Single-Bit-Errors

A Memory Module Supplier's perspective on cause, impact and detection

Larry C. Alchesky, Sr. Manager, Memory Engineering
Smart Modular Technologies, Inc.

INTRODUCTION

DRAMs (Dynamic Random Access Memories) play a key role for computers, servers, routers, workstations and other applications requiring a high-density highly reliable cost effective main memory solution. These applications have been and continue to be plagued by Single-Bit-Errors (SBE) and Multi-Bit-Errors (MBE) that reduce the reliability of those applications. IBM[1] had reported that 98% of failures in memory systems are related to SBE. Although MBE can and do happen the probability of them occurring is much more unlikely.

The occurrence of these SBE and MBE failures have required system designers to use various forms of data protection – parity, ECC and Chipkill. The additions of these data protection schemes have improved the memory system reliability but they have come at an increase in cost and complexity.

A system memory failure occurs when the memory system fails to do what its been designed for. The failure can manifest in the form of a permanent or non-permanent fault[2]. A permanent (or hard) fault won't normally need any special electrical test to identify it, as it will easily fail most write/read operations to the memory. A non-permanent fault can be harder to find as they can occur randomly or require unique conditions, making the detection and localization of the fault difficult.

As defined by van de Goor[2], a non-permanent fault is a non-destructive fault and falls into two categories:

1. Transient Faults (also called Soft Error) caused by environmental conditions – temperature, voltage, humidity, pressure, vibrations, power supply fluctuations, electromagnetic interference, ground loops, cosmic rays, and alpha particles.

2. Intermittent Faults caused by non-environmental conditions – loose connections, aging components, critical timing, resistive or capacitive variations and noise in the system

As the reliability continues to increase in memory devices, the 'hard error' rate has continued to decrease with DRAM vendors reporting 10FIT (FIT = Failure In Time, where time is 10^9 device-hours) device hard failure rates. Therefore, a dominant failure mechanism in memory devices will be Transient Faults, which based on reported soft-error rates can be one to two orders of magnitude higher than the hard failure rates. With typical soft-error rates being in the 100-500 FIT range. These transient faults will reveal themselves as Single-Bit-Errors (SBEs) during testing and/or application execution. These SBEs will impact the chip and module manufacturer directly in the cost of manufacturing and testing the memory devices and modules respectively. The system designer will be affected by the potential of SBEs and the need to take precautions for detecting and correcting these errors, increasing the cost of manufacturing the system.

SBE's cause has been attributed to latent defects that manifest themselves into failing bits after being exposed to temperature and voltage swings during device and/or module assembly as well as device, module or application testing. The absolute cause of a given SBE can be unclear as they can pass all the extensive testing the DRAM and Memory Module suppliers do to detect them and still show up in the application testing done by the customer. However, no matter what the cause, the failures have one thing in common - ***the memory cell's ability to store and retain charge.***

DRAM Memory Cell Review

The DRAM memory cell (Figure 1) is a simple access transistor and storage capacitor structure whose stored data is dependent on how many electrons are stored in the capacitor to represent a 0 or 1. This charge will bleed off over time and needs to be refreshed at a regular interval to retain the correct data. For the data to be retained over the various conditions allowed by the data sheet each cell must have margin to the required refresh rate. This margin can vary from memory cell-to-memory cell, from DRAM supplier-to-supplier and from technology-to-technology.

DRAMs store information (0, 1) as the presence or absence of electron charge on a capacitor. This stored charge decays away over a period of time do to junction and dielectric leakage and must be periodically refreshed to retain the correct information. To replenish the loss of charge, before the correct data is lost, a refresh operation must be done. Refresh is defined as an operation of restoring the charge on the memory cell capacitor. Refresh is specified in the memory device's data sheet by two components -- the numbers of refresh cycles and the time interval in which the refresh cycles must be executed. For example, the 128M x 4 DRAM requires 8192 refresh cycles to be executed every 64ms to insure the DRAM maintains correct data.

Refresh is accomplished by accessing each row (one at a time) in the memory array. When a row is accessed, the word-line goes high and the memory cell data is transferred to the bit-line. The sense amplifier then amplifies the voltage difference on the bit-line. This operation refreshes (rewrites) the charge in the memory cell to its initial state.

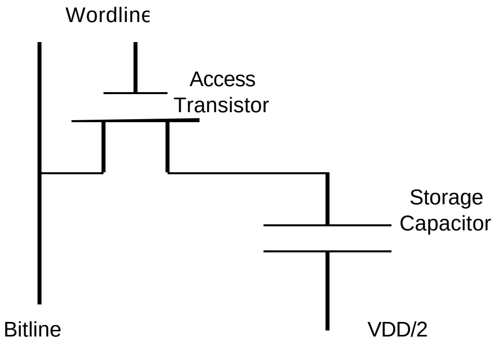

DRAM MEMORY CELL

Figure 1

May and Woods[3] defined that the amount of charge that distinguishes a 1 from a 0 can be measured in the number of electrons stored on the memory cell capacitor. The amount of charge that differentiates between a 1 and 0 was defined as the "critical charge", Q_{crit}. The charge (Q) stored on a capacitor can be calculated as Q=CV. With cell capacitances (C = farads) typically ranging from 30-40fF and the voltage (V = volts) across the cell typically being VDD/2, the amount of charge (for VDD = 5v) ranges from 75-100fC (femto coulombs). This translates to 468,000 to 624,000 electrons being stored in the memory cell capacitor. With the trend of VDD being lowered with each new introduction of DRAMS – 3.3v (SDRAM), 2.5v (DDR1) and 1.8v (DDR2) the amount of charge is reduced even if the capacitance is maintained (see Table 1).

VDD	VDD/2	Q =CV (femto-coulombs)		# Electrons stored (x1000)	
volts	volts	30fF	40fF	30fF	40fF
5.0	2.5	75.0	100	468	624
3.3	1.65	49.5	66	309	412
2.5	1.25	37.5	50	234	312
1.8	0.9	27.0	36	168	225

Table 1

It should be noted this is only an estimation of the charge stored as DRAMs incorporate an internal VDD level that may be somewhat higher (or lower) than the estimate of VDD/2.

Figure 2 illustrates that although the capacitance of the memory cell is being maintained the amount of charge being stored is decreasing do to the reduction of voltage levels. This reduction in charge will make the memory cell more susceptible to upset and/or loss of stored data. This susceptibility may reveal itself in the form of increased single bit errors.

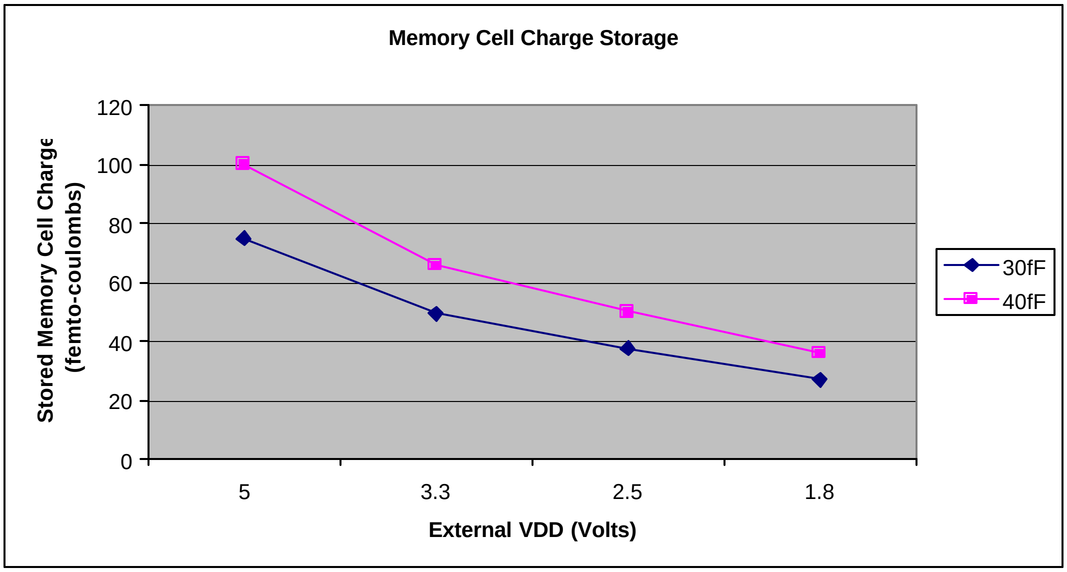

Figure 2

In addition, to the decay of charge over time, the memory cell can be susceptible to loss caused by the electron-hole pairs generated by alpha particles or cosmic rays. The amount of charge stored can be minimized do to latent process defects that impact the ability to store maximum charge or can cause rapid depletion of the stored charge through abnormal leakage paths. Or the charge can be depleted, minimized or affected by external conditions – temperature, voltage, refresh rates, timing, data patterns, etc.

Probability of a Single Bit Error

To determine how much effort (time and money) should be invested in providing the right amount of error detection and protection, the level of reliability of the individual DRAM component and the memory system must be understood. By determining the Mean Time Between Failures (MTBF) for the memory system under design we can determine what the SBE impact may be to the memory system. The following steps can be used to calculate the MTBF for a memory system.

Step 1. Calculate the probability of one single bit error occurring in any bit in the memory device in an hour.

Step 2. Calculate the probability of one single bit error occurring in the memory system under consideration.

Step 3. Calculate the MTBF for the memory system.

Step1. Probability of Single Bit Error

To determine the probability of a single bit error in a memory device, the following formula can be used:

Equation (1)

$$p(1SE) = \frac{FR}{BPD}$$

Where p(1SE) is the probability of a single bit error in a single memory device, FR is the Failure Rate in errors/device hour (FITs/10^9 device-hours), and BPD is Total Bits Per Memory Device.

Example 1: For a 512Mb DRAM (536,870,912 bits) with an assumed 100 FIT Failure Rate for transient or soft-error faults - what is the probability of just one single bit error occurring during one hour of operation? Using Equation (1).

$$p(1SE) = \frac{100/10^9}{536{,}870{,}912}$$

$$p(1SE) = 1.86x10^{-16} \text{ errors/bit - hr}$$

Step 2. Probability of a Single Bit Error in the memory system

Using the binomial distribution the probability of a single bit error occurring in a memory system can be calculated using Equation (2).

Equation (2)

$$P(X = x) = \left(\frac{n!}{(n-x)!x!}\right)\left(p^x\right)(1-p)^{n-x}$$

Where n = number of bits in the memory system, x = the number of errors in the memory system and p = the probability of having one soft error in the memory device.

Equation (2) can be simplified by noting we want to calculate the probability of just one error occurring in the system, P(X=1); the probability (p) of just one error occurring in any single memory device is usually very small (p<<1); and applying the following identity for factorials (n! = n(n-1)!). Equation (2**)** simplifies to become Equation (3**)**.

Equation (3) $$P(X = 1) = np$$

Where n = the number of bits in the system (bits/device x devices/system) and p = p(1SE) or the probability of a single bit error occurring in any one device.

Example 2 For a module with a 72 bit wide word, containing 1GByte of memory, using eighteen 128Mx4 DRAMs having a 100 FIT soft-error rate/device, what is the probability of a single bit error occurring in one hour of system operation? Using Equation (3).

$$n = \left(\frac{536{,}870{,}912 bits}{device}\right)\left(\frac{18 devices}{system}\right)$$

$$n = 9{,}663{,}676{,}416 bits / system$$

And from **Example 1**,

$$p(1SE) = 1.86x10^{-16} \text{ errors/bit - hr.}$$

The probability of single bit error occurring in the 1GB module is:

$$P(X = 1) = (9{,}663{,}676{,}415)(1.86 \times 10^{-16})$$

$$P(X = 1) = 1.80 \times 10^{-6} \text{ errors/system hr.}$$

Step 3. **Calculate the MTBF**

To understand the impact on the system, the probability of a single bit error in a system can be converted to the Mean-Time-Between-Failures (MTBF) using the following equation:

Equation (4) $$MTBF = \frac{1}{P(X = 1)}$$

For the imaginary 1GB system, using 128Mx4 DRAMs, and applying Equation (4) the MTBF would be:

$$MTBF = \frac{1}{1.8 \times 10^{-6} errors / hr}$$

MTBF = 5.56 x 10^5 hrs/error

Or every 23,148 days/error, or every 63.4 years/error.

Figures 3 shows the probability of a single bit error for a single device (512Mbit) for different Failure Rates. The higher the failure rate the more likely a single bit failure will occur. So it's important to use DRAMs that have the lowest possible Failure Rate.

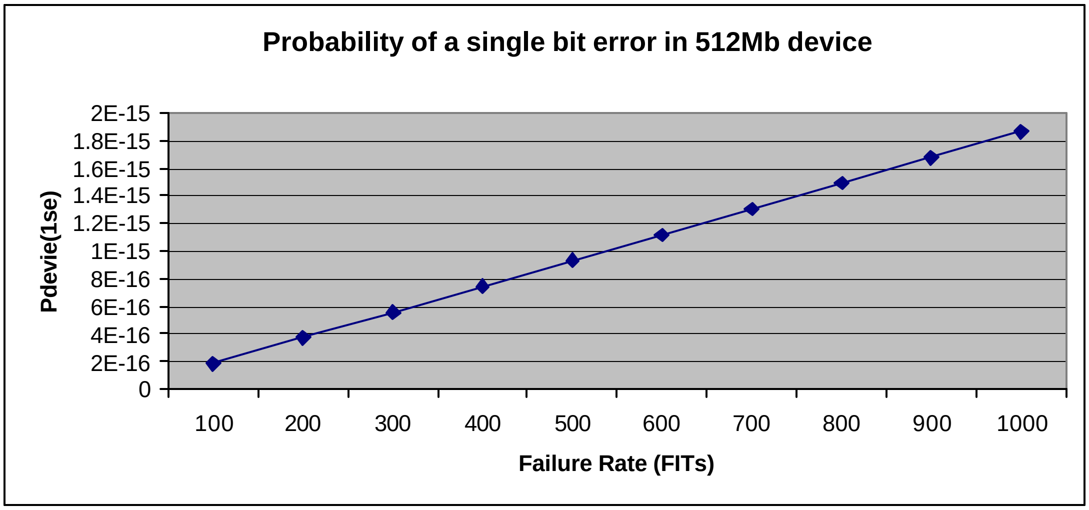

Figure 3

Figure 4 shows the probability of a failure in a 1GB memory module using eighteen 512Mbit (128Mx4) DRAMs. The failure rate for a module is the sum of the individual failure rates of the ICs. For example if the Transient Failure Rate is 100 FITs/device the Failure Rate for the module is then 1800 FITs (100 FITs x 18).

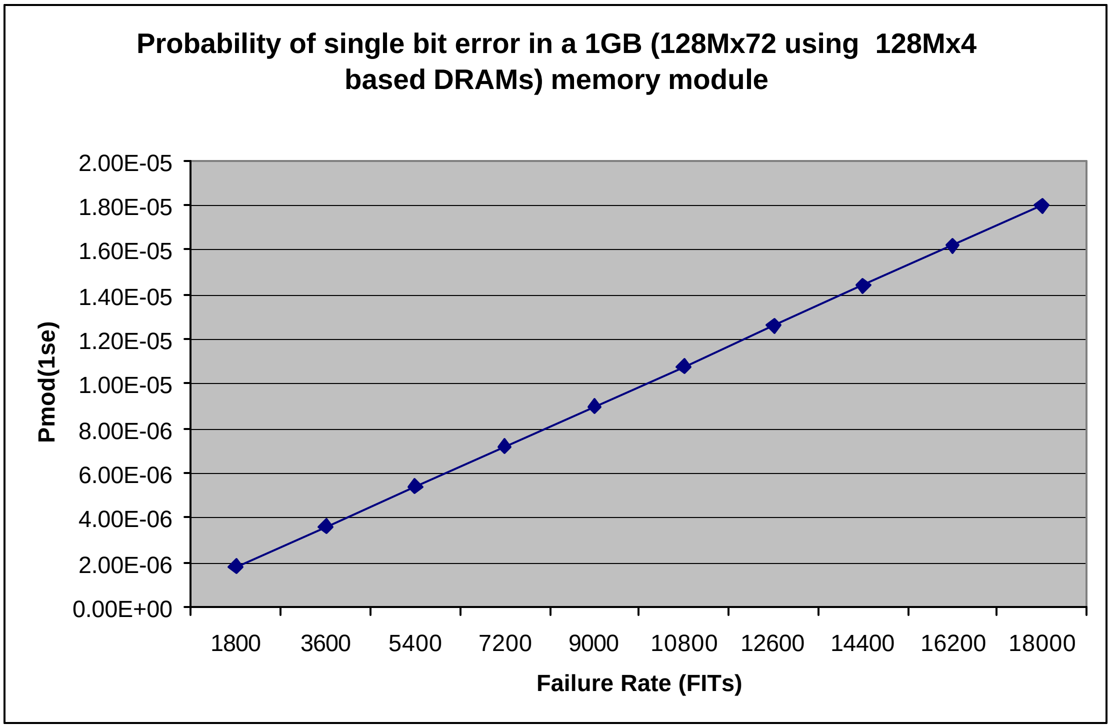

Figure 4

Figure 5 shows the probability of a single-bit error in various memory system configurations (1, 2, 4, 8, 16, 32 DIMMS) using the 1GB module made up of the 512Mb devices. As expected the more memory you have the more likely a single bit error event can occur.

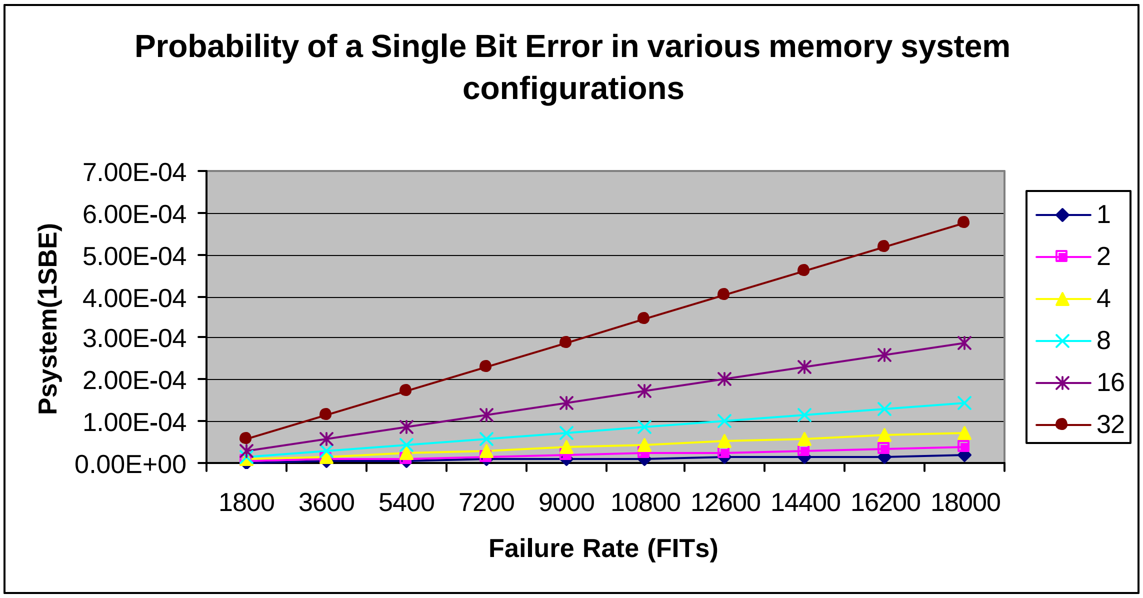

Figure 5

These probabilities of a Single Bit Errors (SBE) are better understood from a MTBF point of view. The more likely the SBE and the more memory being used in a memory system results in a quicker MTBF. Figure 6 shows the impact of memory system density and memory module failure rate on the time to failure.

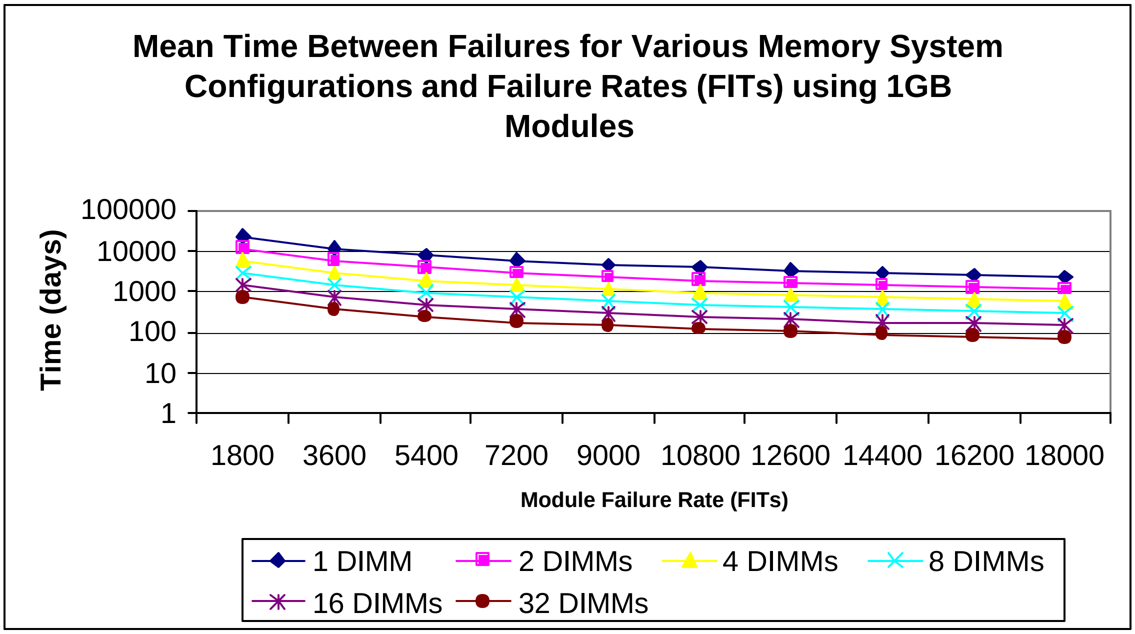

Figure 6

The Enemy – Weak Bits[4]

Figure 6 indicates the chance of failure is remote or at least far and few in occurrence depending on the module failure rate and system memory density. ***Then why is the single bit error phenomenon more prevalent than the theoretical calculations suggest?***

Memory failures don't always fit the classical hard or soft failure definition and therefore it may not be easy to assign a specific (known) failure rate to it to properly predict the occurrence of a SBE. A third class of failure exists, called the "weak bit". It behaves like a hard error (repeatable) and displays characteristics that are very similar to a transient or soft error by failing only specific conditions (VDD, timing, data pattern, temperature, I/O loading, etc.) while passing others.

What causes these 'Weak Bits' and how do the DRAM and Memory Module suppliers miss them during test? Weak bits' are a symptom that can be detected, if we come up with the right test conditions. This might be accomplished by measuring the data retention time (the time a memory cell can hold charge without refreshing) or it might be done with the right timing and data patterns, or it may require voltage and temperature margining or a combination of any and all of these conditions. What we can say is that these 'weak bits' are the proverbial needle in the haystack but harder to find.

The DRAM suppliers try to eliminate the weak bits through a process of burn-in. Burn-in is a process of applying elevated temperature (usually 125C) and voltage (Vdd maximum) during dynamic DRAM IC testing for screening or eliminating marginal devices, those with inherent defects or those with defects resulting from manufacturing aberrations which cause time and stress dependent failures[5]. Without burn-in these failures would show up as infant mortality or early lifetime failures.

Since all DRAM vendors do 100% dynamic burn-in these infant mortality and early lifetime failures should be eliminated. However based on customer's reports of failures when using memory modules in their applications, another defect must be present that is not being detected during burn-in or electrical test by the DRAM suppliers. These undetected failures are called latent defects, which are existing but not detected at the threshold of DRAM burn-in and electrical test used by the DRAM suppliers.

Latent defects that are not caught at burn-in can and do continue to degrade over time do to thermal stresses (like during SMT assembly) and/or do to electrical stress (during customer testing and use).

This failure rate should be low but can be similar to other measurable failure types – hard and soft errors. With typical failure rates for hard failures being reported as 10FITs and soft-error failure rates being at least an order of magnitude higher >100 FITs, the failure rate on a module can be significant since the failure rate of the module is the sum of the individual device failure rates. For example a module made up of 18-DRAM components would have a hard failure rate of 180FITs (18devices * 10 FITs) and a soft-error failure rate of 1800FITs (18 devices * 100FITs).

It has been the industry's experience that the latent defects show up as 'weak' bits being detected by system Error-Correction-Code (ECC) methods.

The fact that these "weak bits" do occur requires the Memory Module Manufacturer to provide a level of test coverage that will minimize the occurrence of these weak bits from exhibiting themselves as SBEs in the customers' applications.

Reports of SBE's

Customers' have reported to Smart the occurrence of SBEs in their applications under various conditions of temperature and voltage. The failures usually materialize during new product introductions or during production testing of the memory in the applications running customer specific memory diagnostics. The failures are typical ECC errors in that they are detected and correctable. Most of these errors occur at elevated temperature - >25ºC.

These high-temperature SBEs fit the normal data retention model where the higher the temperature the quicker the memory cell charge bleeds off through various leakage paths within the memory cell and data can be destroyed.

Figure 7 illustrates a typical DRAM memory cell's data retention behavior as temperature increases. The absolute data retention for a given memory cell will vary depending on DRAM technology and memory design. So supplier-to-supplier variability is very likely.

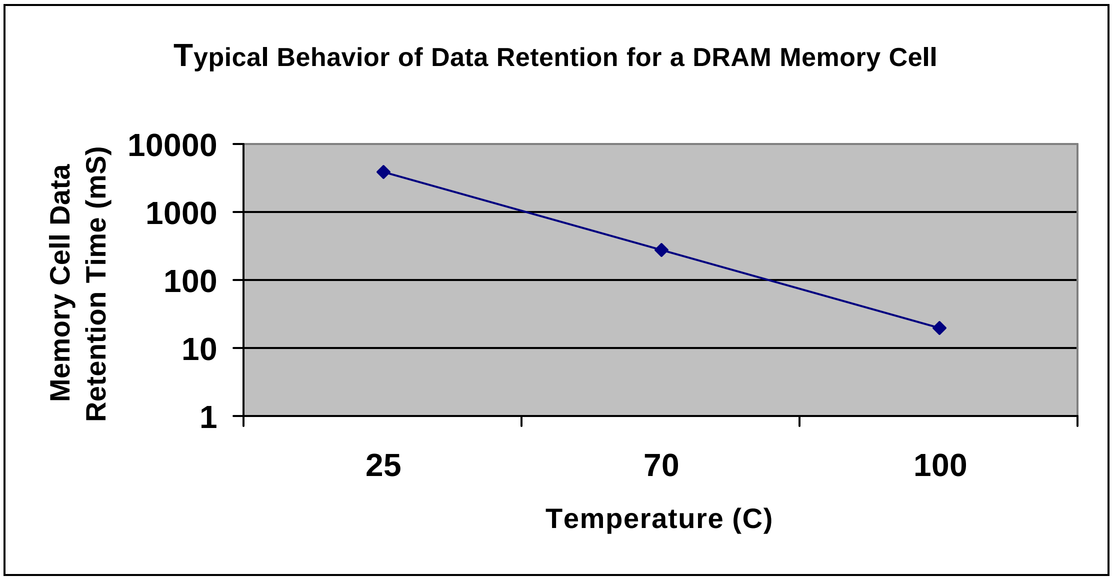

Figure 7

It has been Smart's experience that >95% of the failures reported (that are SBE in nature) fall into the high temperature category. This means the majority of application SBEs could be related to charge storage, which can be screened out at test by using higher temperature or using more severe test algorithms at room ambient temperature to stress the data retention. The later is preferred since high temperature test requires additional capital investment (temperature forcing equipment) and can also impact tester throughput (high temp testing requires time to ramp to the required temperature and some soak time to allow for equalization prior to testing).

Enterprise Class Memory Testing

Smart Modular Technologies has integrated into our production a testing methodology named Enterprise Class Memory Testing or ECM testing to detect the weak bit (SBE) and remove it from the product to be shipped before it can cause the customer a failure during application use.

It must be noted that not all weak bits can be detected and eliminated as this would become test time and cost prohibitive. However the goal is to reduce the customers' observed DPPM rate to a more acceptable level.

Production Test Flow

Smart's typical production test flow provides comprehensive test coverage using proprietary test software carried out on high-end Automated Test Equipment to insure the functionality and quality of all memory modules shipped. Figure 8 highlights the key testing parameters in a typical production test flow.

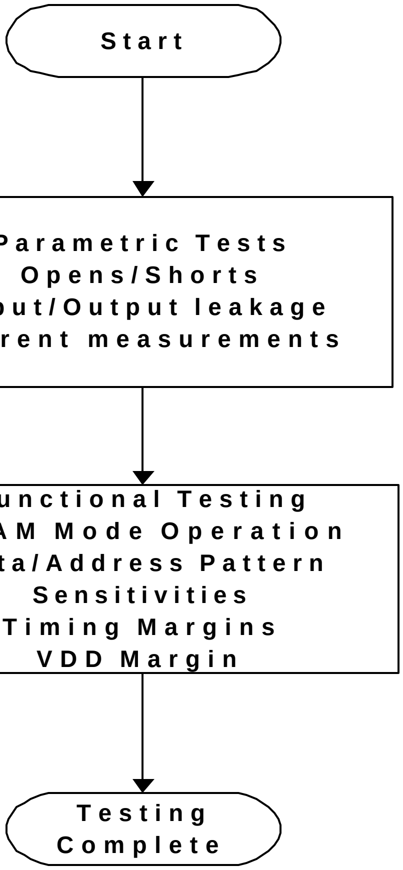

Typical Production Test Flow
Figure 8

ECM Test Flow

An ECM test flow is identical to a simple production test flow, except that more time is spent on detecting the weak bits (SBEs) through evaluating the devices VDD, timing and data pattern sensitivities. These sensitivities will vary from DRAM suppliers and technology. This requires Smart to develop specific test flows and conditions to address these supplier and device differences. The specific tests and conditions are confidential but it takes advantage of traditional and non-traditional tests.

Traditional March Tests[2]

- March A, B, C, C-, C+, LR, X, Y, MATS+, etc…
- These tests can be run with different data pattern, timing and VDD conditions
- These tests can be run with fast row or column address sequencing

The results of the traditional March tests are that they will catch most address, stuck-at, coupling, transition and data sensitivity faults. However they may not detect the weak bit.

To capture the weak bit one must apply non-traditional tests.

Non-Traditional Tests

- Disturb tests
  - Column/row disturb tests
  - Hammer tests (repetitive reads)
- Data retention tests
- These tests can be run with different data patterns, timing and VDD conditions
- These tests can be run with fast row or column addressing sequencing
- Temperature can be used to accelerate the time to failure

The area most impacted by the evolution of DRAM, particularly by the density growth, is the area of test. As Figure 9 shows, the test time is dependent on the density but it is also dependent on the complexity of test coverage. A typical Production Test Flow with limited coverage (for a 128MB module) may take ~10 seconds where as more comprehensive or ECM Test Flow can take almost two orders of magnitude longer to test the same memory density. It is this fine line that is walked by the module manufacturer to carefully balance test time with test coverage. Using proprietary test flows and patterns, Smart has developed the ability to thoroughly test our memory modules, providing our customers the best possible quality while balancing the test time and test costs.

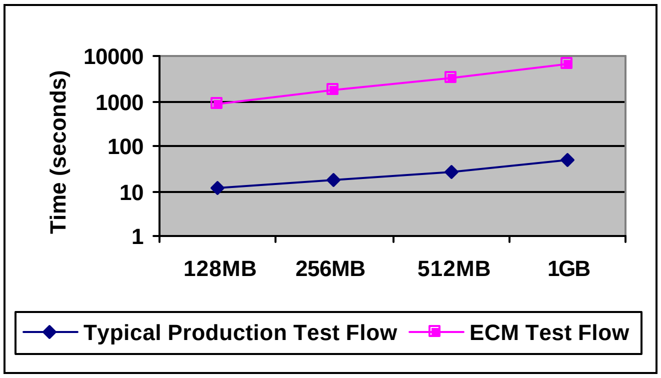

Figure 9

A Case Study of ECM Test Results

Working closely with an OEM customer an ECM test flow was developed to screen out potential system level failing modules containing weak bits. Initially the customer returned 45 1GB Registered DIMMs that had failed during system testing. The modules were tested using the standard outgoing production test program at room temp and then retested using an ECM test flow. Table 2 shows the results of that analysis.

# of Modules returned	# of modules failing standard production testing	# of modules Failing ECM testing	Comments
45	8 (18%)	40 (89%)	All testing was done at room ambient temperature. Test results showed a 61% correlation of the failing memory bit detected with ECM vs the reported system failing bit.

Table 2

Once the ECM development and correlation work was completed the ECM test flow was applied to production test for this 1GB Registered DIMM product and 3% yield loss was measured. The ECM test flow has since been applied to 1GB and 2GB Registered DIMMs for other customers and the yield impact has ranged between 1-5% for these products.

Feedback from the original OEM customer has indicated that DPPM rates have been reduced an order of magnitude when using ECM screened DIMMs.

SUMMARY

Single-Bit-Errors plague the industry by requiring the industry to provide error protection when using DRAMs in their system applications. These SBEs have slipped through the DRAM suppliers' burn-in and test process and the memory module suppliers test after module assembly. SBE's have exhibited characteristics similar to hard and soft errors in the form of a weak bit that materializes repeatedly under specific conditions.

The role and objective of testing memory modules has changed with the new challenges set upon us by new technology and customer requirements. The DRAM vendors can't always identify the device sensitivities first and screen for them, which requires module suppliers to anticipate a possible system level failure and develop tests for them at the module level.

Testing methods and thoroughness is a key differentiator between Smart and its competitors. Working closely with our customers to understand their applications we can and have developed the appropriate test flow that enables us to ship product that meets the customers' high expectations for performance and quality. This is accomplished through the Enterprise Class Memory (ECM) Testing methodology.

REFERENCES

1. IBM 4/14/98, RAM Reliability – Fault Tolerance
2. A.J. van de Goor, Testing Semiconductor Memories – Theory and Practice, 1998
3. C. Mays and M. H. Woods, “A New Physical Mechanism for Soft Errors in Dynamic Memories,” Proceedings 1978 International Reliability Physics Symposium, April 1978.
4. Soft Memory Errors and Their Effects on Sun Fire™ Systems. Sun Microsystems Inc. April 2002.
5. Mil-Std-883E, Method 1015.9, Burn-in Test, June 1993