

Design and Characterization of an Embedded ASIC DRAM

Gershon Birk, Duncan G. Elliott, Bruce F. Cockburn

Department of Electrical and Computer Engineering
University of Alberta, Edmonton, Alberta, Canada T6G 2G7
{birk, elliot, cockburn}@ee.ualberta.ca
<http://www.ee.ualberta.ca/~elliott/fsmemory/>

Abstract

University DRAM research is hindered by the lack of access to specialized commodity DRAM or blended logic-DRAM processes. In this paper we describe the design of an embedded DRAM in the 0.35 μ m TSMC logic process, available through the Canadian Microelectronics Corporation (CMC). Our test chip design used a variation of the HDRAM[®] macro cells developed by MOSAID Technologies Inc. This paper describes the DRAM and gives some preliminary test results.

1 Introduction

As larger systems-on-a-chip are being designed in industry, it is becoming increasingly important to investigate different strategies for embedding large memories along with digital processing circuitry. Static RAM has frequently been used in embedded applications because of its compatibility with standard logic ASIC processes. Embedded DRAM (dynamic random access memory) is currently being considered as an alternative to embedded SRAM because of the higher potential storage density of DRAM over SRAM. Embedding DRAM along with logic is a complex challenge because the processes for the two technologies have been developed to optimize key parameters in often conflicting directions. Blended DRAM/logic technologies will tend to be the best approach for high-performance industrial designs. There are situations, however, where it is desirable to embed DRAMs in a logic process, even if this involves sacrificing some of the density advantage of a blended logic and DRAM process.

In this paper we describe the design of a conventional 1-bit per cell embedded DRAM in the 0.35 μ m TSMC logic process. The DRAM used HDRAM[®] macro cells developed by MOSAID Technologies Inc. [1-2]. Our prototype uses one level of polysilicon and three levels of metal interconnect (although only two metal levels are essential to our design).

Testing and characterizing a DRAM at a University poses challenges. In our design we avoided the complexity of including an on-chip timing generating circuit by bringing all of the control and timing signals out separately to the pads. This also gives us maximum flexibility to help with characterizing the prototype dies. For example, we have complete freedom to test alternative control timing parameters. However, the large number of separately timed edges (over a dozen are required) made testing a challenge. We will describe below how the test and characterization challenges were solved using CMC-supplied test equipment.

This test chip was designed to support both a 1-bit per cell classic DRAM mode and a 2-bit per cell multilevel DRAM [3] mode of operation. Due to a place&route error, the multilevel DRAM mode is not functional. We will therefore only describe the performance characteristics while in 1-bit per cell operation.

1 ASIC DRAM Design Overview

Figure 1 shows two bitline pairs from our design, which included a total of 8 bitline pairs. A sense amplifier is present at the left end of each bitline pair. The sense amplifiers are connected by isolation transistors to a folded bitline pair (parallel bitlines). The bitline pairs

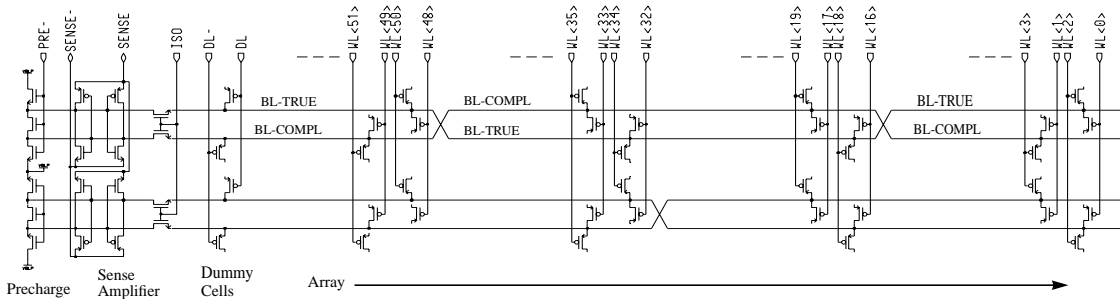


FIGURE 1. DRAM Core Schematic

are provided with one or two twists where the bitlines cross over each other and exchange places. This technique is used to better cancel out capacitively coupled noise. The vertical lines at the right in Figure 1 (denoted WL<X>) are wordlines; storage capacitors are present (but not shown) and cell access transistors are located at alternate intersection of bitlines and wordlines. The even-numbered wordlines control access to cells on the true bitline (BL-TRUE) while odd numbered wordlines control access to the cells on the complement bitlines (BL-COMPL). Two additional dummy-wordlines are provided (DL and DL-). They are connected to one dummy cell per bitline. These dummy-wordlines can be used to cancel charge injection induced by wordline transitions.

The DRAM core is an adaptation of MOSAID's HDRAM. It uses a 0.35 μ m single poly process. PMOS transistors are used for cell capacitors. It is desirable in DRAM to minimize cell leakage by controlling the substrate bias on the cell access transistors. Putting a negative bias on the IC substrate, however, would slow all digital logic that may be present on the same die. Instead, the memory cells are placed in an n-well, requiring PMOS cell access transistors. Putting the cells in an n-well (in this single-well process) has the additional advantage of isolating the memory array from substrate noise.

Wordline signal boosting is a challenge in this process since the wordlines must be driven below V_{SS} to allow the transfer into the storage

cell of a full V_{SS} signal. In HDRAM, a one-shot charge pump is used to deliver the required negative voltage to the wordlines. A large capacitance is required to store enough charge to drive the 2000+ original HDRAM gate load of the wordline. In adapting the design for multilevel DRAM, the charge pump scheme was abandoned altogether. Instead, the basic wordline driver is operated through off-chip control, which also gives full control over the boost voltage. The row decoding and p-channel switch structure of the HDRAM wordline driver were retained, but the terminals of the capacitors were grounded and the charge pump control was removed. There are actually two charge pumps in HDRAM which, when removed, leave two terminals that need separate boosting control and power. These two independent circuits power alternating wordlines and are part of the decoding scheme.

2 Design Details

The peripheral circuitry consists of row and column address predecoders, column decoders, address latches, data I/O, probe pads and bonding pads. All of these blocks were implemented with Cadence cell place & route tools except for a small bit of custom layout for the data bus drivers and the databus precharge switches.

The row (X-address) predecoder uses seven latches to hold the row address. These latches are positive edge sensitive to a clock signal that

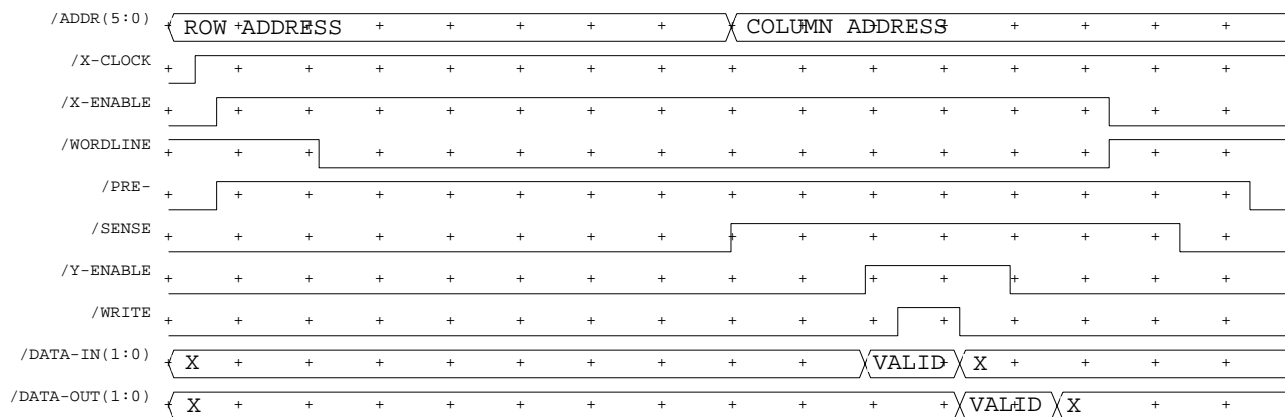


FIGURE 2. DRAM Control Timing

is used solely for these latches. Once latched, the row is pre-decoded and presented to the decoder.

Unlike the row predecoder, the column (Y-address) predecoder does not latch the data. Instead it is combinational logic which continuously feeds the column decoder with the predecoded data from the first four address pins.

The column decoder is not in the pitch of the bitline. Rather, it is placed near the other standard cell logic. The choice of not fitting the column decoder within the bitline pitch was made to reduce design time. Together the predecoder and decoder select one of sixteen possible sense amplifiers. The final control of sense amplifier access is given to the column enable signal Y-ENABLE. This signal must be asserted for any access to any column. This control is achieved with AND gates whose drive is sufficient to drive the parasitic capacitance of the relatively long length of wire connecting the decoder and the sense amplifier access transistors.

A differential data bus connects all eight sense amplifiers to the I/O logic. Due to pad limitations the bus is only one bit wide, but this is sufficient for experimental purposes. The data bus is normally held precharged; but when the Y-ENABLE signal is asserted the precharge is turned off. The read data is always present at the pins but is only valid during a read operation.

The write operation is done through large data bus drivers. These custom drivers overpower the sense amplifier that is selected by the regular column selection scheme.

There are four power supplies for the chip: V_{DD} , V_{SS} , V_{BLP} , and V_{CP} . The standard cell design kit included V_{DD} and V_{SS} pads. A low impedance analog input pad was used for V_{BLP} . A custom pad was made for V_{CP} because it is to be used at least one threshold above V_{DD} . For the custom pad the ESD device and V_{DD} clamping diodes were removed to allow the high voltage to be passed. Three additional special pads are needed for WORDLINE and DUMMY-WORDLINE. These signals are switched below V_{SS} and are for driving the wordlines and dummy-wordlines. Again, to prevent damage and failure the ESD and V_{SS} clamp diodes were removed. All other pads are digital pads supplied in the standard cell package from CMC.

Twelve test pads were included to allow further probing into the circuit. They connect to the test points on the seventh folded bitline only (column address 110). They provide differential amplifier probing of the bitlines on either side of the isolation transistors for both sub-bitlines. They are made from all three metal layers connected with a dense array of via's that were provided to ensure adhesion.




FIGURE 3. Chip Plot

A plot of the IC is shown in Figure 3. Within the padframe, note the test pads, memory array, and standard cell going from top to bottom.

3 DRAM operation

To access the DRAM on the chip all unrelated chip functions are disabled and the remaining pins are forced according to the timing diagram shown in Figure 2. The isolation transistors were not used in our experiments, but are required for multilevel DRAM.

At the start of the cycle the row address is presented to the address pins (ADDR) and latched by the rising edge of X-CLOCK at $t=1.7\text{ns}$. This event triggers the predecoding of the address. Sufficient time is allowed for the predecoding to complete before the X-ENABLE signal is asserted. At the same time, the bitline precharge is turned off, leaving the bitlines floating at V_{BLP} . The value of V_{BLP} is another controlled parameter that we set to $V_{DD}/2$.

The X-ENABLE signal is used to prevent any hazards in the predecoding logic from causing instability on the wordlines. It is the final step in decoding and, when raised, causes one of the wordlines to be isolated from V_{DD} and connected to the WORDLINE pin. Through a bootstrapped switch, the addressed wordline may be driven from a high voltage (V_{DD}) to a low voltage. The low voltage may be set below V_{SS} when boosted wordlines are desired. If this is done, however, charge injection may cause a significant deterioration of the bitline signals. The dummy-wordlines present on the chip can be used to counter this effect. For this experiment we opted to use a V_{DD} to V_{SS} swing.

At $t=9.5\text{ns}$ the WORDLINE signal is driven to V_{SS} and the selected wordline is activated. Charge from the corresponding cells is shared with the bitlines causing a differential voltage at the sense amplifiers. A large time is left for this operation because of the large capacitance and limited current drive of the wordline circuitry. At $t=24.0\text{ns}$ the sense amplifiers are enabled and all eight bits in the column are sensed.

After sensing is complete a column access may be performed. Prior to enabling the column

(through the Y-ENABLE pin) at $t=31.6\text{ns}$ the desired column address must be stabilized on the three least significant ADDR pins. Provided that this has been done, the desired column is accessed and the sense amplifier on that column begins charging the data bus.

For a destructive write, the WRITE pin may be asserted shortly after the Y-ENABLE signal is asserted. The operation of the Y-ENABLE and WRITE pins is restricted to this order of usage because the Y-ENABLE pin directly controls the precharging of the databus. Should the WRITE pin be high while the Y-ENABLE pin is low, a large current would flow between power supplies V_{DD} , V_{SS} , and V_{BLP} causing possible failure. When used in the proper manner, the WRITE pin need only be asserted for a short time (4ns was the minimum pulse width supported by our tester) for the write operation to complete.

For a read operation, the WRITE pin is left low and the data in the addressed column is sent onto the databus and amplified to a digital value before becoming available at the DOUT pin at $t=38.7\text{ns}$.

After the column operation is complete, the data can be restored to the cell by deactivating the wordline. Both the X-ENABLE and the WORDLINE signals are de-asserted and the addressed wordline quickly returns to V_{DD} .

The memory cycle is completed with the deactivation of the sense amplifiers followed the assertion of the PRE- signal. The next cycle can begin after some time is allowed for bitline precharging to complete.

Tests and Results

The test chip was fabricated through CMC (see Figure 3) and tested on a 60MHz IMS tester. The device was fixtured on a custom PGA socket. Care was taken to keep all wires short and equal in length.

Several standard march tests were applied to evaluate the gross functionality of the DRAM arrays. All cells were functional on all five test chips under nominal conditions using a 5n march test [4]. Having functionally verified the chips, we then experimented with the timing to minimize the access time. We had the advantage of direct control over all key timing events, control that is not available to users of commodity DRAMs via the RAS and CAS signals.

The basic strategy for finding the fastest possible timing was to consider the six key intervals within an access cycle (see Figure 2). Beginning with the earliest interval, each was progressively reduced while all other intervals were held at conservative values. Table 1 summarizes the results. An address to data-out access time of 38.7ns was measured on all five test chips.

TABLE 1. Critical Timing Intervals

Interval	Bounding Edges	Value (ns)
11	ADDR row to X-CLOCK rise	1.7
12	X-CLOCK rise to X-ENABLE on	1.3
13	X-ENABLE on to WORDLINE on	6.5
14	WORDLINE on to SENSE on	14.5
15	SENSE on to Y-ENABLE on	7.6
16	Y-ENABLE on to D-OUT valid	7.1

Testing was performed at 3.3V with an ambient temperature of 22C. The memory chips function correctly down to $V_{DD}=2.7\text{V}$.

4 Conclusions

In this paper we described the implementation and preliminary characterization of DRAM test chips implemented in the 0.35 μm TSMC ASIC process. All five dies were verified to be fully functional at nominal timing, corresponding to a 38.7ns access time.

There are further opportunities to reduce the access and cycle times. In particular, the standard cell library and several buffers from the HDRAM library could be resized for our particular design.

We were successful in demonstrating the functionality of a conventional ASIC DRAM in a process that is available through CMC. Our next step is to proceed with the characterization of our multilevel DRAM design.

Acknowledgments

We are grateful for the support of MOSAID Technologies Inc., Micronet, CMC and NSERC under grant OGP0105567. MOSAID has been particularly generous by hosting the first author over two summers, providing access to proprietary design libraries, and answering all manner of questions. Valerie Lines, Randy Torrance, Dan Klein and Peter Gillingham have been exceptional sources of memory design and layout knowledge. We are also grateful to Roger Mah for his preliminary layout work and Albert Kwong for performing place and route.

References

- [1] Peter Gillingham, Betina Hold, Ian Mes, Cormac O'Connell, Paul Schofield, Karl Skjaveland, Randy Torrance, Tomasz Wojcicki and Henry Chow, "A 768k Embedded DRAM for 1.244Gb/s ATM Switch in a 0.8 μ m Logic Process", In *IEEE International Solid-State Circuits Conference*, pp. 262-263, February 1996.
- [2] R.C. Foss, "Implementing Application Specific Memory," 1996 IEEE International Solid State Circuits Conference, San Francisco, CA, Feb 8-10, 1996, pp. 260-261.
- [3] P. Gillingham, "A Sense and Restore Technique for Multilevel DRAM," *IEEE Trans. On Circuits and Systems-II: Analog and Digital Signal Processing*, July 1996, Vol. 43, No. 7, pp. 48
- [4] A.K. Sharma, "Semiconductor Memories Technology, Testing and Reliability," IEEE Press, Piscataway NJ, 1997.