a maximum link length, eventually enforced by applying link pipelining techniques.

4. Robustness to process variations

The first objective of this paper is to compare the inherent robustness of full-swing and PDIFF low-swing signaling schemes

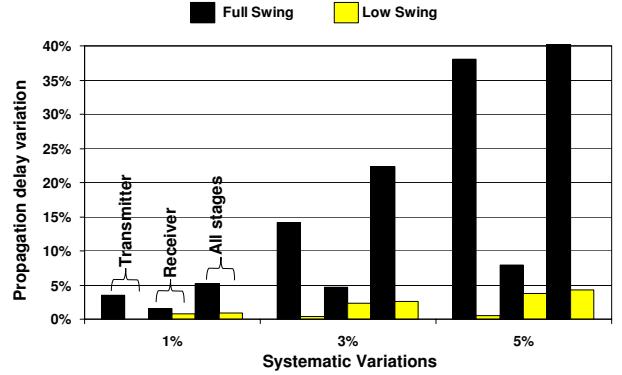


Figure 3. Sensitivity to systematic variations.

respect to process variations, while compensation techniques will be addressed in Section 5. Our focus is on within-die variations, which happen at the length scale of a die, and that can be further divided into two contributors: systematic and random. Systematic variations can be predicted prior to fabrication and exhibit space locality. In contrast, random variations are due to the inherent unpredictability of the semiconductor technology itself. In our tests we inject effective gate length variations, which have implications on the threshold voltage as well, as computed by the SPICE device models of our target library. HSPICE is used as our simulation engine. We ignore variations in wires, in agreement with current variation models (e.g., [23, 29]). Fig.3 shows the sensitivity of the signaling schemes to *systematic variations*. The sensitivity is measured as the variation-induced deviation of the clock propagation time of the receiver from the nominal value. The propagation time goes from the clock sampling edge to the 50% voltage swing of the receiver output, and its nominal value is the same for both full- and low-swing channels, since they were designed to impact the next stage of the design in the same way. Systematic variations have been applied selectively to the transmitter, to the receiver and to the whole channel, so the bars in Fig.3 should be read pairwise. It can be clearly observed that low-swing signaling proves a far more robust scheme to systematic variations. By restricting the analysis to the full-swing channel, its transmitter turns out to be the weak point of this scheme. The reason lies in the high sensitivity of the library flip-flop (i.e., the receiver) to the settling time of its input signal. This latter significantly deviates from nominal conditions when systematic variations affect the transmitter, and this explains the large degradation of the whole full-swing channel performance. In contrast, the receiver seems much more robust, and variations affecting the whole channel introduce only an incremental degradation with respect to the one caused by the transmitter. The only exception occurs for channel-wide 5% systematic variations, where nominal delay is degraded by 90% (height of the last column for full-swing is truncated to preserve the scale). This is much more than one could expect by looking at the transmitter-degraded case, but this is due to the fact that we are working close to the point where full-swing channel operation fails: in this region, delay is highly sensitive to process parameter variations. The opposite holds for the low-swing channel. The PDIFF receiver does a good job in providing a noise margin to the perturbations of its input signal induced by systematic variations in the transmitter. However, when variations affect directly the receiver, the PDIFF scheme suffers from increased switching delay. Clock propagation delay variations are much smaller for low-swing channels with respect to the full-swing ones anyway, and might more easily induce the following stage in the design to fail, since it may be impossible to leave a 90% performance degradation margin for 5% systematic variations, as required by the full-swing channel. We detected a failure of the full-swing channel when the transmitter is affected by 6/7% variations (tolerating a maximum propagation delay degradation by 90%), while the low-swing channel can keep working also under 70% systematic variations affecting the receiver, after that the channel fails. At

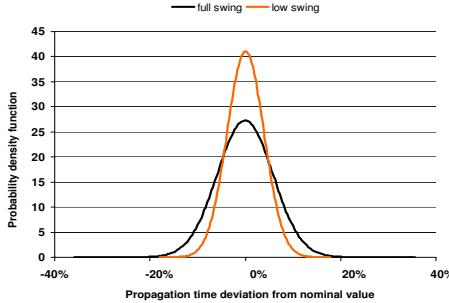


Figure 4. Sensitivity to random variations.

that time, however, propagation delay is degraded by 40%.

The sensitivity of the channels under test to random variations ($3\sigma/\mu=15\%$) is illustrated in Fig.4. Delay variability is similar in the two cases, with a slightly more tightened distribution for the low-swing channel. Again, we found the transmitter to be the most critical part of the full-swing channel, while the receiver is obviously the weak point of the low-swing channel. In fact, its pseudo-differential behaviour makes it very sensitive to random process variations, although we found only a negligible amount of malfunctioning channels with $3\sigma/\mu$ lower than 20%. This indicates that under such variations, the unbalancing of the differential branches remains within the noise margin of the receiver and correct 1/0 sampling takes place in due time. Delay variations pointed out in Fig.3 and Fig.4 indicate that compensation is apparently more challenging in full-swing channels, though the effectiveness of compensation depends not only on the entity of delay variations, but also on the sensitivity of channel delay to the different channel sub-blocks and also to the interaction among them, as illustrated hereafter.

5. Post-silicon compensation

Next, we explore the effectiveness of ABB (and forward body bias, FBB, in particular) vs ASV in bringing channel instances slowed down by process variations back within nominal performance. Compensation is applied to both the driver and the receiver for channel-wide tuning, but also selectively to individual sub-circuits to capture sensitivity of channel performance to that of these sub-circuits and eventually come up with lower-cost compensation techniques.

The ideal performance tuning range of each technique is investigated, without regarding of implementation issues, to justify an investment on the most suitable technique for each kind of communication channel later on.

5.1 Experimental framework

Our experiments encompass the compensation of a representative subset of variation scenarios. Similarly to [7, 23], worst-case systematic variations of $\pm 5\%$ of parameter nominal value are assumed and superimposed to random variations. For these latter, the $3\sigma/\mu$ of channel length distribution is varied from 10, 15 to 20%, thus giving rise to three scenarios featuring the same amount of worst-case systematic variations and an increasing parameter spread associated with random variations.

Recently, advanced modeling frameworks have been proposed to propagate variation information from the transistor compact model up to the system level, offering a correlated view on yield, timing, dynamic and static energy [28]. They also improve the traditional Monte Carlo statistical static timing analysis techniques by accounting for rare events in variability distributions. Since this paper focuses on a relatively small yet critical amount of logic, we developed an ad-hoc and simplified methodology based on Monte Carlo analysis to study the impact of systematic and random variability and how effectively it can be compensated. For each signaling scheme, variation scenario and compensation technique,

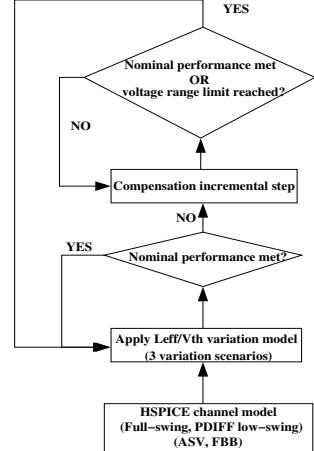


Figure 7. Framework for assessing the effectiveness of variability compensation techniques.

we perform Monte Carlo simulations with a statistically significant sample set. Each Monte Carlo run (i.e., a channel instance with different random variability injections) goes through the compensation methodology illustrated in Fig.7. At first, we check for nominal performance requirements. If met, a new instance is analyzed. If not, a compensation step is applied. In practice, if FBB is under test, an incremental reduction step of the body bias is applied so to improve performance. Similarly, the supply voltage is increased when ASV is assessed. Decrement/Increments are applied with steps of 100 mV both for ASV and FBB. This choice stems from the conclusion of previous works [21] and from considering realistic resolutions of low-cost voltage regulators. After the compensation step, performance is re-evaluated and eventually an additional compensation step is applied. The process completes when nominal performance is finally met OR when the voltage range limit is reached: 500 mV for forward body bias (to avoid turning on the source pn junction of transistors) and 200mV for ASV (for reliability and technology library constraints). Since our target 65nm manufacturing process does not provide a triple well, we apply forward body biasing only to PMOS transistors. Our analysis aims to capture whether this lower cost solution suffices for compensation purposes in on-chip communication channels. In addition, it is not possible to selectively apply ASV only to the receiver of a full-swing channel, since this would require a voltage level shifter which is not there. In contrast, such level shifter comes for free in a low-swing channel, which therefore allows PDIFF receiver selective compensation with ASV. FBB does not have any kind of constraints in any signaling scheme.

Effectiveness of a technique is expressed as the percentage of the sample set that can be brought back within nominal performance by the compensation technique under test. We denote those effectively compensated samples as *working samples*. Nominal performance means correct sampling at 1.68 GHz with clock propagation time constraints met at the output of the receivers. Moreover, the average power overhead for compensating channel instances with the highest power supply value (lowest PMOS body bias value) is measured, denoting *power efficiency* of the compensation techniques. For low-swing signaling, we also explore adaptive voltage swing as an additional and built-in compensation technique by raising the voltage swing in increments of 100mV. Finally, *systematic variations* were applied to the *whole channel* but also selectively to the receiver and to the transmitter to account for place&route effects. In fact, transmitter and receiver might be far apart from each other, thus suffering from systematic variations to a different extent, or they might be placed close to each other. In this latter case, physical parameters of the whole communication channel would be skewed by the same amount. For lack of space, we hereafter report only this latter case and the differences (if any) with the other variation scenarios are discussed in the text. We also

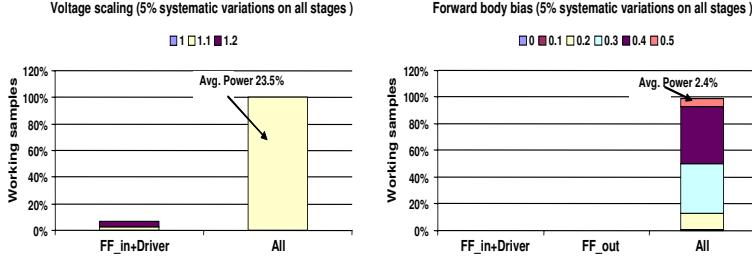


Figure 5. Working samples after compensation of *full-swing* channels. x-axis indicates the channel (sub-)circuits to which compensation has been applied.

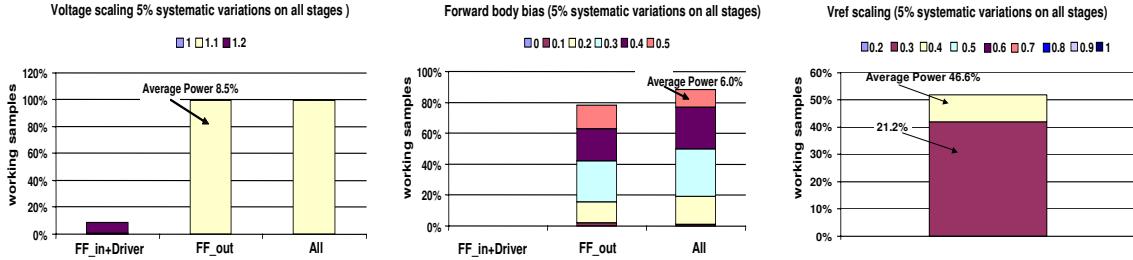


Figure 6. Working samples after compensation of *PDIFF* low-swing channels. x-axis indicates the channel (sub-)circuits to which compensation has been applied.

recall that *random variations were always applied* to the circuits of the *whole channel*, and in the first set of experiments $3\sigma/\mu$ is assumed to be 15%. See subsection 5.4 for different values. When systematic variations were injected in the entire channel (like random ones), we found almost no channel instances in the sample set working without compensation, both for full-swing and low-swing channels. So, in the experiments that follow, the entire sample set needs to be compensated.

5.2 Compensation efficiency in full-swing links

As can be observed from Fig.5, neither ASV nor FBB are able to restore functionality of all working samples by only acting upon the transmitter or (for FBB) the receiver. The compensation in this case would be totally ineffective. Variability can only be compensated by tuning all the circuits of the channel. In fact, under 5% systematic variations performance of full-swing channels is highly sensitive to the interaction between the signal provided by the transmitter and the requirements imposed by the receiver. Moreover, such variations (recall Fig.3) significantly impact both the transmitter and the receiver. As a consequence, an effective compensation can only be carried out by acting upon both modules at the same time. However, while ASV requires a single voltage step to reach 100% working samples, FBB needs its entire voltage range to achieve the same objective. Even the large variations taking place in full-swing channels can be offset by FBB in spite of its inherently weaker performance tuning capability by exploiting the sensitivity of channel performance to the circuits compensation is applied upon. The main difference lies however in the power efficiency of the techniques. When ASV raises the supply voltage to 1.1V, the communication channel instances on average exhibit a 23% power overhead with respect to the variation-free scenario. In contrast, a 500mV forward body bias incurs only an average power overhead of 2.4%, almost negligible. When we applied systematic variations only to the transmitter (flip-flop and driver), we observed that tuning only the transmitter circuits only partially solved the problem. ASV could restore about 80% of the samples, while FBB about 60% by remaining in the voltage range limits. This indicates the impact of random variations, which require a tuning of the receiver as well to restore 100% working samples. The situation is even worse when only the receiver is affected by

systematic variations: while no selective tuning of the flip-flop is feasible with ASV due to a lack of a voltage level shifter, only 20% of working samples were achieved by selective FBB. Again, the only option was to tune the entire channel, finding again the same power efficiency gap between FBB and ASV.

5.3 Compensation efficiency on low-swing links

Quite different considerations hold for variability compensation in low-swing channels. This time, ASV can be selectively applied to the receiver since the level shifter is built-in in the signaling scheme. Fig.6 clearly shows that a selective tuning of the receiver with both ASV and FBB reaches a high percentage of working samples. With just one voltage increment step applied to the output flip-flop, ASV can restore performance of all slow samples. More interestingly, the average power overhead is limited to 8.5%, much lower than in a full-swing channel. In low-swing channels, the transmitter is marginally impacted by systematic variations (recall Fig.3). At the same time, receiver performance is much less sensitive to the perturbations of the input signal than in full-swing channels. Therefore, acting upon the receiver proves an effective compensation method. Unfortunately, FBB cannot reach 100% working samples with a selective compensation at the receiver, and neither a channel-wide compensation can (90% is the best result achieved with a 500mV FBB). This is essentially due to weak performance knob represented by FBB, which is not boosted by any circuit level property in this case (for instance, no high sensitivity of channel performance to transmitter-receiver interaction). The worst-case average power overhead incurred by FBB is around 6%, comparable with that of ASV. Considering the cases where systematic process variations affect only the transmitter or the receiver, we found that FBB is not able again to reach 100% of working samples (best coverage is 90%). ASV instead works effectively. However, in all cases and for both ASV and FBB, selective compensation at the receiver turns out to be as effective as full channel compensation. Power overhead for ASV is around 7 and 8%, while for FBB is around 3%. There is a slightly higher power overhead of ASV which is the price to pay to achieve a higher compensation efficiency and, in the end, a higher yield.

Fig.6 also shows the efficiency of an intuitive compensation technique which stems from the possibility to tune the voltage

swing in the low-swing channel. Although intuitive, this technique proves highly ineffective to restore channel performance. By increasing the voltage swing from 200mV to 400mV, only 50% of the slow samples can be saved. Interestingly, by further increasing the swing proves useless, and no further improvements can be achieved, thus spending power uselessly. This is due to the fact that compensating process variations is not just an issue of speeding up signal propagation across the link, but to restore correct functionality at the transmitter and at the receiver. Only when the transmitter is impacted by systematic variations while the receiver is not, then speeding up the link with a swing of 400mV achieves 82% working samples. Compensating receiver variability proves more difficult (about 60% working samples). Another argument against reference voltage scaling is power. The measured average power overhead for the worst case compensations (those at 400mV) amounts to a significant 46%. This confirms the results of the work in [6], showing that using the voltage swing to speed up a low-swing link is highly power inefficient.

5.4 Role of random variations

When we repeated the experiments with a $3\sigma/\mu = 10\%$ and below, the minor role played by random variations translated into a better compensation efficiency of FBB in low-swing channels, since working samples were always close to 100%. The lower delay spread makes the worst-case compensation scenario affordable also for the tuning capability of FBB, so that this latter can be considered also for low-swing signaling as the impact of random variations decreases. Finally, $3\sigma/\mu$ was set to 20%. In this case, even for full-swing channels FBB could not bring all samples within nominal performance bounds, although still achieving around 95% working samples. Interestingly, in low-swing channels the effectiveness of FBB was as low as 70% working samples.

6. Conclusions

This work explores the effectiveness of ASV and FBB as post-silicon variability compensation techniques for on-chip communication channels. Our work shows that FBB is effective for tuning performance of full-swing channels with minimum power overhead. In contrast, when applied to low-swing channels, FBB proves not capable of compensating all variation patterns, since its limited performance tuning capability is not amplified by any circuit property. On the other hand, ASV can exploit the built-in voltage level shifter in low-swing channels and achieve an effective and low power-overhead compensation.

The results of this paper point out the superior robustness of low-swing channels to process variations. After considering a realistic range of systematic and random WID process variations and exploring all the possible countermeasures based on FBB and ASV, it is evident that low-swing channels (i) can better cope with systematic variations (lower delay deviations and functional correctness guaranteed over a wider range of variations), (ii) feature a lower delay spread under random variations, (iii) can be compensated with success against delay variability at a low power cost. These features add up to the reference characteristic of low-swing channels, which is their inherent low power consumption. Future work aims at analyzing how capacitive cross-talk impacts the efficiency of variability compensation techniques.

7. Acknowledgements

This work has been partially supported by the GALAXY European Project (FP7-ICT-214364) and partly by the Hipeac Network-of-Excellence (Interconnect Cluster).

8. REFERENCES

- [1] D. Sylvester and K. Keutzer, "Getting to the bottom of deep sub-micron II: A global paradigm", Proc. IEEE Int. Symp. Physical Design, pp.193-200, 1999.
- [2] M.Karlsson, M.Vesterbacka, L.Wanhammar, "Low-Swing Charge Recycle Bus Drivers", ISCA '98, pp.117-120, 1998.
- [3] Rjoub, A.; Koufopavlou, O.; "Efficient drivers, receivers and repeaters for low power CMOS bus architectures", ICECS '99, pp.789 - 794 vol.2, 1999.
- [4] Byung-Do Yang; Lee-Sup Kim; "High-speed and low-swing on-chip bus interface using threshold voltage swing driver and dual sense amplifier receiver", ESSCIRC '00, pp.105 - 108, 2000.
- [5] R. Dobkin, A. Morgenshtern, A. Kolodny, R. Ginosar "Parallel vs. serial on-chip communication", SLIP 2008, pp.43-50.
- [6] S.Medardoni, M.Lajolo, D.Bertozzi, "Variation tolerant NoC design by means of self-calibrating links", DATE '08, pp.1402-1407, 2008.
- [7] E.Humenay, D.Tarjan, K.Skadron; "Impact of process variations on multicore performance symmetry", DATE '07, pp.1653 - 1658, 2007.
- [8] Garcia, J.C.; Montiel-Nelson, J.A.; Nooshabadi, S.; "High performance bootstrapped CMOS low to high-swing level-converter for on-chip interconnects", ECCTD 2007, pp.795 - 798, 2007.
- [9] Mensink, E.; Schinkel, D.; Klumperink, E.A.M.; van Tuijl, E.; Nauta, B.; "Optimal Positions of Twists in Global On-Chip Differential Interconnects", IEEE Transactions on VLSI Systems, Volume 15, Issue 4, April 2007, Page(s):438 - 446.
- [10] Kangmin Lee; Se-Joong Lee; Hoi-Jun Yoo; "Low-power network-on-chip for high-performance SoC design", IEEE Transactions on VLSI Systems, Volume 14, Issue 2, pp.148 - 160, 2006.
- [11] A.Narasimhan, B.Srinivasaraghavan, R.Sridhar, "A low-power asymmetric source driver level converter based current-mode signaling scheme for global interconnects", Int.Conf.on VLSI Design, 4 pp., 2006.
- [12] F.Worm, P.Ienne, P.Thiran, G.De Micheli, "A robust self-calibrating transmission scheme for on-chip networks", IEEE Trans. on VLSI Systems, pp.126 - 139, 2005.
- [13] Jeong, W.; Paul, B.C.; Kaushik Roy; "Adaptive supply voltage technique for low swing interconnects", ASP-DAC 2004, pp.284 - 287, 2004.
- [14] Chang-Ki Kwon; Kwang-Myoung Rho; Kwyro Lee; "High speed and low swing interface circuits using dynamic over-driving and adaptive sensing scheme", ICVC '99, pp.388 - 391, 1999.
- [15] H.Zhang, V.George, J.M.Rabaey, "Low-swing on-chip signaling techniques: effectiveness and robustness", IEEE Trans. on VLSI Systems, pp.264-272, Vol.8, no.3, June 2000.
- [16] von Arnim, K.; Borinski, E.; Segebrecht, P.; Fiedler, H.; Brederlow, R.; Thewes, R.; Berthold, J.; Pacha, C.; "Efficiency of body biasing in 90-nm CMOS for low-power digital circuits", IEEE Journal of Solid-State Circuits, Volume 40, Issue 7, July 2005, Page(s): 1549 - 1556.
- [17] Venkatraman, V.; Anders, M.; Kaul, H.; Burleson, W.; Krishnamurthy, R.; "A Low-swing Signaling Circuit Technique for 65nm On-chip Interconnects", International SOC Conference, pp.289 - 292, 2006.
- [18] Garcia, J.C.; Montiel-Nelson, J.A.; Nooshabadi, S.; "High performance CMOS symmetric low swing to high swing converter for on-chip interconnects", IEEE Northeast Workshop on Circuits and Systems, 2007, pp.566 - 569, 2007.
- [19] Joonsung Bae; Joo-Young Kim; Hoi-Jun Yoo; "0.6pJ/b 3Gb/s/ch transceiver in 0.18 um CMOS for 10mm on-chip interconnects", ISCAS 2008, pp.2861 - 2864, 2008.
- [20] Meijer, M.; Pessolano, F.; Pineda De Gyvez, J.; "Technology exploration for adaptive power and frequency scaling in 90nm CMOS", ISLPED '04, pp.14 - 19, 2004.
- [21] Tschanz, J.; Kao, J.; Narendra, S.; Nair, R.; Antoniadis, D.; Chandrakasan, A.; Vivek De; "Adaptive body bias for reducing impacts of die-to-die and within-die parameter variations on microprocessor frequency and leakage", IEEE Journal of SSCs, pp.1396 - 1402, vol.37, no.11, 2002.
- [22] Tschanz, J.W.; Narendra, S.; Nair, R.; De, V.; "Effectiveness of adaptive supply voltage and body bias for reducing impact of parameter variations in low power and high performance microprocessors", IEEE Journal of SSCs, pp.826 - 829, vol.38, Issue 5, 2003.
- [23] Bonesi S., Bertozzi D., Benini L., Macii E., "Process variation tolerant pipeline design through a placement-aware multiple voltage island design style", DATE 2008, pp.967 - 972, 2008.
- [24] Chen, T.; Naffziger, S.; "Comparison of adaptive body bias (ABB) and adaptive supply voltage (ASV) for improving delay and leakage under the presence of process variation", IEEE Transactions on VLSI Systems, Volume 11, Issue 5, Oct. 2003, Page(s):888 - 899.
- [25] Gregg, J.; Chen, T.W.; "Post silicon power/performance optimization in the presence of process variations using individual well adaptive body biasing (IWABB)", 5th International Symposium on Quality Electronic Design, Page(s):453 - 458, 2004.
- [26] I.Hatirnaz, S.Badel, N.Pazos, Y.Leblebici, S.Murali, D.Atienza, G.De Micheli; "Early wire characterization for predictable network-on-chip global interconnects", SLIP'07, Page(s):57-64, 2007.
- [27] D.Bertozzi, L.Benini, B.Ricc, "Parametric timing and power macromodels for high level simulation of low-swing interconnects", ISLPED 2002: pp.307-312.
- [28] A.Papanicolaou et al., "At Tape-out: Can System Yield in Terms of Timing/Energy Specifications be Predicted?", IEEE Custom Integrated Circuits Conference, pp.773-778, 2007.
- [29] E.Humenay, D.Tarjan, K.Skadron; "Impact of Parameter Variations on Multi-Core Chips", Int.Workshop on Architectural Support for Gigascale Integration, 2006.
- [30] H. C.Wan et al., "Channel doping engineering of MOSFET with adaptable threshold voltage using body effect for low voltage and low power applications", Int. Symp. VLSI Technology, Systems, and Applications, 1995, pp.159-163