Design and Failure Analysis of Logic-Compatible Multilevel Gain-Cell-Based DRAM for Fault-Tolerant VLSI Systems

Pascal Meinerzhagen[1], Onur Andıç[2], Jürg Treichler[2], and Andreas Burg[1]
[1]Telecommunications Circuits Laboratory, EPFL, Lausanne, Switzerland
[2]Integrated Systems Laboratory, ETHZ, Zurich, Switzerland
pascal.meinerzhagen@epfl.ch, onur.andic@espros.ch, treichle@iis.ee.ethz.ch,
andreas.burg@epfl.ch

ABSTRACT

This paper considers the problem of increasing the storage density in fault-tolerant VLSI systems which require only limited data retention times. To this end, the concept of storing many bits per memory cell is applied to area-efficient and fully logic-compatible gain-cell-based dynamic memories. A memory macro in 90-nm CMOS technology including multilevel write and read circuits is proposed and analyzed with respect to its read failure probability due to within-die process variations by means of Monte Carlo simulations.

Categories and Subject Descriptors: B.3.1 [**Memory Structures**]: Semiconductor Memories—*Dynamic memory (DRAM)*; B.8.1 [**Performance and Reliability**]: Reliability, Testing, and Fault-Tolerance

General Terms: Design, Reliability

Keywords: Embedded memory, high density, multilevel storage, gain cell, process variations, read failure

1. INTRODUCTION

Embedded memories consume an increasingly dominant part of the overall area of application specific integrated circuits (ASICs) and VLSI systems-on-chip (SoCs) [1]. Hence, increasing the storage density of embedded storage macros is key to achieving area efficiency in future ASIC and VLSI SoC designs.

Due to increasing process variations and higher defect levels it becomes challenging to design reliable VLSI systems in modern sub-100-nm deep-submicron (DSM) CMOS technologies. Shifting the design paradigm from the traditional requirement of 100 % correctness for devices and interconnects to fault-tolerant VLSI systems [3, 7] can result in smaller area, lower power consumption, and reduced cost. Random within-die process variations such as line edge roughness (LER) and random dopant fluctuations (RDFs) affect memory cells more than logic since the transistors are of minimum size in memory cells for higher density requirements [7]. It is thus possible to trade the reliability of embedded memories in DSM technologies for less area. There is also a tradeoff between retention time and area, as will be expatiated on below. As an example for the use

of unreliable memories in fault-tolerant systems, we mention the work in [14], where the effect of unreliable storage of log-likelihood ratios on the performance of wireless communication transceivers is investigated. The system under consideration in [14] requires retention times below $10\,\mu s$ and it is shown that even error rates up to a few percent can be tolerated. Such encouraging results motivate the consideration of area-efficient albeit unreliable storage devices for fault-tolerant VLSI systems.

The most common type of embedded memory used in today's VLSI systems, regardless of the actually required data retention time, are static random access memory (SRAM) macrocells. These macrocells are compatible with standard digital CMOS technologies, but they suffer from a relatively large area of the 1-bit storage cell. In order to increase the storage density, embedded dynamic random access memory (eDRAM) macrocells are an interesting alternative to SRAMs. However, conventional 1-transistor-1-capacitor (1T-1C) eDRAMs typically require special process options and 4 to 6 extra masks to accommodate high-density 3D capacitors [10], which adds cost to standard digital CMOS technologies. From a functional perspective, dynamic memories usually require refresh cycles that are costly in terms of bandwidth and power consumption. However, in DSP systems requiring only short data retention times the refresh cycles may be skipped and in some cases retention time can even be compromised further for the benefit of higher storage density.

In order to further increase the storage density of eDRAM at the cost of compromised reliability and reduced data retention times, storing more than one bit per cell has been proposed in various multilevel DRAM (MLDRAM) designs [2, 6, 8, 15] which are all based on the conventional 1T1C storage cell. The noise margin in an n-level MLDRAM is reduced by a factor of $1/(n-1)$ compared to the noise margin in a conventional single-bit-per-cell (two-level) DRAM [5] which implies that MLDRAMs are less reliable.

Gain cells (GCs) are an alternative to SRAMs and to 1T1C-based eDRAMs. They can be smaller than the SRAM 1-bit storage cell and, as opposed to conventional 1T1C eDRAMs, they can be fully compatible with standard digital CMOS technologies and they allow for a non-destructive read operation. Various GC-based DRAMs with different basic cell structures have been proposed during the last decade [4, 9, 12, 13, 16]. However, to the best of our knowledge, the possibility of storing multiple levels (more than one bit) per GC has not been exploited yet.

Contribution: In this paper, an 8-kbit multilevel GC-based DRAM architecture storing 2 bits per GC is presented (Sec. 2) and it is shown how the storage density is increased at the expense of reduced reliability and retention time compared to commercially available SRAM macrocells (Sec. 4).

With a view toward fault tolerant VLSI signal processing systems, we investigate the dependency of the read failure probability on the *time upon write*, i.e., the time that passes between writing and reading back the data from the storage array (Sec. 3). The results serve as a link to the area of fault-tolerant system performance analysis and design where the knowledge about the degree of data integrity for a given retention time can be taken into account.

2. MULTILEVEL GAIN CELL MEMORY

The basic idea behind GC-based memories is to store data in the form of charge on a capacitive storage node (SN) formed by the gate terminal of a dedicated storage transistor (ST). In *multilevel* GC-based memories, many different voltage levels must be generated and transferred to the SN during the write operation. During the read operation, the transconductance gain of the ST is exploited to yield different sensing currents which can be compared to reference currents to yield a decision on the information stored in the cell. In summary, a multilevel GC-based memory comprises the following key components: an array of storage cells, a circuit for the generation of storage and reference levels, and a read circuit.

2.1 Gain Cell

We shall first compare the GCs of previously proposed logic-compatible *single-bit-per-cell* macro memories [4, 9, 11–13, 16] to find the GC topology that is best suited for *multilevel* operation. Subsequently, different transistor configurations are discussed to optimize the area of the storage array while maintaining good reliability.

2.1.1 Comparison of GC Topologies

All previously proposed single-bit-per-cell GC topologies have a ST whose gate terminal is connected to the SN, a write port consisting of a write word line (WWL) and a write bit line (WBL) terminal, and a read port consisting of a read word line (RWL) and a read bit line (RBL) terminal. Except for the conventional 3T GC [11] all GCs also have a coupling capacitor between the SN and the RWL either in form of an explicit MOS capacitor or solely in form of the gate-over-diffusion overlap capacitor of the ST. This coupling capacitor is used to boost the SN voltage once the RWL is pulled high during read operation to yield a larger sensing current for a faster read operation.

The 2T1MOSCAP gain cell requires a large MOS capacitor and also two additional process steps to implement very high threshold voltage transistors [9]. The 2-transistor-1-diode (2T1D) GC [12] uses the ST also as read transistor (RT), thereby saving silicon area. However, the number of words which can be connected to the same RBL is seriously limited, as the sum of the leakage currents drawn from the RBL by unselected cells quickly masks the sensing current of the selected cell to such an extent that the read operation fails. This problem is mitigated in the 3-transistor-1-diode (3T1D) GC [13] by adding a separate RT to the cell, at the price of a larger silicon area which is now close to the one of a 6-transistor-SRAM cell. The 2PMOS GC [16] is very compact, but has a low retention time and complex peripheral circuitry is necessary to overcome the masking problem associated with the missing read transistor. Adding a third PMOS transistor results in the circuit considered in [4].

2.1.2 Extension to Multilevel Operation

In *single-bit-per-cell* storage arrays only an on- and an off-state of the ST, corresponding to two intervals of the SN voltage, must be distinguished. In *multilevel* GCs, the drain current of the ST is modulated by means of its gate voltage to distinguish between multiple levels during the read
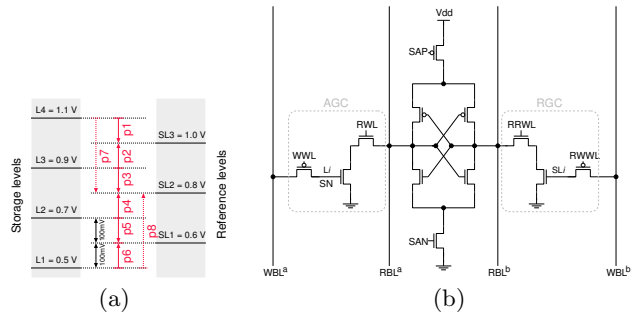

Figure 1: (a) Allocation of storage and reference levels. (b) SA connected to AGC and RGC.

operation. To this end, the dynamic range of the voltage on the SN is partitioned into multiple non-overlapping regions corresponding to the individual symbols stored in a cell. This more fine-grained partitioning of the available dynamic range of the SN voltage increases the sensitivity of the GC to leakage, which causes the SN voltage to drift, and therefore limits the retention time of the circuit. Furthermore, for multilevel sensing smaller differences in sensing current must remain distinguishable compared to single-bit-per-cell storage arrays.

For our implementation, we chose the 3T GC (see gray box in Fig. 1(b)) for reasons of its area efficiency compared to most of the other topologies. The additional, separate RT compared to the more area-efficient 2T GC was chosen to avoid the masking issues during read operation that were already critical in previous single-bit-per-cell implementations [16].

The 3T GC topology can be implemented using different combinations of pMOS and nMOS devices. Clearly, an all-pMOS or an all-nMOS configuration yields the most compact cell layout. Unfortunately, the drawback of such a configuration is that the gate voltage of the write transistor (WT) must be boosted to be able to transmit the maximum available dynamic range for which the ST is turned on to the SN during write operation in order to maximize the available margin between different levels. This implies the use of level shifters and a second power supply to generate the boosted WWL voltage. Furthermore, the correct functioning of the memory might be difficult to guarantee due to excessive gate tunneling and the long-term reliability might be compromised without a proper power-up sequence which ensures that the maximum voltage between the terminals of the WT does never exceed the specifications of the technology.

To avoid the above described problems, we chose a configuration in which the WT is implemented as pMOS transistor while the ST and RT are implemented as nMOS transistors (vice versa would also be possible). The drawback of this solution is the area overhead required for the spacing between nMOS and pMOS devices. In our *mixed* GC configuration, this overhead is minimized by sharing the n-well on 3 sides between neighboring cells. Since the cell area is mostly limited by the contacts, the overall cell area increases only by a very small amount. As for the entire memory macro, requiring neither level shifters nor the generation of an additional boosted supply voltage, our mixed GC configuration results in much smaller overall area than the nMOS- or pMOS-only configuration.

2.2 Level Generation

Fig. 1(a) shows the 4 storage levels on the left-hand side and the 3 reference levels on the right-hand side which must be generated for storing and reading back 2 bits per cell.

In order to locally generate these levels within the macro,

we follow the area-efficient approach proposed in [5, 8] by using charge sharing between bitline segments (subbitlines) which are precharged to either $0\,\mathrm{V}$ or to the supply voltage (V_{DD}) and then shorted together. Fig. 2 shows one column of the memory macro and highlights the switches connecting two subbitlines.

2.3 Multilevel Sensing

As shown in Fig. 2, each column of the macro memory contains not only the actual GCs, but also a reference GC (RGC). The sense operation starts by writing a reference level to such a RGC in an unselected column of the storage array. Subsequently, the current drawn by the active GC (AGC), i.e., the GC being read, is compared to the current drawn by the RGC. To distinguish between multiple levels, one storage level must be compared to several reference levels. These comparisons can be done either sequentially [8] or in parallel [6]. For sequential 4-level sensing implementing a successive approximation, one storage level must be compared to two reference levels. As opposed to DRAMs based on the conventional 1T1C cell, a storage level can easily be sensed multiple times in GC-based memories due to the non-destructive read access to the GCs. Using sequential rather than parallel multilevel sensing leverages this advantage to keep the area of the readout circuits small.

Fig. 1(b) shows the sense amplifier (SA) together with the AGC and the RGC. After storing the mid-range reference level (SL2 in Fig. 1(a)) to the RGC, the RBL of the active and the reference column are precharged to V_{DD} and equalized by the bit line equalizer shown on the right-hand side of Fig. 2. The RWLs associated with the AGC and the RGC are then enabled at the same time which causes the RBLs to be discharged. Since the voltage levels stored in the GCs are different, the two RBLs are discharged unequally fast. The SA is triggered by the control logic after a short delay that is chosen long enough to allow for the development of a sufficient voltage difference between the two RBLs. The sense operation is then repeated with a second reference level that is chosen depending on the outcome of the first comparison.

3. RELIABILITY ANALYSIS

The dynamic storage mechanism combined with the reduced margin between the levels representing different symbols for the multilevel storage capability compromise the integrity of the data stored in the memory array. In the following, we presume a fault-tolerant application that can tolerate unreliable, but still mostly functional circuit behavior and we analyze the reliability of the proposed storage array for different operating conditions and process corners.

3.1 Read Failure Analysis

The two main reasons for not being able to read back the content of a memory cell correctly in the described storage array are 1) *within-die* process parameter variations that give rise to mismatch between the transistors on the active branch and the reference branch of the readout circuit and 2) the sum of leakage components from and to the SN that alters the voltage on the SN. The second effect causes a shift of the SN voltage in the direction of one of the neighboring levels which reduces the sense margin that is available to compensate for process parameter variations. Hence, the percentage of errors due to process parameter variations depends on the *time upon write* which defines the time between the read operation and the last write operation to the corresponding GC.

3.1.1 Impact of Within-Die Process Variations

We shall first investigate the impact of process parameter variations alone, without also explicitly considering the de-

pendency of the error rate on the time upon write. To this end, we consider the voltage difference ΔV between the SNs of the AGC and the RGC as a parameter that we can set to emulate the voltage drift of the SN. A read failure can occur due to mismatch between the corresponding transistors in the active and the reference branches of the GCs and of the SA. The smaller ΔV, the higher the sensitivity of the sensing scheme to mismatch. For the GCs, the corresponding STs as well as the corresponding RTs should match, while in the SAs the nMOS (pMOS) transistors in the cross-coupled inverter pair should match (see Fig. 1(b)). Transistors in the GCs are of minimum size and can be far apart. They can therefore hardly be matched and process parameters must be considered to be independent and identically distributed (i.i.d.). The opposite is true for transistors in the SA which can be placed in close proximity to each other and can be sized generously to improve matching. Nevertheless, i.i.d. process variations between the AGCs and the RGC as well as within the SAs are considered for all following analyses.

We evaluate the failure probabilities using Monte Carlo circuit simulations with back-annotation of all relevant layout parasitics. Depending on the level being stored in the AGC and depending on the state of the successive approximation algorithm (first or second comparison), 8 sense operations, labeled p1 ... p8, are distinguished as shown in Fig. 1(a). The sense operations p7 and p8 have a much greater margin than the other sense operations (p1 to p6). We can therefore limit the analysis of the read failure probability to the sense operations p1 to p6. Fig. 3(a) shows the corresponding empirical failure probabilities p_{fail} for 1000 within-die process parameter realizations under *worst*-case conditions, corresponding to the fast-fast process corner at $85\,^{\circ}\mathrm{C}$. As expected, the read failure probabilities increase as the margin ΔV decreases and reach $50\,\%$ for $\Delta V = 0\,\mathrm{V}$. We also observe that the failure probabilities depend mostly on ΔV and not much on the absolute SN voltage levels and are thus very similar for the six relevant sense operations.

3.1.2 Impact of Time Upon Write t_{w}

As discussed previously, ΔV for a particular sense operation can change over time due to leakage from and to the SN. This effect is negligible for the RGC which is set immediately before the read operation, but the time upon write t_{w} needs to be taken into account to determine the SN voltage of the AGC during the read operation.

Fig. 3(b) shows the sensing failure probabilities $p_{\mathrm{fail}}(t_{\mathrm{w}})$ as a function of t_{w}, again obtained through Monte Carlo simulations, for the fast-fast process corner at $85\,^{\circ}\mathrm{C}$. For each sense operation we have constructed a worst-case scenario that keeps the WBL constantly at a level that maximizes the subthreshold current of the WT pulling the SN voltage of the AGC toward the reference level of the respective sense operation. We observe from Fig. 3(b) that the sense operations p1, p3, and p5 are less likely to fail than p2, p4, and p6. The reason for this difference is that for the more reliable sense operations (p1, p3, and p5), the gate-induced drain leakage (GIDL) current of the WT charges the SN, while the subthreshold current of the WT discharges the SN. For the less reliable sense operations (p2, p4, and p6) both the GIDL current and the subthreshold current of the WT charge the SN. The worst situation occurs for p6 due to the largest drain-to-source voltage of the WT.

For practical systems, the above worst-case assumption on the state of the WBL is highly unrealistic. In fact, during the idle state of the memory, the voltage on the WBL can be controlled and can be kept in the middle of its dynamic range. As can be seen in Fig. 3(c), $p_{\mathrm{fail}}(t_{\mathrm{w}})$ decreases significantly under this new assumption, as the subthreshold
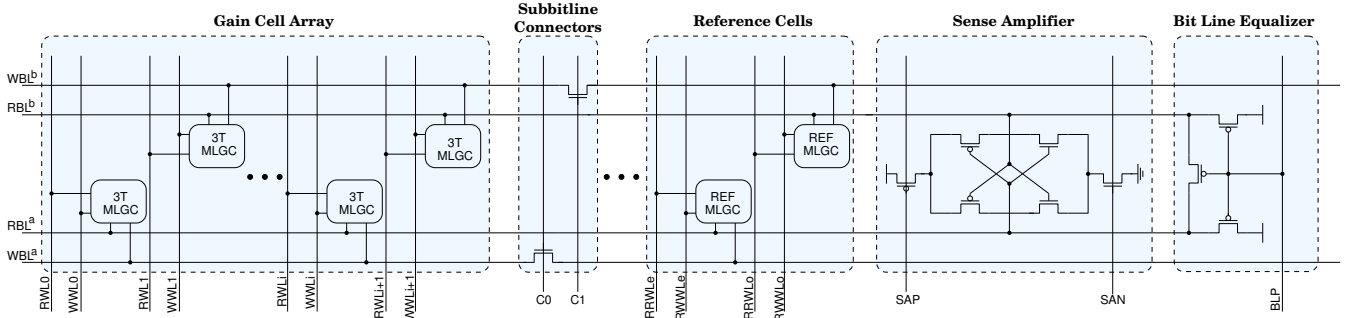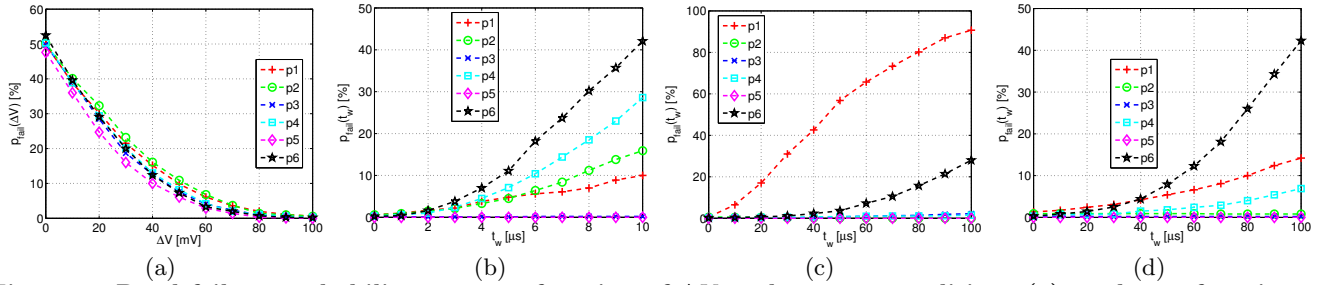
Figure 2: Macro memory architecture.

Figure 3: Read failure probability p_fail as a function of ΔV under *worst* conditions (a), and as a function of the time upon write t_w under *worst* conditions (b), *bad* conditions (c), and *typical* conditions (d).

conduction of the WT is smaller. Fig. 3(c) also shows that now the highest failure probabilities occur for the sense operations p1 and p6 due to the largest drain-to-source voltage values of the WT. p6 has a smaller failure probability than p1 as the pMOS WT has higher gate-to-source and gate-to-drain voltages and thus a smaller subthreshold current.

Keeping the same assumption on the WBL state, and for the typical-typical process corner at $25\,°C$, the maximum read failure probability among all possible sense operations $10\,\mu s$ ($50\,\mu s$) after writting is $1.7\,\%$ ($7.9\,\%$), as shown in Fig. 3(d).

4. IMPLEMENTATION RESULTS

The implemented macro memory has a storage capacity of $8\,192$ bits. With an area of $86\times138.1\,\mu m^2 = 11\,877\,\mu m^2$, the proposed 4-level GC-based macro memory is only $54.8\,\%$ the size of a corresponding commercially available single-port SRAM macrocell ($152.6\,\times\,141.9\,\mu m^2\,=\,21\,654\,\mu m^2$) with the same storage capacity, even though the SRAM macrocell contains smaller than minimum-size features (e.g., narrower contacts) and also violates other design rules (e.g., minimum diffusion enclosure of contact, minimum poly to diffusion spacing) for higher density.

5. CONCLUSIONS

In this paper, the concept of storing many bits per basic memory cell has been applied to fully logic-compatible gain cells, in order to trade reliability and retention time for higher storage density in future fault-tolerant VLSI systems in deep-submicron CMOS technologies. An 8-kbit macro memory including multilevel write and read circuits was presented and analyzed regarding its failure mechanisms. The read failure probability at a given time upon write was shown to depend quite heavily on the state of the write bit lines and is significantly decreased if the write bit lines are kept in the middle rather than on either bound of their dynamic range during the idle state of the memory. Under typical operating conditions, the maximum failure probability among all possible sense operations $10\,\mu s$ ($50\,\mu s$) after writting is less than $2\,\%$ ($8\,\%$), which can be tolerated by some fault-tolerant VLSI systems. The area of the proposed macro memory is only $54\,\%$ of the area of a commercially available single-port SRAM macrocell of equal storage capacity.

6. ACKNOWLEDGMENTS

7. REFERENCES

[1] ITRS 2009. http://www.itrs.net.
[2] M. Aoki et al. A 16-Level/Cell Dynamic Memory. *JSSC*, 22(2):297–299, 1987.
[3] M. Breuer. Let's Think Analog. In *Proc. ISVLSI*, pages 2–5, 2005.
[4] K. Chun et al. A Sub-0.9 V Logic-compatible Embedded DRAM with Boosted 3T Gain Cell, Regulated Bit-line Write Scheme and PVT-tracking Read Reference Bias. In *Proc. VLSIC*, pages 134–135, 2009.
[5] B. Cockburn. A Multilevel DRAM with Hierarchical Bitlines and Serial Sensing. In *Proc. MTDT*, pages 14–19, 2003.
[6] T. Furuyama et al. An Experimental 2-bit/Cell Storage DRAM for Macrocell or Memory-on-Logic Application. *JSSC*, 24(2):388–393, 1989.
[7] S. Ghosh and K. Roy. Parameter Variation Tolerance and Error Resiliency: New Design Paradigm for the Nanoscale Era. *Proc. of the IEEE*, 98(10):1718–1751, 2010.
[8] P. Gillingham et al. A Sense and Restore Technique for Multilevel DRAM. *TCAS-II*, 43(7):483–486, 1996.
[9] N. Ikeda et al. A Novel Logic Compatible Gain Cell with two Transistors and one Capacitor. In *Proc. VLSIT*, pages 168–169, 2000.
[10] H. Kaeslin. *Digital Integrated Circuit Design: From VLSI Architectures to CMOS Fabrication*. Cambridge University Press, 1st edition, 2008.
[11] S. Kang and Y. Leblebici. *CMOS Digital Integrated Circuits: Analysis and Design*. McGraw-Hill, 3rd edition, 2003.
[12] W. Luk and R. Dennard. 2T1D Memory Cell with Voltage Gain. In *Proc. VLSIC*, pages 184–187, 2004.
[13] W. Luk et al. A 3-Transistor DRAM Cell with Gated Diode for Enhanced Speed and Retention Time. In *Proc. VLSIC*, pages 184–185, 2006.
[14] C. Novak et al. The Effect of Unreliable LLR Storage on the Performance of MIMO-BICM. In *Proc. ACSSC*, 2010.
[15] T. Okuda and T. Murotani. A Four-Level Storage 4-Gb DRAM. *JSSC*, 32(11):1743–1747, 1997.
[16] D. Somasekhar et al. 2GHz 2Mb 2T Gain-Cell Memory Macro with 128GB/s Bandwidth in a 65nm Logic Process. In *Proc. ISSCC*, pages 274–613, 2008.