

Feedback Based Droop Mitigation

Salvatore Pontarelli, Marco Ottavi, Adelio Salsano
Dept. Electronic Engineering
University of Rome "Tor Vergata", Rome, Italy
{pontarelli,ottavi,salsano}@ing.uniroma2.it

Kamran Zarrineh
Advanced Micro Devices
Boxborough, MA, USA
kamran.zarrineh@amd.com

Abstract—A strong dI/dt event in a VLSI circuit can induce a temporary voltage drop and consequent malfunctioning of logic as for instance failing speed paths. This event, called power droop, usually manifests itself in at-speed scan test where a surge in switching activity (capture phase) follows a period of quiescent circuit state (shift phase). Power droop is also present during mission mode operation. However, because of the less predictable occurrence of the switching events in mission mode, usually the values of power droop measured during test are different from those measured in mission mode. To overcome the power droop problem, different mitigation techniques have been proposed. The goal of these techniques is to create a uniform current demand throughout the test. This paper proposes a feedback based droop mitigation technique which can adapt to the droop by reading the level of VDD and modifying real time the current flowing on ad-hoc droop mitigators. It is shown that the proposed solution not only can compensate for droop events occurring during test mode but also can be used as a method of mission mode droop mitigation and yield enhancement if higher power consumption is acceptable.

Keywords: droop, mitigation techniques, ATPG, power supply;

I. INTRODUCTION

Modern VLSI devices provide multiple functional modules into a single die. These modules perform diverse functions and can be activated at different times depending on the overall workload of the system, substantially varying the system power consumption over a short period of time. The occurrence of these peak power demands gives rise to transient current events also known as dI/dt events. Examples of dI/dt events are the switching from shift to capture phase in test mode, the activation of a core put in a low power state, the access to on chip memories (such as caches in microprocessors) and in general the activation of big functional blocks previously in idle state. These events can cause power droop and overshoot of up to 20% of the nominal voltage [1],[2]. Different techniques have been proposed to measure and mitigate the occurrence of power droop during test, some techniques propose a modification of the test vector [3] other propose to measure the magnitude of the power droop event and the activation of inducers to counter balance and mitigate the initial effect [1]. The droop problems in mission mode are also related to the need of detecting when the droop will occur. Some prediction methods based on architectures analysis [4] and/or signatures [5] have been proposed. Detecting droop in mission mode can also be done by using suitable detectors like proposed in [6],[7],[8]. After prediction and/or detection of droop event, countermeasures should be applied such as

presented in [8],[9]. Also frequency throttling can be efficiently used to mitigate the effects of the droop event [10]. As discussed in [5] a countermeasure against droop is more effective the earlier its activation occurs. Droop prediction is therefore a very promising method to mitigate power droop events, however it is mostly unfeasible to predict with a high accuracy the exact time and magnitude of a droop event, since its occurrence dependson too many concurrent causes.

Two methods for testing the random logic in high performance and complex digital designs have been proposed: 1) structural test and 2) functional test. A typical structural test could have the following series of events: a series of slow occurring shift clock cycles, e.g., 100MHZ $\rightarrow$ a quiet time $\rightarrow$ two or more 2GHZ+ clock pulses. The topology of the structural test increases the probability of power droops that could cause perfectly good designs failing the test.

The functional tests share the very same behavior as the mission mode of operation of the design, thus, techniques and methods developed for the mitigation of power droop in the mission mode are also applicable to the functional tests. Although, our proposed technique could be extended to cover the functional test, we mainly concentrate on the structural test. This paper proposes a novel mitigation technique based on a feedback controlled mitigator transistor that is able to quickly react to the droop event providing a dI/dt of opposite sign with respect to the dI/dt event that causes the droop. The proposed technique easily solves the issues related to the magnitude of the droop compensation while also partially resolving the need of an accurate definition of the droop activation event. In fact, as it will be shown later, also an approximate indication of the droop event would allow enabling the proposed mitigation technique.

The rest of the paper is organized as follows: section 2 describes feedback controlled transistor based mitigation technique and shows its performances with respect to other techniques and by varying some parameters. Section 3 shows the application of the proposed technique both in test mode and mission. Finally conclusions are drawn in section 4.

II. PROPOSED MITIGATION TECHNIQUE

The model we used to simulate a droop event is composed of an RLC network, which is triggered by a pulse generated by an NMOS, which simulates the occurrence of many concurrent switching events. The scheme of the proposed model is presented in Fig. 1a. The VDD wire is the power line affected by the droop, the trigger signal activates the droop event and

978-3-9810801-7-9/DATE11/©2011 EDAA

This research was partially funded by the Italian Ministry for University and Research; Program "Incentivazione alla mobilità di studiosi stranieri e italiani residenti all'estero", D.M. n.96, 23.04.2001

the external VDD is the theoretical value of VDD. Fig. 1b shows a droop event triggered by a sequence of capture events as modeled by the circuit in Fig. 1a.




Fig.1: a) Circuit modeling power droop b) Waveform of a droop event

The main problem of the currently used techniques for voltage stabilization is that they are slow or cannot be easily integrated by using the standard VLSI technology process. To overcome these limitations, [1] has proposed the use of a mitigator transistor that is used to provide a current variation of opposite sign with respect to the one produced by the concurrent activation of current sinking circuit blocks of the system under control. The mitigator transistor can be directly inserted between the VDD wire to control and the ground on the chip. During the activation of these current sinking blocks, the mitigator should reduce its current consumption, in such a way that the overall current consumption seen by the external VDD power supply changes more gradually than the current flowing in the activated blocks alone. When the blocks under control draws less current, the mitigator increases the amount of current flowing into it, always providing a di/dt of opposite sign with respect to the rest of the circuit. This technique has two positive aspects: it can be easily integrated because it is based on a wide transistor providing additional current consumption, and is very fast. However this technique has also two main issues: adding a mitigator actually increases power consumption, and it needs an external signal to activate the mitigator before the droop event. [1] proposes to activate these devices by a programmable and flexible logic controlled by scan chains. The values that control when and how much the mitigator must be activated, have to be determined by running an initial analysis of the droop effects observed on the circuit both in testing and mission mode. The efficient use of the mitigator therefore requires a complex set of measurement campaigns in order to be effective. The mitigation techniques should be able to compensate the droop event of the testing mode in an accurate way. In fact, an overcorrection of the droop event could give some false negative testing results for chips that are actually unable to satisfy the timing requirements in mission mode. When overcorrection occurs, chips pass the testing phase only because the droop mitigation provides a higher voltage value with respect to the realistic one available for those chips in mission mode. On the other hand, also an under-correction of the droop event is undesirable since they would discard some chips that in mission mode are completely able to satisfy the requirements.

Summarizing, only a precise knowledge of the chip behavior with respect to the droop events during the test allows correctly

programming the control logic to activate/deactivate the mitigator transistors. Finally, the mitigation technique presented in [1] increases the test time by requiring a two pass test pattern application: 1) test pattern to determine the time and magnitude of the voltage droop and 2) test pattern with the activated mitigation technique.




Fig.2: Feedback based droop mitigation

The solution proposed by this paper is to bypass this characterization by adding a feedback logic that controls the inducers allowing them to follow real time the droop event and mitigate it. A simple way to implement this feedback is shown in Fig. 2 where the gate of an NMOS mitigator is controlled by the same voltage affected by the droop event.

When VDD is at its nominal value M2 is conducting a constant current. When VDD lowers as a consequence of a droop event the mitigator becomes less conductive thus reducing the current flowing through it. On a di/dt point of view the sum of the currents flowing through the inductance L1 has a lower derivative with respect to the derivative of the circuit in Fig. 1a. Therefore the added mitigator works effectively as a feedback system on the flow of current drained from the power supply. Differently from [1] this solution does not require an external input for the control of the inducer/mitigator, however, since the control signal could be of low magnitude the size of the inducer/mitigator has to be larger than the one proposed in [1]. This does not pose a serious issue since the mitigator will be turned off completely during the mission mode, thus, there will be no power overhead to this technique only a small area overhead increase. The feedback transistor simplifies the characterization phase discussed previously with respect to two issues: first it is not anymore necessary to know exactly when the droop events occur, because before the droop events the feedback transistor draws a fixed amount of current, while as soon as the droop is activated by an external event, the




Fig. 3: comparison of droop mitigation techniques

feedback transistor starts to work. Also when the droop event terminates, the feedback transistor gradually returns to draw a constant current. Moreover, also the magnitude of the droop event, and consequently of its countermeasure, is taken into account by the feedback, avoiding the cases in which the droop event is overcompensated or in undercompensated. To check the efficiency of the proposed mitigation approach a set of simulations has been performed by using PTM models in order to validate the proposed approach on a 45 nm technology [11]. According to the specifications of the PTM models, the supply voltage for all the following simulation is 0.9 V.

Figure 3 shows a comparison of the proposed droop mitigation technique with respect to [1] and to the un-mitigated droop event. While the [1] circuit requires the droop inductor to be triggered concurrently to the foreseen droop event (i.e. when the first capture is started) with the proposed technique this constraint is not necessary since the mitigation technique is triggered by the droop event itself. In the simulations it is assumed that both the transistors used by [1] and the feedback circuit have the same channel width and that the deactivation of the inducer in the case of [1] mitigation technique occurs at the same time of the trigger signal that causes the droop event. The waveform shows that both methods mitigate the droop events, even if the externally controlled inducer gives better results in terms of the minimum value of VDD. In particular, the value of VDD for the unmitigated, externally mitigated and the proposed solution are reported in table I.

TABLE I:
MINIMUM VDD OF UNMITIGATED/MITIGATED DROOP

VDD values		
<i>Unmitigated</i>	<i>[1]</i>	<i>proposed</i>
0.6634 V	0.7337 V	0.7081 V

As it will be shown later, the minimum VDD is dependent on the width of the transistor, therefore by correctly dimensioning the feedback mitigator, the proposed method can obtain the same values of the externally controlled mitigator.

TABLE II:
MIN. VDD OF UNMITIGATED/MITIGATED DROOP

Time shift (ns)	VDD (V)	Time Shift (ns)	VDD (V)
-5	0.6761	+1	0.7306
-4	0.6934	+2	0.7211
-3	0.7091	+3	0.7062
-2	0.7221	+4	0.6873
-1	0.7307	+5	0.6706
0	0.7337	-	-

It is interesting to observe what happens when the signal that activates the externally controlled mitigator occurs not exactly at the same time of the droop event. While the exact time of the droop event is known in test mode, in mission mode it can be only predicted with a large uncertainty. The deactivation of the mitigator has been shifted with respect to the activation of the droop event, in an interval from -5ns to +5ns and in Table 2 the minimum value achieved by VDD are reported. The results show that with an uncertainty of ± 3 ns the proposed technique gives the same results of the external controlled one, whereas if

the uncertainty of activation of the inducer is as high as ± 5 ns the externally controlled mitigator gives the same results of the unmitigated case, therefore becoming useless.

Different simulations have been performed varying the transistor width W to evaluate the minimum VDD during the droop event and length in time of the droop event itself. The drooplength has been defined as the time interval during which the VDD is below 10% of the nominal voltage value.

Figure 4 shows a series of simulations in which the width of the mitigating transistor varies between 300nm and 1 μ m. It can be seen that with an increase of W the mitigation effect is stronger and that the overshoot quickly disappears. The obtained results are also summarized in Table III. The measure of the droop length can be useful supposing that the system under control can operate correctly until the VDD variation is less than 10%




Fig. 4: comparison of droop mitigation at different W

and can be suspended by using frequency throttling techniques when the VDD variation is more than 10%. Therefore reducing the droop duration allows reducing the interval of time in which the frequency throttling has to occur.

TABLE III:
MIN. VDD AND DROOP LENGTH WITH DIFFERENT W

W (nm)	Vdd _{min} (V)	T (ns)
300	7.081	7,9
600	7.415	7,1
800	7.586	6,3
1000	7.726	5,5

III. APPLICATIONS OF THE PROPOSED TECHNIQUE

The proposed droop mitigation technique is designed primarily to be used during the test mode. However, the scheme could be extended to mission mode if it is possible to estimate even with a certain level of uncertainty the occurrence of droop events. These events usually correspond to the activation of some functional blocks of a system, which have high current requirements, such as the activation of second and third level caches, the passage of the system (or part of it) from a power state to another, use of complex co-processing elements (e.g. floating point operation). During test mode it is assumed that

power consumption is not a main concern therefore the proposed mitigation technique can be active throughout test. Fig. 5 shows a simple scheme of the circuit enabling the mitigation transistor. During the test initialization phase, the JTAG register shown in figure is loaded with the '1' value to enable the droop mitigation during test. This bit enables all the droop mitigator spread across the chip. The activation bit closes the M2 transistor thus enabling current flow through the feedback-controlled mitigator.




Fig. 5: Circuit enabling mitigation in test mode

The activation of the droop mitigator M1 implies the occurrence of a wide dI/dt event, which does not cause any malfunction because in the initialization phase the actual test process has not already been started. At the end of the testing phase, the JTAG register is reset, thus disabling the mitigator. Also in this case, the droop event induced by the mitigator itself is not a concern, because the system is at the end of the testing phase, but not yet in functional mode.

If the droop mitigator is used also in mission mode, the mitigation circuit includes the possibility of being activated gradually by using a suitable control circuitry. In this case the activation/deactivation of the mitigator occurs during the critical phases of test and mission modes. Therefore the dI/dt event produced by the activation/deactivation of the mitigator should be reduced as much as possible. This can be achieved by extending the time to activate/deactivate the mitigator.

In Fig. 6 three narrow parallel transistors have replaced the activation transistor. When all of the transistors are closed, their resistance is the same of transistor M2 of Fig. 6. Since the activation of these transistors happens serially, the behavior of this circuit is similar to the one in Fig. 6, but with a difference in the magnitude of the dI/dt event. The Δt in this case is three times the Δt of the activation of the circuit in Fig. 6. Therefore, the corresponding droop event caused by the activation/deactivation of the mitigator can be easily shaped by selecting an appropriate number of parallel transistors allowing the gradual activation of the M1 transistor.

In the test mode, the capture event is preceded by hundreds of clock cycles during the shift phase of a new test vector.




Fig. 6: Circuit for gradual activation of mitigator

Therefore, the last clock cycles of the shift phase can be used to gradually activate the mitigator feedback transistor. Since the occurrence of a droop event can be easily predicted during the test, the mitigator can be activated and deactivated any time is needed, thus reducing the overall power consumption also in test mode. Differently from other droop compensation techniques, if the activation of the droop compensation occurs too early, this does not compromise the efficiency of the proposed method but only increases the amount of current drawn by the mitigator. Finally, the feedback-controlled transistor allows solving the issues related to definition of the droop compensation magnitude, avoiding overcompensation.

IV. CONCLUSIONS

This paper proposed a novel mitigation technique based on a feedback controlled mitigator transistor. The mitigation technique is able to easily modulate the compensation with respect to the magnitude of dI/dt event that cause the droop, and can start as soon as the droop event occurs. The simulations performed show the efficiency of the proposed method with respect to the minimum VDD value and the droop duration achievable if the mitigation technique is used. Moreover it is shown that, differently from other techniques, also if the mitigator is activated before the droop event, its efficiency is not compromised. The characteristics of the proposed technique allow using this compensation both in test and in mission mode. Finally a modification of the circuit providing a gradual activation of the mitigator has been presented. This extension allows to gradually activate/deactivate the mitigator, limiting the droop induced events activator by the mitigator itself.

REFERENCES

- [1] Rex Petersen, et al. "Voltage Transient Detection and Induction for Debug and Test" ITC 2009
- [2] Polian, I.; Czutro, A.; Kundu, S.; Becker, B.; "Power Droop Testing", Design & Test of Computers, IEEE, Vol. 24, n. 3
- [3] P. Pant, et al., "Understanding Power Supply Droop During At-Speed Scan Testing", VLSI Test Symposium, 2009.
- [4] V.J.Reddi, et al. "Predicting Voltage Droops Using Recurring Program and Microarchitectural Event Activity," IEEE Micro, vol. 30, no. 1.
- [5] V.J.Reddi, et al. , "Voltage emergency prediction: Using signatures to reduce operating margins", High Performance Computer Architecture, 2009. HPCA 2009. IEEE 15th International Symposium on,
- [6] A. Muhtaroglu, et al., "On-Die Droop Detector for Analog Sensing of Power Supply Noise", IEEE Journal of Solid-State Circuits, Pg. 651-660, Vol. 39, No. 4, April 2004.
- [7] E. Alon, V. Stojanovic, Mark A. Horowitz,, "Circuits and Techniques for High-Resolution Measurement of On-Chip Power Supply Noise", IEEE Journal of Solid-State Circuits, pp. 820-828, Vol. 40, No. 4.
- [8] M. Holtz, S. Narsimhan and S. Bhuniya, "On-die CMOS Voltage Droop Detection and Dynamic Compensation," in Proc. Great Lakes Symposium on VLSI (GLSVLSI'08), pp. 35-40, 2008
- [9] R. Joseph, D. Brooks, and M. Martonosi. "Control Techniques to Eliminate Voltage Emergencies in High-Performance Processors," in Proc. International Symposium on High-Performance Computer Architecture (HPCA-9), pp. 79-90, 2003.
- [10] J. Tschanz, et al. "Adaptive Frequency and Biasing Techniques for Tolerance to Dynamic Temperature-Voltage Variations and Aging", IEEE International Solid-State Circuits Conference, 2007. ISSCC 2007.
- [11] W. Zhao, Y. Cao, "New generation of Predictive Technology Model for sub-45nm early design exploration," IEEE Transactions on Electron Devices, vol. 53, no. 11, pp. 2816-2823, November 2006.