A 128kb High Density Portless SRAM Using Hierarchical Bitlines and Thyristor Sense Amplifiers

Michael Wieckowski, Gregory K. Chen, Daeyeon Kim, David Blaauw, Dennis Sylvester
Department of Electrical Engineering and Computer Science
University of Michigan, Ann Arbor, MI

1. Abstract

A 128 kb portless SRAM is presented with 1024 rows per hierarchical bitline and CMOS thyristor-based local sense amplifiers. Each portless cell is 0.317 μm^2 in 45nm CMOS and consumes 50.8 fJ of energy per access at a 17.86 ns cycle time. A 65% read SNM improvement and a 33% leakage power reduction is achieved over a conventional 6T design. The thyristor-based sense amplifier occupies 2.4 μm^2 and provides variation tolerant sensing within the hierarchical column pitch.

2. Keywords

SRAM, memory, low-power, noise-margin, stability

1. Introduction

Process scaling and the aggressive specialization of SRAM design rules have been accompanied by a trend towards shorter bitlines as a way to cope with cell leakage and cell stability variation [1, 2]. Shorter bitlines inhibit high array efficiency since the relative area overhead of peripheral support circuitry remains large. This is especially true for sense amplifier circuits where column pitch layout compatibility cannot be achieved without sacrificing mismatch tolerance. One alternative solution is to employ segmented or hierarchical bitlines to effectively amortize peripheral area overhead across multiple cell columns [3]. While feasible for logic design rules and non-6T cell structures, this technique is difficult to implement with high-density SRAM cells because of a limited number of available metal tracks per cell width. Even though the design rules can be aggressively optimized for the SRAM devices, little can be done to increase metallization density. This results in a typical value of four metal tracks per column pitch including one dedicated power track. Therefore, when implementing a hierarchical design, local I/O circuitry within each column must be able to communicate read and write data to and from the global bitlines with low area overhead and without requiring additional metallization layers [4]. Achieving this within the standard 6T column pitch while ensuring local sense amplifier robustness poses a major challenge. In this work, a new portless memory cell is presented along with a CMOS thyristor-based local sense amplifier (LSA) to enable hierarchical bitline organization with 1024 rows per global bitline using industrial SRAM design and metallization rules.

2. Hierarchical Portless SRAM

The canonical portless cell uses only five-transistors and communicates read and write data through its power rails instead of the traditional wordline devices seen in 6T designs. Previous portless memories have used only one of their supplies (power or ground) for both read and write, resulting in the same metal track limitations as in a standard

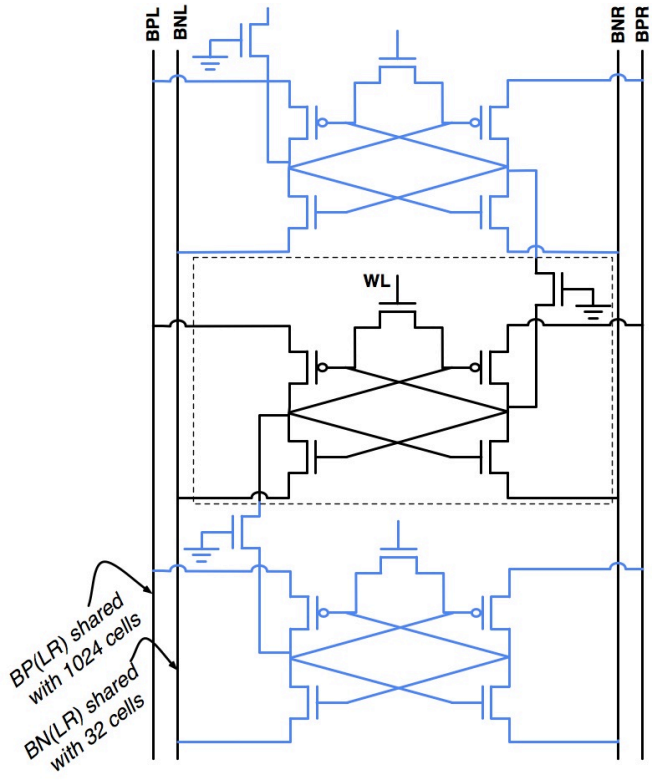

Figure 1: Portless cell schematic within the column bitlines

6T cell [5]. The cell proposed in this work, shown in Fig. 1, uses both supplies, reading data through differential ground rails and writing data through differential power rails. A local column can therefore be implemented using two metal layers for the cell routing and a single four-track metal layer that acts as the local bitlines and the power supply. Writes are performed in this new scheme by skewing the P bitlines (BPL and BPR) and pulsing the WL with Vdd as shown in Fig. 2. This effectively shorts the cell nodes together, allowing them to return to a state corresponding to the

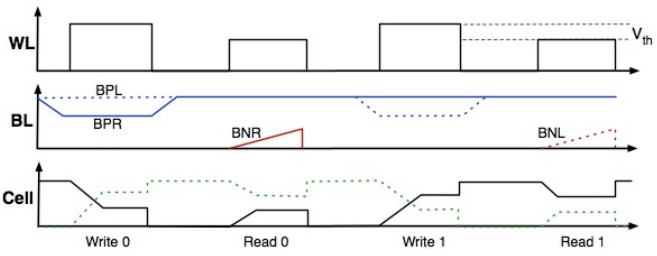

Figure 2: Read and write timing of a portless cell in the local column hierarchy

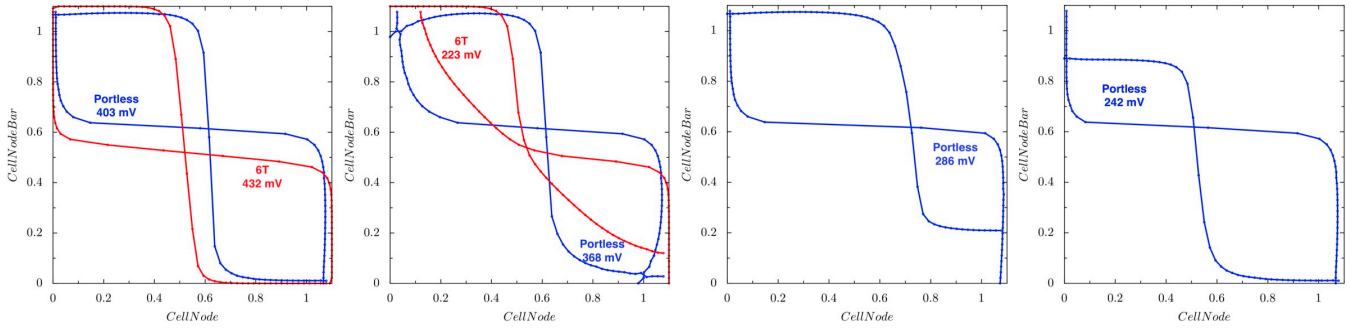

Figure 3: Standby, read, pseudo-read, and pseudo-write noise margins of the portless memory cell in a local column

bitline skew. Reads are performed by floating the N bitlines (BNL and BNR) and pulsing the WL with a voltage that is one Vth below the supply voltage. The read voltage and cell sizing are chosen to prevent the data nodes from fully shorting, thus allowing the cell to force a differential onto its N bitlines without corrupting the cell's contents. By effectively combining the power, ground, read, and write bitlines, a hierarchical organization can be realized with no additional metallization layers.

It is important to note that a sixth transistor is included in the proposed cell design which is permanently turned off as seen in Fig. 1. The presence of this device is the result of using a conventional 6T "form factor" to take advantage of industrial SRAM design rules while guaranteeing manufacturability. An existing 6T cell metallization layout was effectively modified such that one of its wordline devices becomes the portless wordline device and the other is permanently grounded. Interestingly, by adding a second wordline decoder, this unused transistor can be turned controlled to enable a variety of cell-to-cell operations including reads, writes, and logical combinations. Further research is ongoing to evaluate the benefits of such a scheme.

3. Noise margin and cell stability

Fig. 3 depicts the noise margins extracted using Arrows [6] for the hierarchical portless cell and an industrial low-power 6T cell. In standby, the portless and 6T wordline devices are all off resulting in large noise margins as expected for a cross-coupled inverter pair. During read, the noise margin of the portless cell is 65% larger than the 6T cell due to the reduced wordline gate overdrive voltage. Increasing this voltage results in a faster cell response with a lower noise margin. Unaccessed portless cells in the same column can experience pseudo-read and pseudo-write conditions due to the shared bitline configuration, which is a situation not present in 6T designs. When one cell is being accessed for read or write, the supply voltage of the other cells in the same column is effectively reduced. For 200 mV read and write bitline differentials, the noise margins of the unaccessed portless cells are still higher than the 6T noise margin by 28% and 9%, respectively.

4. Thyristor based local sense amplifier (LSA)

To sense and propagate data from the local N bitlines to the global P bitlines, a sense amplifier based on a cross-coupled pair of CMOS thyristors was developed as shown in Fig. 4. When pre-charged, all cross-coupled devices are turned off resulting in extremely low gain. When released from pre-charge, the circuit enters a meta-stable "dead zone" where OUT and OUTbar are both low. Only a large voltage differential can move out of this zone and into the high gain region. In Fig. 5, a vector plot extracted using Arrows [6]

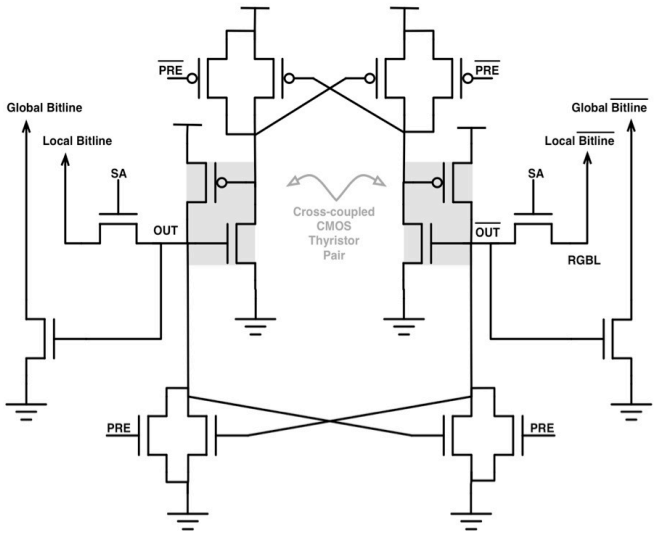

Figure 4: Thyristor based local (within column) sense amplifier with global column interface

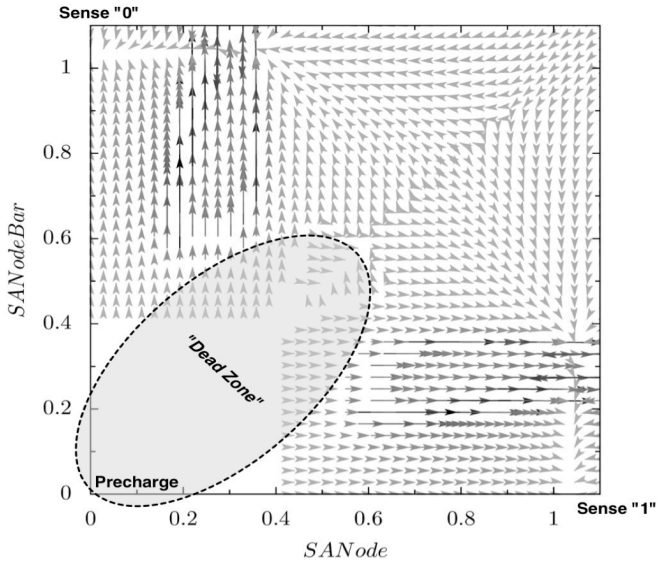

Figure 5: State space gain plot of the thyristor sense amplifier depicting mismatch tolerant dead zone

shows the direction and gain of the LSA output nodes for each point in its state space. The dead zone is clearly visible on the lower left quadrant of the state space. When a significant voltage differential is applied by the local bitlines, the output nodes of the sense amplifier will transition into the high gain regions and create a differential on the global bitlines. This property makes the LSA highly tolerant to process variation and enables layout of the LSA using SRAM design rules and minimum sized devices within a single column pitch. In 1000 Monte Carlo simulations with random Vth mismatch, an industrial sense amplifier with the same area constraint exhibits an 11% failure rate whereas the proposed LSA is 100% functional. Without area constraints, the industrial sense amplifier can sense a minimum bitline differential of 72 mV including random Vth mismatch, while the proposed LSA can sense 82 mV, a 10 mV increase for only ~7x reduction in area.

5. Bank organization

Fig. 6 shows the global column organization using hierarchical bitlines within a single memory bank. Each local column contains 32 cells sharing a common N bitline and LSA. Each global column contains 32 local columns sharing a common P bitline and global read bitline. Each LSA is effectively tri-stated onto the global read bitline through its reset mechanism, allowing only one LSA in the hierarchy to latch at a time. The sense amplifiers for the global read bitline (GSA) use a standard current mode design. To generate the two distinct WL voltages required for reads and writes, an NFET output stage is used in the wordline driver for read operations and is put in parallel with a conventional PMOS pull-up for write operations. This ensures that the WL will not rise above Vdd - Vth during the read operation, thus guaranteeing cell stability.

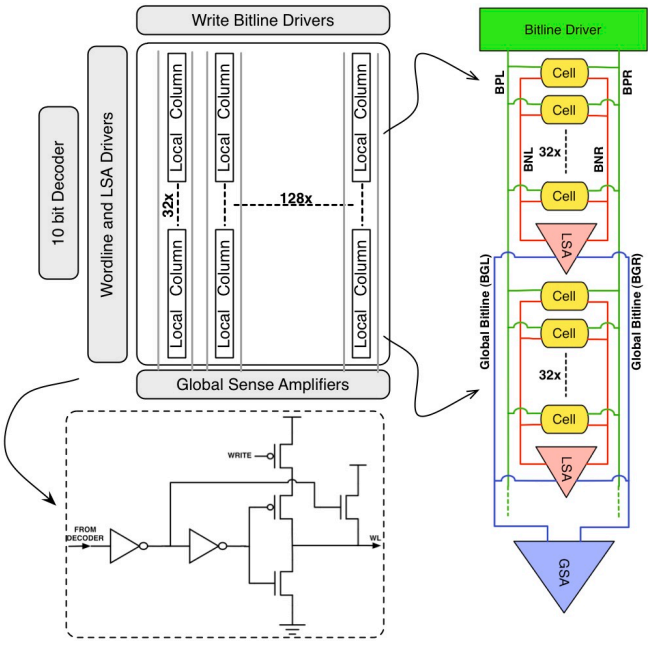

Figure 6: Array organization, dual voltage wordline driver, and column hierarchy with local and global SA's

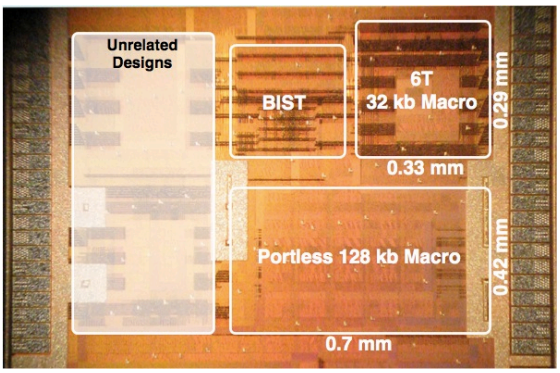

Figure 7: Micrograph of the fabricated test chip

6. Measurements And Results

A test chip was fabricated in a low-power 45 nm CMOS process containing a 128 kb hierarchical portless array and a 32 kb commercial 6T array (Fig. 7). Pushed design rules (SRAM) were used in both arrays including the hierarchical LSAs and dummy cells. To leverage these design rules, the five-transistor portless cell uses a standard 6T form factor with all minimum-sized devices. A metal2 wire is used to create the portless access device from one of the 6T wordline (WL) devices as shown in Fig. 8. The remaining 6T WL device is permanently turned off and left in place to guarantee manufacturing yield. The total cell area is 0.317 μm^2.

The portless array, with LSAs, dummy cells, and tie cells, occupies 0.058 mm^2. The average dead zone time measured for the LSA is 10x longer than required for functionality and demonstrates the robustness concept

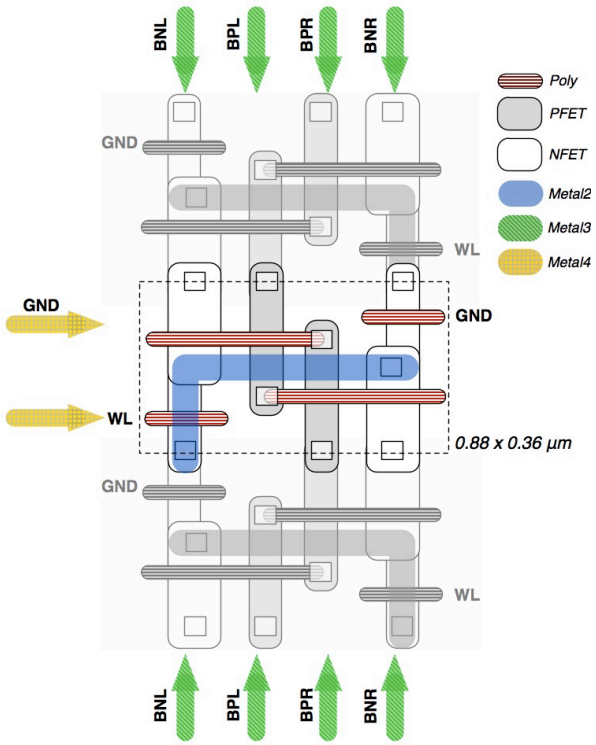

Figure 8: Layout cartoon of three portless cells in a column using the standard 6T form factor and four metal layers

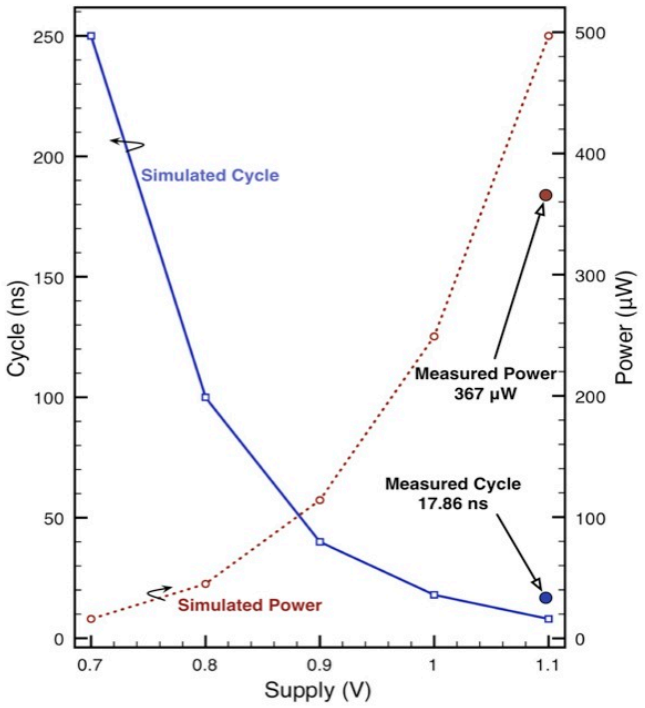

Figure 9: Measured and simulated cycle time (left axis) and power (right axis) vs supply

(bitline differentials are generated on the ns-scale yet the measured LSA's can tolerate µs). At 1.1 V, a period of 17.86 ns was measured with energy consumption of 50.8 fJ per bit per access (Fig. 9). Due to the small size of the portless bitcell, the measured leakage per bit at 1.1 V is 33% lower than the comparison 6T array. This improvement includes amortized leakage overhead for the LSAs and the bitline drivers that power the cells. Table 1 provides an overall performance summary of the hierarchical portless design.

Table 1: Performance Summary

Design Parameter	Measured Results @ 1.1V
Technology	45nm, 7 Metal CMOS
Design Rule	SRAM Pushed Rules
Supply Voltage	1.1 V
Cell Size	0.317 µm^2
Organization (row x column)	1k x 8 words x 16 = 128 kb
Bitline Hierarchy (local x global)	32 bits x 32 sub-bitlines
Array Size (cells, LSA's, dummies)	0.058 mm^2
LSA Dead Zone Time	10.84 µs
Bank Cycle Time	17.86 ns
Leakage Per Cell (Amortized to include LSA's and drivers)	48.1 pW
Energy Per Access Per Bit	50.8 fJ

7. Acknowledgment

The authors thanks ST Microelectronics for support in chip fabrication.

8. References

[1] A. Kawasumi et al., "A Single-Power-Supply 0.7V 1GHz 45nm SRAM with An Asymmetrical Unit-β-ratio Memory Cell," Proceedings of the IEEE International Solid State Circuits Conference, pp.382-622, Feb. 2008.

[2] Y. Ye et al., "A 6-GHz 16-kB L1 cache in a 100-nm dual-VT technology using a bitline leakage reduction (BLR) technique," IEEE Journal of Solid-State Circuits, vol. 38, no.5, pp. 839- 842, May 2003.

[3] K. Kushida et al., "A 0.7V single-supply SRAM with 0.495um^2 cell in 65nm technology utilizing self-write-back sense amplifier and cascaded bit line scheme," Proceedings of IEEE Symposium on VLSI Circuits, pp.46-47, June 2008.

[4] S. Ishikura et al., "A 45 nm 2-port 8T-SRAM Using Hierarchical Replica Bitline Technique With Immunity From Simultaneous R/W Access Issues," Journal of Solid-State Circuits, vol.43, no.4, pp.938-945, April 2008.

[5] M. Wieckowski et al., "Portless SRAM - A High-Performance Alternative to the 6T Methodology," Journal of Solid-State Circuits, vol.42, no.11, pp. 2600-2610, Nov. 2007.

[6] M. Wieckowski et al., "A black box method for stability analysis of arbitrary SRAM cell structures," Design, Automation & Test in Europe Conference & Exhibition (DATE), 2010, pp. 795-800, March 2010.