A Comparative Simulation Study of Four Multilevel DRAMs

Gershom Birk, Duncan G. Elliott, and Bruce F. Cockburn
Department of Electrical and Computer Engineering, University of Alberta
Edmonton, Alberta, Canada T6G 2G7
{birk,elliott,cockburn}@ee.ualberta.ca http://www.ee.ualberta.ca/~elliott/memory/

Abstract

Multilevel DRAM (MLDRAM) attempts to increase storage density by recording more than one bit per cell. Several different two-bit-per-cell schemes have been described in the literature; however, it is difficult to compare them directly because the original papers use different technologies and operating conditions. This paper presents a detailed simulation study that compares three published MLDRAM schemes, along with a new MLDRAM scheme that combines the speed of a MLDRAM proposed by Furuyama et al. and the noise cancellation techniques of a MLDRAM proposed by Gillingham. Our SPICE simulation models use the same array size and process models for each to allow us to make direct comparisons.

1. Introduction

It is becoming increasingly expensive to raise the storage density of DRAM by reducing the physical size of the cells and using complex three-dimensional cell capacitor structures. One additional dimension, which has yet to be successfully exploited in commercial parts, is to store more than one bit per cell. In a DRAM cell this involves storing and then subsequently sensing more than two distinct voltage levels on the cell capacitor. Figure 1 illustrates how one might encode the four logic pairs 00, 01, 10, and 11 as four equally-spaced voltage levels in the range V_{SS}=0 to V_{DD}. To extract the two bits from a cell, the cell voltage must be compared to at least two of the reference levels. For example, the most significant bit (MSB) can be determined by comparing the cell voltage with an MSB reference of V_{DD}/2. If the result of this comparison yields a logical 1 (i.e. $V_{CELL} > 3V_{DD}/6$) then a second comparison is done with a least significant bit (LSB) reference level of $5V_{DD}/6$. Otherwise, a logical 0 result from the first comparison implies the use of an LSB reference of $V_{DD}/6$. In this way both bits can be extracted sequentially.

One important challenge in MLDRAM that is evident in Figure 1 is that the differential signal strength, and thus the noise margins for a 4-level MLDRAM, will be only one third those of a conventional 2-level DRAM with the same supply voltage. This makes MLDRAM more vulnerable to soft errors, leakage current from the cells, and sense amplifier offsets. It is still unclear whether or not the advantages of MLDRAM can be achieved despite these technical challenges.

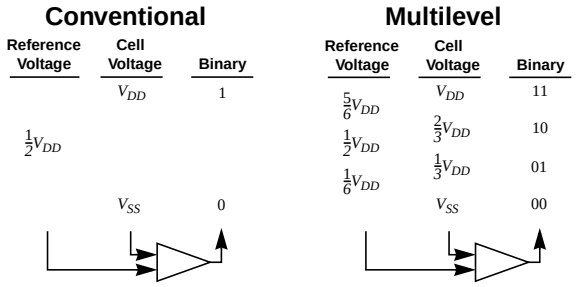

Figure 1. A two-bit-per-cell storage scheme.

Several different MLDRAM schemes have been proposed in the literature [1-4]. They differ in the techniques used to store the four analog cell voltages and/or the techniques used to sense and restore the cell signals. It is difficult to compare these schemes directly because they were described for different processes at different times. To better gauge the real advantages and disadvantages of the alternative schemes, it is important to compare them fairly using the same process parameters, array size, and device models. In this paper we describe a simulation study that attempts to provide a fair comparison of the MLDRAM schemes described by [2], [3], [4], and a new scheme that we are proposing. We did not consider a 4-level version of the 16-level MLDRAM scheme proposed by [1] because fundamental limitations (e.g. long access time) render it uncompetitive with the other designs.

This paper is organized as follows: In the next section we review the three published MLDRAM schemes that were considered in our study. Section 3 introduces a new MLDRAM that attempts to combine the advantages of the MLDRAMs of [2] and [3]. Section 4 details the assumptions and conditions under which our simulations were performed. Section 5 presents our simulation results, and the implications of these results are discussed in Section 6. Section 7 makes concluding remarks and suggests appropriate directions for further research.

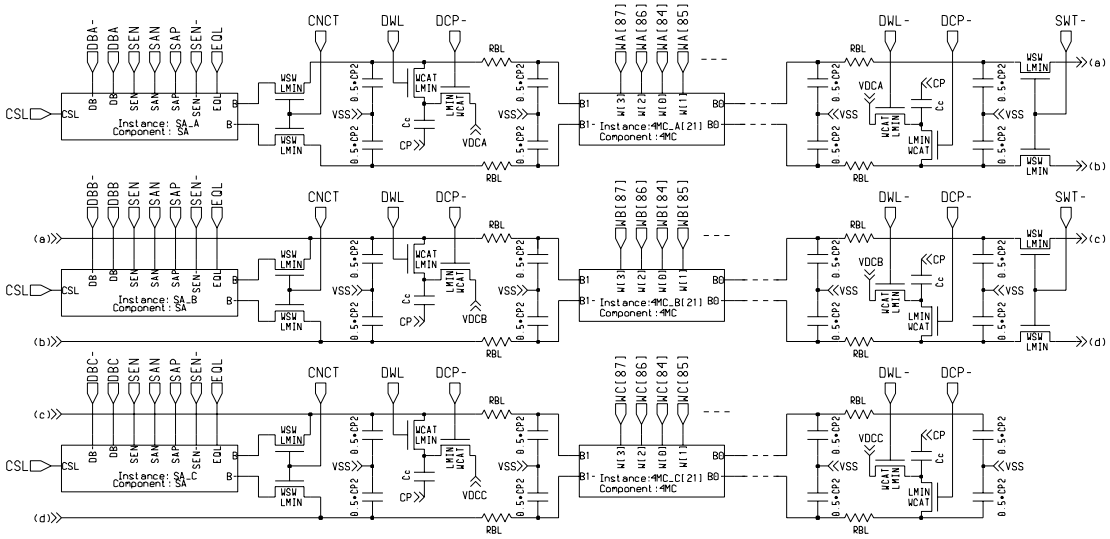

Figure 2. Schematic used for simulating the MLDRAM from [2].

2. A Review of Published MLDRAMs

The key internal operations in DRAMs involve sensing and restoring cell voltages. The sensing and restoring techniques for conventional two-level DRAM are well known and do not vary much between designs [5]. When implementing MLDRAM, it is the sense and restore scheme that becomes the most difficult circuit to design. The multilevel sense and restore scheme must provide the capability to extract data encoded as one of many allowed voltage ranges on a capacitor (the sense operation), and have the capability to take multiple bits and convert them to one of many nominal voltage levels (the restore operation).

To analyze the various sensing schemes that have been proposed, the strengths and weaknesses of each potential candidate must be addressed. First off, circuitry size is critical. The sense and restore circuitry must at least fit in the width of one or two columns of memory cells (the bitline pitch). Second, insensitivity to process variation is highly desirable. In the charge sharing operations used in MLDRAM, the capacitances involved depend on many process parameters, each of which impacts differently the cell and bitline capacitances. Third, noise insensitivity is required. DRAMs are electrically balanced so that the common mode noise rejection in the sense amplifiers is very high. This requirement is important because, as noted earlier, in MLDRAM noise margins are decreased considerably. Fourth, MLDRAM should be comparable in speed to DRAM. This requirement is difficult to meet because of the potential increase in sequential steps needed to extract and then write back multilevel encoded data.

2.1. Furuyama's Multilevel DRAM

T. Furuyama *et al.* propose a multilevel sense and restore method in [2]. In this scheme four voltage levels are mapped to two bits as described in Figure 1. To read multilevel data, the cell charge is shared with a bitline. The bitline itself is modified so that it can be split via signal SWT- into three equal parts called sub-bitlines (see Figure 2). After sharing the cell signal charge equally with the three sub-bitlines, they are isolated from one another via the SWT- switches and a sense amplifier is connected to each via the CNCT switches. The multilevel data is compared in parallel to the three references (those of Figure 1 diluted by the ratio of the cell capacitance to the bitline capacitance) that are available through the dummy wordline (DWL). A data-bus carries the three sense amplifier outputs to a buffer after which the three logic values are converted into two bits according to the function shown in Table 1.

Table 1: Conversion Function for [2]

Sense-Amplifier Results	Two-Bit Data Value
000	00
001	01
011	10
111	11

Restoring the data involves simply disconnecting the sense amplifiers from the three sub-bitlines and then reconnecting the sub-bitlines together. Charge sharing will give a final voltage on the bitline equal to the correct value (V_{DD}, $2V_{DD}/3$, $V_{DD}/3$, or V_{SS}). The wordline is then brought to a low voltage thus capturing the desired multilevel voltage in the addressed cell.

The advantages of this design are that it is fast and relatively simple. Using three sense amplifiers means that the two bits are available at the same time. This also means more area is devoted to supporting circuitry rather than to storage cells, reducing the potential density gain that MLDRAM technology offers. The main disadvantage of

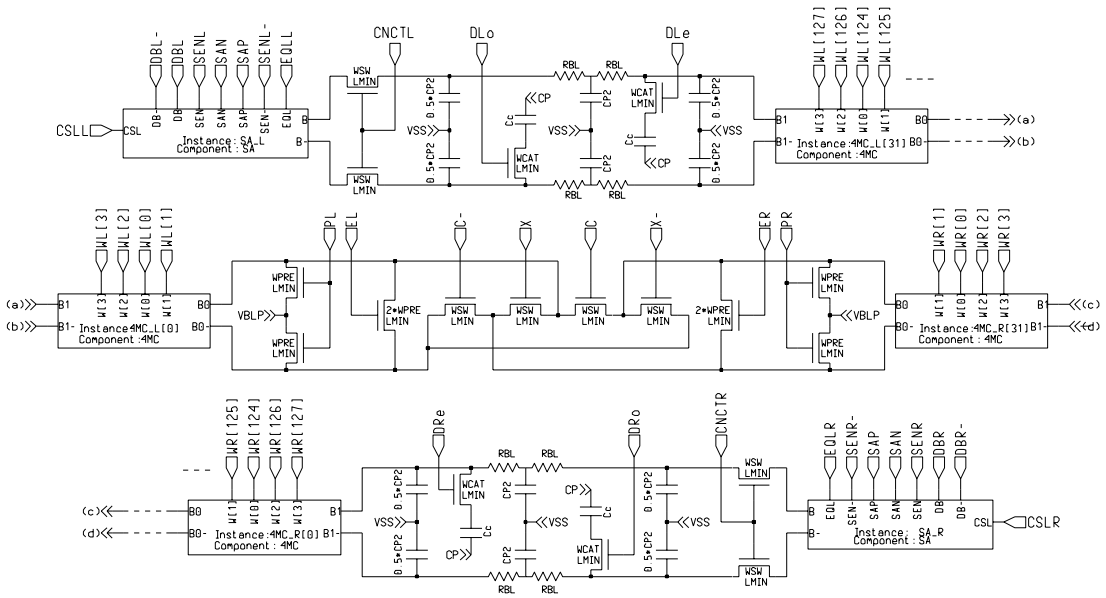

Figure 3. Schematic used for simulating the MLDRAM from [3].

the scheme is that it is susceptible to sensing errors due to improper reference values. The potential for such errors arises from the use of global reference voltages (VDCA, VDCB, VDCC) that need to be generated on chip and distributed across the array. A slight inaccuracy in the global reference voltage levels would be enough to introduce errors.

2.2. Gillingham's Multilevel DRAM

Another MLDRAM scheme is proposed by Gillingham [3]. The same restore technique as [2] is employed, but the sensing method is quite different. The architecture consists of cell columns implemented using two pairs of sub-bitlines, with each pair having a sense amplifier at one end (Figure 3). The four sub-bitlines can be connected in six different ways using a transistor switch matrix. As well, each sense amplifier can be disconnected from its respective sub-bitline pair through yet another set of switches. Unlike the parallel operation of the sense amplifiers used in [2], this scheme uses sequential sensing. The choice of reference level used in the second sensing operation comes from the result of the first sensing operation. Referring to Figure 3, an initial MSB sensing operation compares the multilevel data to $V_{BLP}=V_{DD}/2$. The reference for the final sensing operation is generated by storing the sensed MSB value (either V_{DD} or V_{SS}) back in the accessed cell, precharging three sub-bitlines to $V_{DD}/2$, and then sharing the MSB charge from the cell onto the three sub-bitlines. If the first sensing operation reveals that the cell voltage is above $V_{DD}/2$ then the second sensing operation will compare (after dilution) the cell voltage to an LSB reference of $5V_{DD}/6$. Conversely, if the first sensing operation reveals that the cell voltage is below $V_{DD}/2$ then $V_{DD}/6$ is chosen as the LSB reference. The results of the two sensing operations produce the MSB and LSB values for the one addressed cell. At this point the data is latched at the sense amplifiers and ready for reading and writing. As in [2], these techniques involve exploiting the carefully matched sub-bitline capacitances.

The major advantage of [3] is that this circuit uses local components to do sense and restore operations. Specifically, the scheme uses local generation of reference voltages. This means that the reference signal is created using the cell which is being read, rather than special reference cells and power supplies as in [2]. Any gradual parametric differences across the chip will not affect the reference quality. All of this should lead to higher yield and greater reliability. Another redeeming quality is that the circuit is composed of proven standard DRAM sub-circuits. Thus the departure from proven techniques is limited to the operational and configuration of these rather than the optimized cell array.

A drawback to [3] is the time penalty implied by sequential sensing. There are also some required operations involving the charging of the bitlines between sensing and restoring that add time to the cycle. Although there are only two sense amplifiers per column, significant area is taken by the central switch matrix area. Another drawback is the complexity of the control logic required. In particular, the scheme calls for multiple transitions on the wordline and the dummy wordline (used for noise cancellation and capacitive balance). Multiple transitions may be diffi-

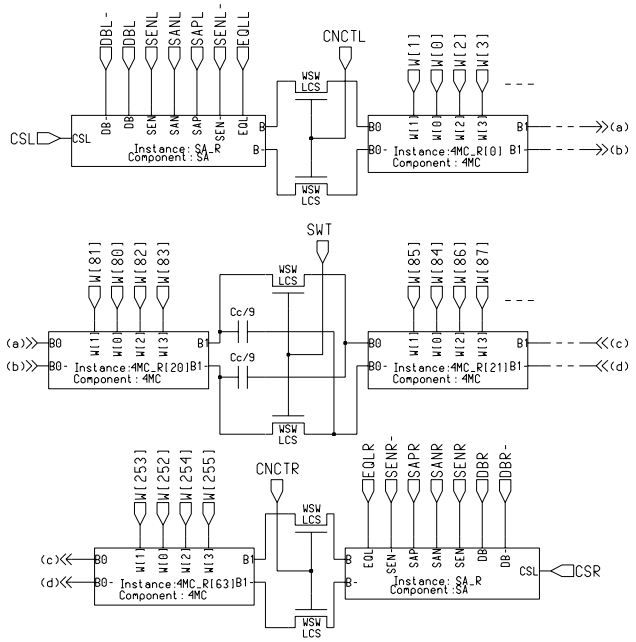

Figure 4. Schematic used for simulating the MLDRAM from [4].

cult to generate and may require larger drivers and charge storage areas for wordline boosting.

2.3. Okuda's Multilevel DRAM

The third MLDRAM scheme that we are considering was proposed by T. Okuda *et al.* [4]. For this experiment we are only interested in the multilevel sense and restore method so the bitline time-multiplexing and offset cancellation aspects of the design were left out. Although the amplifier arrangement has been demonstrated [6] to be an effective method of increasing device reliability, it could be incorporated into any of the schemes considered here.

As in [3], the cell data is extracted using sequential sensing. However, rather than using local charge sharing to generate the LSB reference voltages, a capacitive coupling method is used. The value of the most significant bit is placed onto a capacitor which bumps the second reference value up or down from the middle reference to one of the two possible LSB references. The restore scheme is similar to the previous two, but the architectural configuration is not inherently balanced so great care is required to ratio the bitline capacitances properly.

The obvious advantage to this scheme is its simplicity. The extra circuitry needed is minimal and the control timing is relatively simple.

Perhaps the most vulnerable aspect of this design is the use of the coupling capacitors to produce the LSB reference. These capacitors must be exactly proportioned to the cell capacitance since any error in the ratio translates directly into a reduction in the noise margins.

3. A New MLDRAM Scheme

It is possible to combine the fast access advantage of [2] with the local reference generation ideas of [3]. We propose a new MLDRAM in which, like [2], each bitline pair is divided into three equal-length sub-bitline pair segments, where each segment is provided with a sense-amplifier circuit. This allows fast, single-step flash-conversion sensing. Instead of using globally-generated reference voltages for sensing, the new scheme uses charge-sharing techniques between three adjacent sub-bitlines, reminiscent of [3], to locally generate the three necessary reference levels.

Figures 5 and 6 illustrate the new MLDRAM. Figure 5 shows the two different sub-bitline pair configurations used and Figure 6 shows how nine sub-bitline pairs are organized as a 3-by-3 array. The position of each sub-bitline pair is identified by a horizontal co-ordinate {L, C, R} and a vertical co-ordinate {T, M, B}. The sub-bitlines can be connected together horizontally, L to C to R, via switches controlled by signals SWT0 and SWT1. In addition, the sub-bitlines can be connected vertically, T to M to B, via switches controlled by signals REF0 and REF1.

The M sub-bitline pairs differ slightly from the T and B sub-bitline pairs in connections made to generate word lines GW0 and GW1 and reference wordlines RW0 and RW1. These special wordlines are identical to normal wordlines in all respects except that the cell access transistors for GW0 and GW1 are missing in the M sub-bitline pairs (following configuration SBL_RW in Figure 5), and the cell access transistors for RW0 and RW1 are missing in the T and B sub-bitline pairs (following configuration SBL_GW in Figure 5).

The sub-bitline pairs also differ in connections made to transistors controlled by the signal GEN. This signal is used to generate the reference voltages after a restore in preparation for the next sensing cycle. As shown in Figure 6, signal GEN connects the sub-bitlines to either V_{DD}, V_{SS} or V_{BLP}. The organization of these connections used is critical, though there are some possible variations. The signal GEN connects the sub-bitlines TL and BL to V_{SS}, TR and BR to V_{DD}, and all others to V_{BLP}.

Reference cells are provided in the ML, MC and MR sub-bitline pairs. Consider ML first. Sub-bitlines in TL, ML, and BL are precharged separately to V_{SS}, V_{BLP}, and V_{SS}, respectively, by asserting the GEN signal. To ensure that the capacitances of all sub-bitlines are equal, the RW0, RW1, GW0 and GW1 signals are all asserted causing each sub-bitline to have the same capacitance equal to the parasitic capacitance of the bitline (C_B) plus one memory cell (C_C). After charging, the sub-bitlines are shorted together

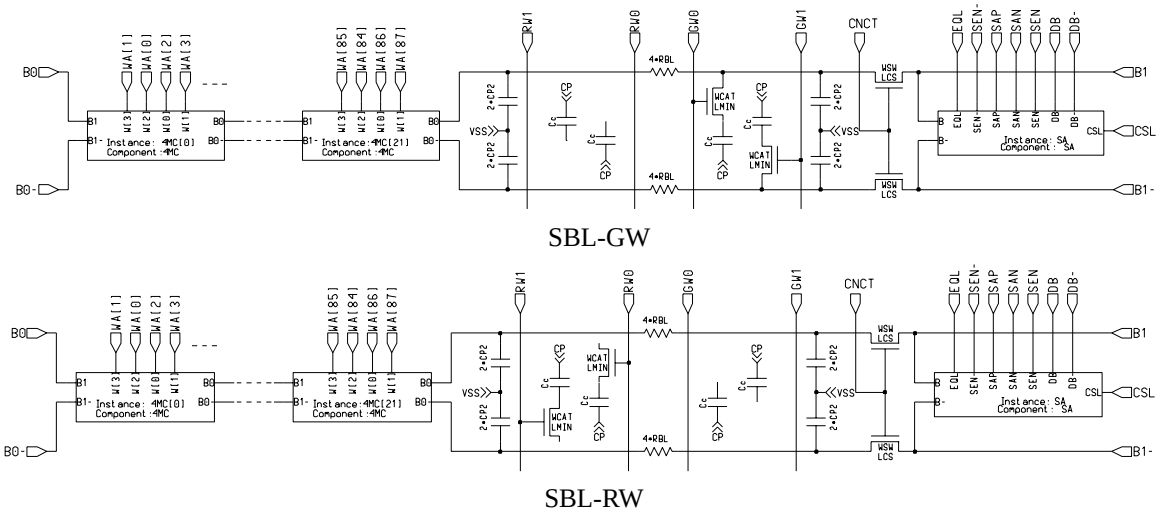

Figure 5. “Generate” and “Reference” sub-bitlines for the proposed MLDRAM scheme

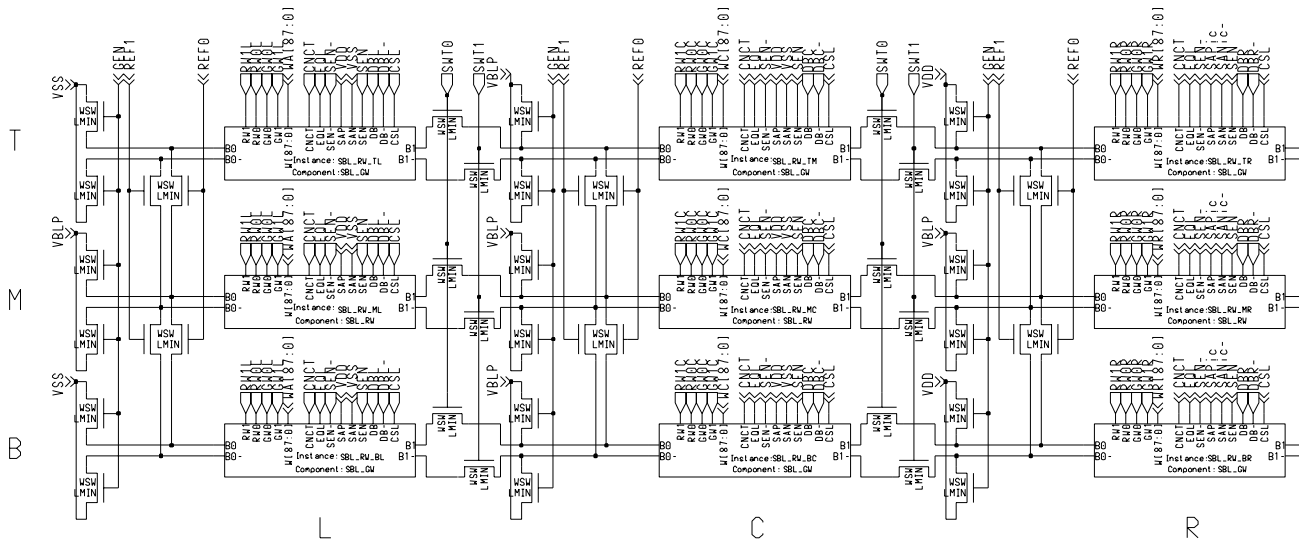

Figure 6. Schematic for the proposed MLDRAM scheme

(T to M to B) creating the final voltage of $V_{DD}/6$. The same operation is performed on the C and R groups of sub-bitlines but the charged values for these cause the resulting voltage to be $V_{DD}/2$ for TC, MC, and BC and $5V_{DD}/6$ for TR, MR, and BR. Having created the three required reference voltages, the reference voltages are stored in the reference cells by de-asserting signals RW0 and RW1 in L, C, and R.

At the start of the sense/restore cycle all sub-bitlines are shorted together through REF0, REF1, SWT0 and SWT1, and precharged to V_{BLP} through the precharge devices in the sense amplifiers. Consider a read to a cell on an even wordline WA[0] in TL. In Figure 7 below it can be seen that the cell signal will be shared with even bitline BL0. Thus an appropriate reference cell signal must be shared with the odd bitline (BL1). After disconnecting all of the sense amplifiers from the sub-bitlines through signal CNCT, the cell signals and reference cell signals are distributed as follows. Referring to Figure 6, the odd sub-bitlines are disconnected in the horizontal direction by de-asserting SWT1, leaving the even sub-bitlines connected horizontally. At the same time, the even sub-bitlines are disconnected from each other in the vertical direction by de-asserting REF0, leaving only the odd sub-bitlines shorted vertically. After this operation the wordline WA[0] and the reference wordlines RW1L, RW1C and RW1R are asserted. The cell signal is shared horizontally along three even sub-bitlines for each of T, M and B, and the reference cell signals are shared vertically along three odd sub-bitlines for each of L, C and R. Each cell signal and each ref-

erence cell signal is diluted onto the same bitline capacitance (three sub-bitlines) and so the resulting differential signals are correct. The final step in sensing is to isolate all sub-bitlines and then connect and power up the sense amplifiers.

Restore is performed using charge sharing among L, C and R sub-bitlines in essentially the same way as [2]. Following restore the reference voltages are generated and stored as described above.

4. Experimental Setup

The simulations were performed using a process based on a composite of typical 250 nm CMOS processes. For each scheme, a minimum of 256 wordlines and one bitline pair were modeled (three were required in the new scheme). For a control case, a DRAM bitline pair was modeled and simulated. This can be used to compare the various advantages and disadvantages of MLDRAM in general.

Ideal capacitors were used for the cells and a distributed RC model was used to model the bitline parasitic resistance and capacitance. For all of the schemes, the same organization of memory cells was used, as can be seen by the symbol "4MC" in Figures 2, 3, 4 and 5. The basic cell array model is shown in Figure 7. This fundamental unit of cells requires that schemes [2] and the proposed scheme have greater than 256 wordlines since the number of sub-bitlines is not a factor of 256. It should be noted that the C_{GS} of the cell access transistors and the switches account for most of the bitline capacitance. C_{P2} is added as a way to control the cell-to-bitline capacitance ratio and is intended to only loosely resemble the bitline poly capacitance encountered in a DRAM process.

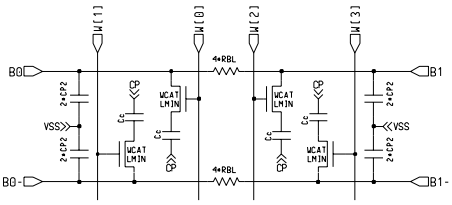

Figure 7. The model for four memory cells (4MC).

The same sense amplifier model was used for all simulations, as can be seen in the figures above. Each scheme uses a generic DRAM sense amplifier as a building block so this is a parameter that can be controlled without biasing the performance of any scheme. The sense amplifier used in the simulations is shown in Figure 8. To get realistic sensing times in a DRAM environment, a resistor was added in series with each of the SAP and SAN- voltage sources to provide transient behavior similar to a worst case in a conventional DRAM.

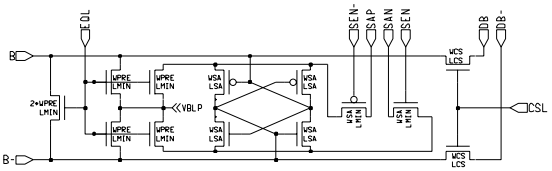

Figure 8. The model used for the sense amplifiers.

The other controlled parameters are given in Table 2. The resistances, capacitances, and transistor lengths and widths pertain to Figures 2-8.

The SPICE simulations were performed as described in their respective papers. Wherever possible, timing waveforms and circuit configurations are true to their paper descriptions, but some modifications were required to fit the schemes into the experimental setup. For example, the sense amplifier in [2] is slightly different from the one that we used (Figure 8), but both work for the MLDRAM scheme.

For MLDRAM, as in regular DRAM, read and write operations occur once the sense amplifiers are turned on. The modeling and operation of the data-bus and column access was not included in the simulations. Each scheme could be equipped with the same data-bus interface and it is assumed that the use of this interface is not crucial to MLDRAM operation.

Table 2: Simulation Conditions

Parameter	Description	Value
V_{DD}	Main supply voltage	2.50 V
V_{BLP}	Bitline precharge	1.25 V
V_{CP}	Cell capacitor common node	1.25 V
V_{BOOST}	High value for boosted NMOS switches	3.50 V
T	Temperature	27 C
t_r, t_f	Control signal rise and fall time	2.0 ns
t_w	Wordline rise and fall time	5.0 ns
t_{skew}	Time allowance between edges	0.5 ns
RBL	Bitline resistance per cell	25 ohms
C_{P2}	Bitline poly parasitic capacitance per cell	100 aF
Cc	Cell capacitance	25 fF
Cb/Cc	Full Bitline to cell capacitance ratio	~5.33
WCAT	Cell access transistor width	300 µm
LMIN	Transistor minimum length	250 µm
WSA	Sense amplifier transistor width	2.0 µm
LSA	Sense amplifier transistor length	0.5 µm
WSW	NMOS switch width	1.0 µm
WPRE	Precharge transistor width	0.5 µm

Simulations were performed in two steps. First, each scheme was simulated with relaxed timing to allow nodes to reach steady state between switching operations. After the steady state values were determined, the timing was optimized by compressing the operations in time as much as possible without compromising MLDRAM operation or

signal quality. The criteria used to optimize the timing are shown in Table 3. Wherever a charge sharing operation was performed, the worst case of the nodes involved was allowed to meet the criterion before the next step of the critical timing path occurs. On top of this convergence time allowance, a constant delay, t_{skew} (Table 2), was added between each edge to accommodate any skew between signal edges that might occur in a real chip.

Table 3: Timing Optimization Criteria

Switching operation	Criteria
Switch turning off	None (only rise/fall time + t_{skew})
Charge sharing	99% of their steady state value
Sensing	Within 5mV of the power supply voltage
Precharging	Within 1mV of V_{BLP}

5. Results

Our results are summarized in Table 4. Sense amplifier (SA) activation is the time when the SEN signal (Figure 8) begins to rise. The bitline (BL) 0.5V split is the time it takes the sense amplifier to cause a differential voltage of 0.5V to develop on the bitlines (or sub-bitlines). This parameter is important because the different sensing operations take different times and the 0.5V split gives a fair indication of the delay required for a read operation. For [3] and [4] the MSB and LSB are available at different times so there are separate rows in Table 4 for the two bits. The cycle time is measured from the first event to the time when the final precharge criterion (from Table 3) is met.

Table 4: Simulation Results

Parameter	DRAM	[2]	[3]	[4]	New
SA activation, 1st bit (ns)	8.00	14.06	18.33	15.27	14.25
BL, 0.5V split, 1st bit (ns)	9.08	16.73	20.16	17.00	17.03
SA activation, 2nd bit (ns)	N/A	14.06	70.25	34.04	14.25
BL, 0.5V Split, 2nd bit (ns)	N/A	16.73	72.50	35.40	17.03
Cycle time (ns)	37.87	57.75	152.48	82.09	77.74
Energy per bit (fJ/bit)	809	1101	1490	767	1605

An example of the sense and restore waveforms from the simulations is shown in Figure 9. All bitlines with all possible data values are plotted for each scheme, including the one-bit DRAM case. Each plot shows the full sense and restore scheme. Unlike the others, the new scheme is shown starting with the restore, then precharging, and finishing with the sensing. This was done to avoid assuming perfect reference cell signals at the start. Rather, the reference cell signals are generated and captured in the reference cells. The precharge follows just before the actual start of the sense and restore cycle at t=45ns.

From examining the waveforms one can see that for [2] and the new scheme there is only one sensing operation, indicative of the flash-style sensing. For [3] and [4] the two sequential sensing operations can be seen.

Another observation is that the quality and placement of the restore values varies somewhat. This is an indication of the degree to which the sub-bitlines used in the restore are matched and balanced, and the degree to which charge injection effects are cancelled.

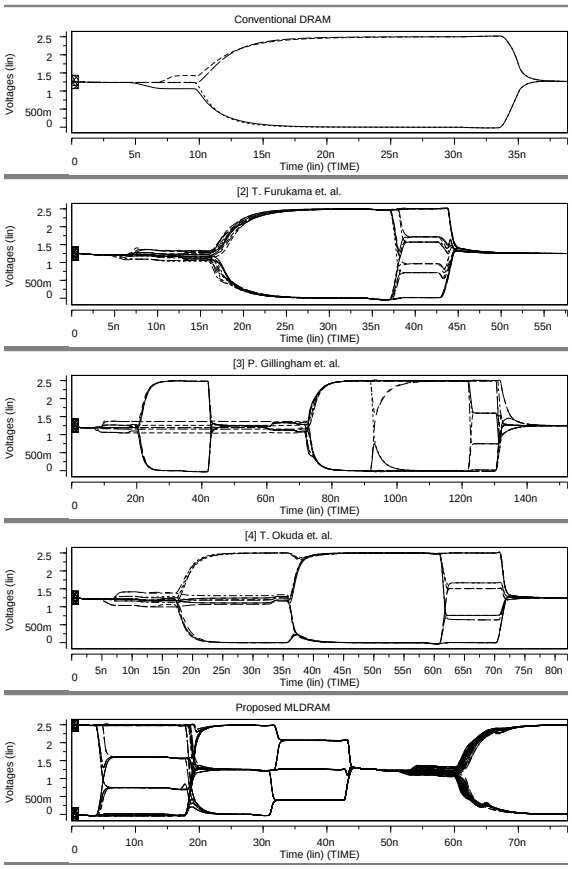

Figure 9. Sense and restore waveforms.

6. Discussion

The quantitative and qualitative aspects of the MLDRAM schemes are compared briefly and summarized below in Table 5, 1=Best Case and 5=Worst Case.

As compared with conventional two-level DRAM, all MLDRAMs suffer in terms of speed. From the results it is clear that the parallel sensing used both in [2] and the proposed method provide the fastest access and cycle times.

The energies per bit are all higher than for DRAM. The reason for this is as follows. For [2] and the proposed method there are three sense amplifiers that operate per cycle. For each sense amplifier, significant crowbar current leaks from V_{DD} to V_{SS} during the first few moments of sensing, thus increasing the energy used despite the fact that approximately the same capacitance is being charged to the same voltage as in a DRAM. This is also true for the other MLDRAM schemes, but they have fewer sense

amplifiers leaking charge to ground. The other reason for increased energy per bit is the use of control switches. The energy used to switch devices on and off was significant, especially for those signals which needed to be boosted above V_{DD}.

The energy of [3] is higher than [2] because there are more charging operations per cycle. Specifically, during restore in [3] one sub-bitline is completely charged from V_{SS} to V_{DD}. As well, the control switching is the most complex so the energy required for this is highest. The energy of the proposed method is higher than [2] because the reference generation procedure requires extra energy to charge the required bitlines to V_{DD}. The energy of [4] is actually lower than that of DRAM because it has minimal extra control circuitry and sparse use of boosted control signals in order to recover two bits.

In terms of area, [4] has the least increase over DRAM because it uses a minimal amount of circuitry and only two sense amplifiers. [3] is a close second to [4] because of the extra area used by the switch matrix. [2] and the proposed scheme have the worst overall area gain because they use three sense amplifiers per bitline.

In terms of robustness, the new scheme is ranked the best among the MLDRAMs because of the use of local reference generation and local charge sharing based restore. In terms of signal quality it is better than [3] because it has balanced control switch usage, which means charge injection errors are not a concern. [2] has good robustness because it uses charge sharing for the restore voltage, but it may suffer from the accuracy and impedance of the voltage supplies that are required to charge the reference cells. [4] ranks last because, although it uses charge sharing, LSB sensing depends heavily on the capacitor used to generate the reference. The exact size of this capacitor is critical and this raises reliability concerns.

The final criterion on which the MLDRAMs are compared is the control complexity. In general, the more complex the control timing, the more room there is for skew errors that affect the soft error rate. Furthermore, the more control logic that is required, the more area and power is consumed when generating the signals and driving the long control wires. The ranking shown in Table 5 is given according to the number of edges required for the cycle.

Table 5: MLDRAM Ranking

Attribute	DRAM	[2]	[3]	[4]	New
Access time	1	2	5	4	3
Cycle time	1	2	5	4	3
Cycle energy per bit	2	3	4	1	5
Overall robustness	1	4	3	5	2
Area per bit	5	3	2	1	4
Control logic simplicity	1	3	5	2	4

7. Conclusions

MLDRAM has been proposed by several authors as a way of increasing storage density without reducing the physical feature size. As a result of the disadvantages of reduced noise margin, increased circuit complexity and longer access and cycle times, MLDRAM has yet to enter production parts. The simulation study presented in this paper is a step towards fulfilling the promise of MLDRAM by attempting to clarify the real strengths and weaknesses of proposed MLDRAM schemes. Another contribution of this paper is the introduction of a new MLDRAM scheme that reduces the access time disadvantage of [3] while retaining the advantages of charge-sharing restore and local reference voltage generation.

We believe that the results of our simulation study are sufficiently promising to justify the implementation of test chips. Samples of a small test chip for [3] were available, but not yet tested, at the time of writing. We also plan to implement our new MLDRAM. Characterization of the test chips should help resolve remaining concerns, such as soft error rates, noise margins, and refresh requirements.

8. Acknowledgments

This research was supported by Micronet R&D, MOSAID Technologies Inc., the Canadian Microelectronics Corporation, the Natural Sciences and Engineering Research Council of Canada and the University of Alberta. We thank Valerie Lines, Randy Torrance, Roger Colbeck, Dan Clein and Peter Gillingham of MOSAID Technologies Incorporated for their technical and financial support of MLDRAM research at the University of Alberta. We would also like to thank Ralph Nevins for his invaluable help in keeping our CAD tools and workstations running smoothly.

9. References

[1] M. Aoki et al., “A 16-Level/Cell Dynamic Memory,” IEEE J. Solid-State Circuits, Vol. SC-22, No. 2, pp. 297-299, April 1987.

[2] T. Furuyama et al., “An Experimental 2-bit/Cell Storage DRAM for Macrocell or Memory-on-Logic Application,” IEEE J. Solid-State Circuits, Vol. 24, No. 2, pp. 388-393, April 1989.

[3] P. Gillingham, “A Sense and Restore Technique for Multilevel DRAM,” IEEE Trans. On Circuits and Systems-II: Analog and Digital Signal Processing, Vol. 43, No. 7, July 1996.

[4] T. Okuda et al., “A Four-Level Storage 4-Gb DRAM,” IEEE J. Solid-State Circuits, Vol. 32, No. 11, pp. 1743-1747, Nov. 1997.

[5] B. Prince, “Semiconductor Memories: A handbook of Design Manufacturing & Applications,” 2nd Ed., John Wiley & Sons, 1996.

[6] T. Kawahara et al., “A High-Speed, Small-Area, Threshold-Voltage-Mismatch Compensation Sense Amplifier for Gigabit-Scale DRAM Arrays,” IEEE J. Solid-State Circuits, Vol. 28, No. 7, pp. 816-823, July 1993.