

FY12 End of Year Report for NEPP DDR2 Reliability

Dr. Steven M. Guertin
Jet Propulsion Laboratory
Pasadena, California

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

JPL Publication 13-1 01/13



FY11 End of Year Report for NEPP DDR2 Reliability

NASA Electronic Parts and Packaging (NEPP) Program
Office of Safety and Mission Assurance

Steven M. Guertin
Jet Propulsion Laboratory
Pasadena, California

NASA WBS: 724297.40.49.11
JPL Project Number: 103982
Task Number: 03.02.02

Jet Propulsion Laboratory
4800 Oak Grove Drive
Pasadena, CA 91109

<http://nepp.nasa.gov>

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, and was sponsored by the National Aeronautics and Space Administration Electronic Parts and Packaging (NEPP) Program.

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

Copyright 2013. California Institute of Technology. Government sponsorship acknowledged.

TABLE OF CONTENTS

Abstract.....	1
1.0 Introduction	2
1.1 Reliability of DRAMs in Use	2
1.2 Change of Direction	3
1.3 Establishing a Useful Approach	3
1.4 Background and Examples	4
1.5 Hardware Development	4
2.0 Findings From FY11.....	5
2.1 25°C and 125°C Test Results.....	5
2.2 Sample Size.....	6
2.3 Functional Testers.....	6
3.0 Test Planning	8
3.1 Approach.....	8
3.2 Test Devices	8
3.3 Parametric Studies	8
3.4 Cell Data Storage.....	9
3.5 Limited Life Testing	10
4.0 hardware development.....	11
4.1 DIMMs.....	11
4.2 Eureka 2 Tester.....	11
4.3 Functional Tester DIMM Upgrade	12
4.3.1 Upgrades to Hardware	12
4.3.2 Upgrades to Firmware	13
4.3.3 Software Updates.....	15
4.4 Credence D10 System	15
5.0 Testing	16
5.1 Eureka 2.....	16
5.2 Functional Tester	17
6.0 Partnering and Injection	20
6.1 Leverage from MSL Effort	20
6.2 Use for Flight Screening.....	20
6.3 Community Partnering	20
7.0 Future Work	21
7.1 FY13 DDR2 Work	21
7.2 DDR3 Capability.....	21
7.3 Migration to New Hardware Platform	21
8.0 References.....	23
Appendix A. Acronyms and Abbreviations	24

ABSTRACT

This document reports the status of the NASA Electronic Parts and Packaging (NEPP) Double Data Rate 2 (DDR2) Reliability effort for FY2012. The task expanded the focus of evaluating reliability effects targeted for device examination. FY11 work highlighted the need to test many more parts and to examine more operating conditions, in order to provide useful recommendations for NASA users of these devices.

This year's efforts focused on development of test capabilities, particularly focusing on those that can be used to determine overall lot quality and identify outlier devices, and test methods that can be employed on components for flight use. Flight acceptance of components potentially includes considerable time for up-screening (though this time may not currently be used for much reliability testing). Manufacturers are much more knowledgeable about the relevant reliability mechanisms for each of their devices. We are not in a position to know what the appropriate reliability tests are for any given device, so although reliability testing could be focused for a given device, we are forced to perform a large campaign of reliability tests to identify devices with degraded reliability. With the available up-screening time for NASA parts, it is possible to run many device performance studies. This includes verification of basic datasheet characteristics. Furthermore, it is possible to perform significant pattern sensitivity studies. By doing these studies we can establish higher reliability of flight components.

In order to develop these approaches, it is necessary to develop test capability that can identify reliability outliers. To do this we must test many devices to ensure outliers are in the sample, and we must develop characterization capability to measure many different parameters. For FY12 we increased capability for reliability characterization and sample size. We increased sample size this year by moving from loose devices to dual inline memory modules (DIMMs) with an approximate reduction of 20 to 50 times in terms of per device under test (DUT) cost. By increasing sample size we have improved our ability to characterize devices that may be considered reliability outliers.

This report provides an update on the effort to improve DDR2 testing capability. Although focused on DDR2, the methods being used can be extended to DDR and DDR3 with relative ease.

1.0 INTRODUCTION

During FY12 efforts to ascertain DDR2 device reliability, we expanded capability for testing and critically reviewed our test capability. The focus for this year was to identify places where it is possible to improve on manufacturer efforts (where we have limited knowledge of what manufacturers do for reliability testing), and to identify methods that can be carried out on components being considered for flight. The key capabilities we identified where increased focus can result in improved reliability data are related to having significantly more time for characterizing parts than the manufacturer does. We also determined that we do not have and cannot get the die, design, and process-level reliability information that the manufacturer has. Thus we moved away from focused testing of specific reliability qualities, and focused more on using our time to develop more pre-flight characterization, with focus on methods to identify and remove parts with reduced performance.

This year's work also reflects findings in collaboration with both flight users and the results from FY11 work. We have increased our focus on pattern sensitivity of bit errors due to flight observations—where preflight characterization was insufficient and flight anomalies are less understood than desired (due to a lack of data about pattern sensitivity on the flight devices). We also determined that sample sizes used for earlier study were simply insufficient for these highly commercialized devices. Although it may be possible to get test structures and examine the basic failure mechanisms of DDR2 devices, it is simply not a viable path towards understanding the reliability issues of potential flight devices. A solid knowledge base on any particular generation of DDR2 cell structures cannot be certain to reflect devices used for flight. Instead we decided to focus on increasing sample size and increasing the variety and number of characterization tests that are run.

This report provides a review of the FY12 efforts. We will first present a rough overview of the entire task, including justification based on available literature and test results, and the efforts carried out in support of this overview. We will also briefly review the FY11 results. The updated approach also requires modifications of test planning, which will be covered. We will then discuss hardware development and results from testing devices with the developed hardware.

1.1 Reliability of DRAMs in Use

The relevant approach to determining the reliability of DDR2 devices for flight missions depends on a large number of factors, ranging from the manufacturer's efforts to improve reliability to the actual failure mechanisms in the field. We must also take into account where the greatest benefit can be gained for flight missions. Development of reliability data for NASA missions gains the most benefit by focusing on the amount of time available for testing of parts. We will focus here on performing reliability testing utilizing the time benefit available for NASA missions.

CMOS devices can experience reliability failures due to several mechanisms. The most commonly discussed mechanisms are electromigration, time-dependent dielectric breakdown, and hot carrier injection. The appropriate models of each of these mechanisms are not simple and require considerable study to determine the right dependencies. This material is outside of the scope of the effort that this NEPP task, which is focused on packaged commercial devices, can accomplish. However, the worst-case stress conditions can generally be listed as: maximum and minimum bias, maximum and minimum temperature, and switching and constant electric field.

Failures or errors in devices are only relevant to study in up-screening if the failures are permanent, or if the error rate of the device is related to its construction, and not due to random processes in the field. Focusing on errors observed in a laboratory setting can remove any reduced reliability devices before they are put in the field. The error rates for devices in the field were carefully studied by Schroeder [1] where the Google computer fleet was examined over a 2.5-year period. They showed that DIMMs have an approximately 10% chance for a correctable error (CE) in a year, and about a 1% chance for an uncorrectable error in a year. Figure 1.1-1 shows the relationship between errors in a given month and

errors in the previous month. One interesting finding in the Google data is that the FIT/Mb rate was found to be 25,000–70,000, which is up to 40 times higher than previous studies found.




Figure 1.1-1: Correlation between error rates in DIMMs from one month to the next. The left panel shows how CEs in a month is related to those in the previous month. The right panel shows the autocorrelation [1].

1.2 Change of Direction

For FY11, significant amounts of data were collected on bare 78 nm DDR2 SDRAMs. Devices were costly to prepare (approximately \$500 each). Data collection was limited to nine functional testers testing one device each. The results showed essentially no change in any relevant parameters after 1000 hours of life testing. The actual collected data was also limited in scope. We focused narrowly on cell retention and a couple of the operational currents. Our testing also used only one data pattern on the devices—both for retention scanning and for the stress data pattern used on the DUTs during life testing. The collected data, impacted by these testing limitations, suggested the need to alter the test methods for this task in order to improve applicability to flight projects.

The FY11 findings prompted the need to increase the number of DUTs, the number of characterization methods, and the number of datasheet parameters that were directly tested. This pushed the need to achieve datasheet operating frequency and dramatically decrease the cost of DUTs. In addition, testing showed that multiple data patterns were needed to provide useful data and stress conditions. Life testing and thermal acceleration are still considered important, but we have chosen to apply those to devices that already show signs of outlier behavior. This way, testing performed as initial characterization can be used to identify outliers, both for research reasons and for flight project up-screening. Potential impacts on flight parts can be assessed through targeted life testing of outlier devices as well, without the need to perform long term life testing on devices that would not be identifiable during up-screening.

1.3 Establishing a Useful Approach

Given our resources and the difficulty of extracting information from manufacturers, we determined that it is not possible to perform a complete reliability test regime on potential flight parts within the constraints of limited manpower and technical information on device production. Nor is it possible generally to target key reliability concerns on potential flight parts because the likely failure mechanisms of a particular device are not available to the flight project or to NEPP.

Manufacturers are much better positioned and must have better reliability engineering on average devices than can be achieved by laboratory up-screening. Thus we must determine how to ensure that any devices selected for a mission are at least as well performing as average devices, while also ensuring that any potential reliability issues that can be up screened are examined.

1.4 Background and Examples

Reliability testing of DDR2 devices generally consists of verifying datasheet parameters, examining changes in or violations of those parameters as a function of various types of life testing and duration, and verifying device operation in mission-specific conditions. The final of these is expected to be verified by ensuring the DUT meets the others; however, this is not guaranteed.

A recent example from NASA missions illustrates the difficulty of using parts that were not fully up screened. Recent spacecraft anomalies include bits that show unexpected loss of data, isolated to a specific address or set of addresses. These data losses are occurring at approximately 50°C in systems that use refresh rates of approximately 32 ms (this is better than the specification which is 64 ms). Ground-based testing of these devices included checkerboard and inverse checkerboard, but did not include complex algorithms such as March-X or multiple pseudo-random patterns. So it is not known if in-flight anomalies are intrinsic to the devices or are a problem picked up during or after assembly.

1.5 Hardware Development

As indicated earlier, the loose device approach used in FY11 is not sufficient for the predicted quantity of devices necessary to obtain a sample relevant for reliability studies. The material above suggests a failure rate between 0.1 and 1% for individual parts. And at a price of \$500/device, the indication is that the cost may be as high as \$500,000 in test parts before a subject with interesting reliability issues is found. By using DIMMs, with prices on the order of \$1/device, we can reduce the average price per device with poor reliability to around \$1,000 (i.e. every 1000 devices purchased is expected to have one device with poor reliability).

In order to test DIMMs it was necessary to develop hardware. We developed two specific solutions for DIMMs. The first is to ensure industry-level basic reliability through a lot-acceptance tester with some additional ability to perform many industry-standard reliability measurements. This was done with the Eureka 2 tester, which is designed for DDR2 DIMMs. The second development was to build a DIMM adapter for our existing functional DDR2 Reliability Tester (D2RT). These are discussed later in this report.

As an exploratory option we also examined the use of a Credence D10 tester for this work. We concluded that due to the limited device throughput this tester would be useful for flight part screening, but not viable for testing hundreds of devices as needed for reliability studies (given that we do not have sufficient knowledge of each type of device to know which failure mechanisms and stresses are relevant). Because of the limitations, we have not injected the Credence D10 into our reliability test flow.

2.0 FINDINGS FROM FY11

This section provides a quick review of the work in FY11 in order to set the stage for this year's work. We focus on the test effort, developments, and the lessons learned.

2.1 25°C and 125°C Test Results

For FY11, life testing was performed with temperatures of 25°C and 125°C. This testing was performed with individual test devices and was intended to show changes in the retention curve with duration of life testing and parameters of life testing. The findings are more completely reported in [2] and the basis for the test approach can be found in [2] and [3].

In Figure 2.1-1, we see that there is minimal change in operating currents during life testing. This result may be limited due to testing at 125 MHz, which is not the correct frequency for obtaining the true IDD3N current for the test devices. Also, in Figure 2.1-2 (for Micron devices), we see that there is significant change in the room temperature (25°C) retention curve. This is not believed to be true device sensitivity but rather highlights the need to control test temperature more closely. At elevated temperatures (85°C), the temperature is better controlled and the DUT shows a slight improvement in cell retention in the mid-range. This change in the mid-range retention is indicative of possible imprinting because the test pattern was fixed (storing the same value for the entire duration of testing).

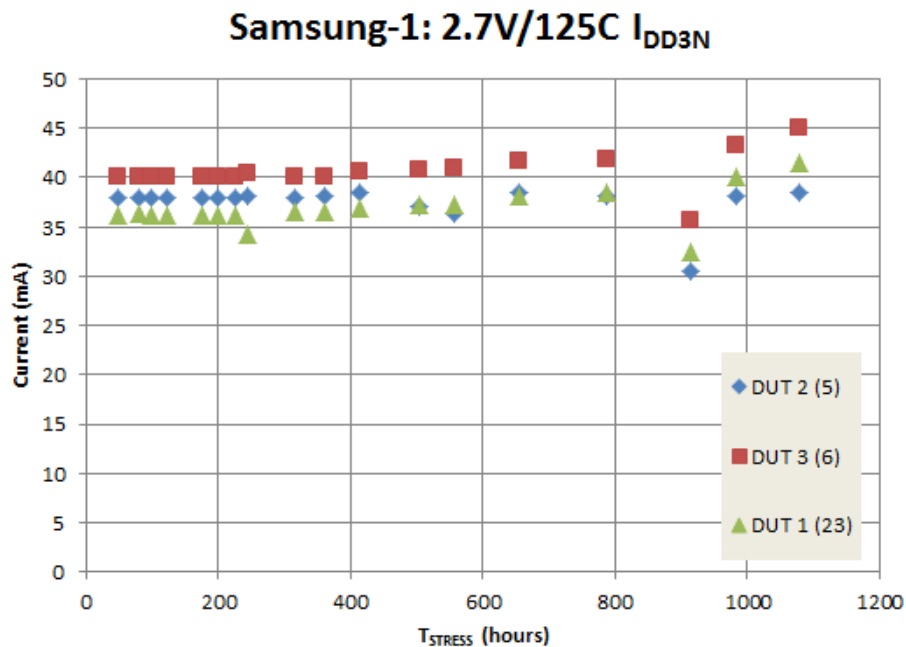


Figure 2.1-1: The Samsung 2.7V/125°C test point provided the most significant change in operating currents during life testing [2].

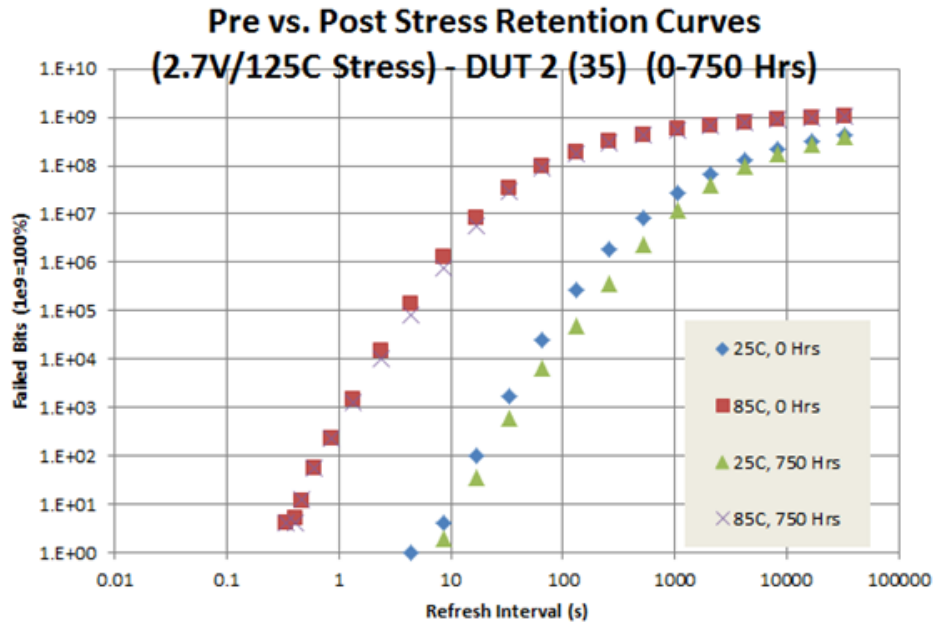


Figure 2.1-2: The Micron devices stressed at 2.7 V/125°C showed the largest change in data retention. However the change was towards more robust data storage. The findings also indicated the refresh approach required attention, and room temperature measurements were not well-controlled in terms of temperature.

The main things these plots show is that the characterization effort was not sufficient for our main goals of understanding the devices well, and the sample size was inadequate to ensure some outliers are in the test samples and would provide interesting reliability results. The number of parameters measured at each characterization point was not sufficient to identify significant changes, and the set of measurements was insufficient for drawing general conclusions about the test samples that would be useful for flight projects.

2.2 Sample Size

FY11 testing utilized test samples of three devices each. The lack of relevant reliability data clearly indicated this was not a large enough sample size. The difficulty is that in reliability testing there are two types of data one can collect. The first is the type of data from stress that leads to failure, i.e. observing failure types. The second is to look at how an ensemble of devices responds to stress in order to identify changes in failure rates. In the former, you must test the majority of devices until a reliability failure occurs. In the latter you must test enough devices to observe a failure rate. Unfortunately the 1000-hour life stress, even at 125°C and 2.7 V operating current, did not result in degraded devices. And this means the only viable test data that could be derived was from failure rates—but the number of DUTs was insufficient for establishing failure rates. Thus sample size in FY11 was simply insufficient.

2.3 Functional Testers

The FY11 testing did highlight successful implementation of the JPL functional tester. Nine boards were successfully employed simultaneously using three laptops, each connected to an Opal Kelly USB adapter. The connection diagram for implementing the test system is shown in Figure 2.3-1. For FY12 this approach was adapted by changing the D2RT to a DUT mezzanine card to facilitate connection of a DIMM instead of a loose DDR2 device.

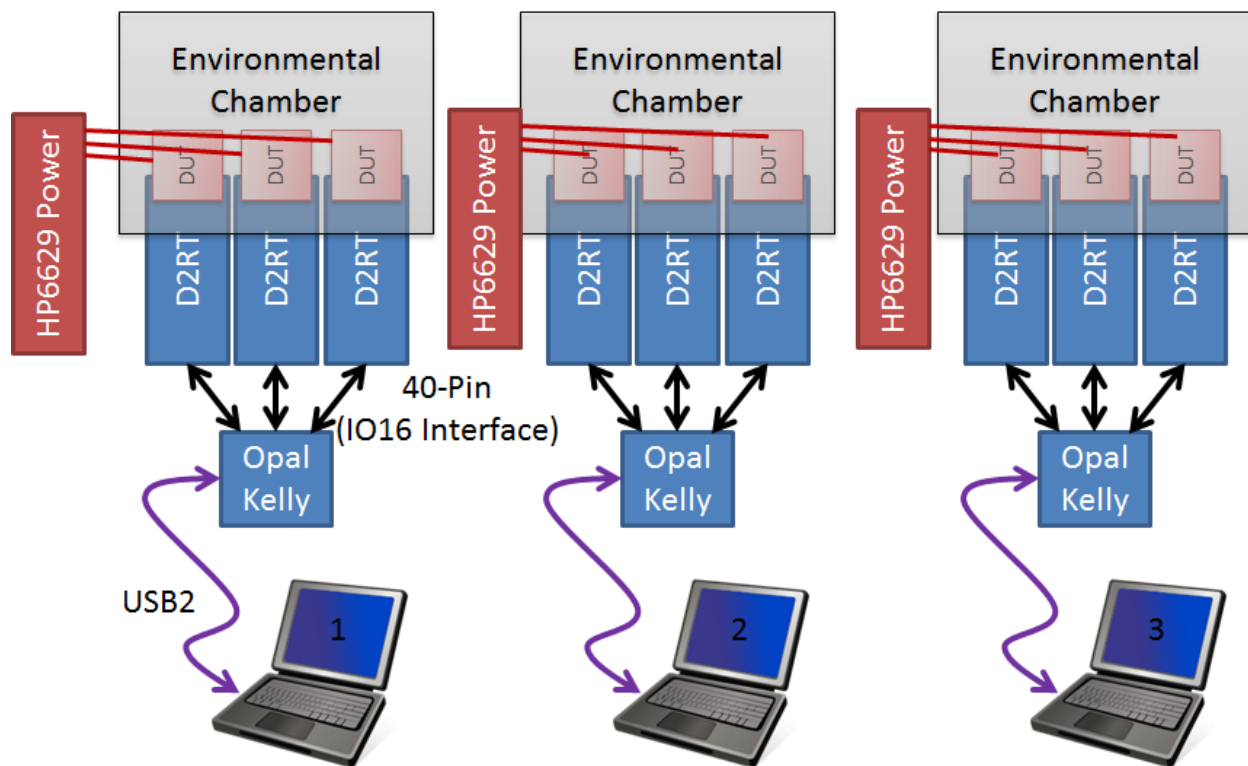


Figure 2.3-1: The layout of functional tester used in FY11 testing [2].

3.0 TEST PLANNING

Given the significant effort on development of test capability, it is relevant to discuss test planning based on the increased test capability coming on line. This section presents the test plan as implemented for the new equipment brought on line in FY12, and as envisioned for future testing under the upcoming DDR2 reliability test effort.

3.1 Approach




The main limitations for testing DDR2 devices were identified as: IDD surveys require full operating speed; limited IDD survey is insufficient; testing only three devices is not good enough for having a reasonable probability for outlier devices in the sample; and a single data pattern is not sufficient for finding weak cells. Given these issues, the key items of interest in the testing of DDR2 devices under this NEPP task are the following:

1. Verify functionality of devices across as much of the datasheet as practical.
2. Measure all IDD parameters utilizing standard test equipment.
3. Examine cell data storage using multiple methods including march tests and multiple data patterns, including the measurement of retention curves.
4. Where appropriate, perform limited life testing.

3.2 Test Devices

DIMMs provide an excellent source for DDR2 devices. For this year we obtained Samsung, Micron, and Hynix 2GB DIMMs that were produced using 16 1-Gb devices. Each device type was obtained in a set of 10 DIMMs, totaling 160 DDR2 devices for each manufacturer. All test devices have 14 row bits, 10 column bits, and 3 bank bits. Devices all have an 8-bit data word. Device details are given in Table 3.2-1.

Table 3.2-1: DDR2 devices in DIMMs for FY12 development.

Manufacturer	Part Number	Device Photo	Number of Parts	Feature Size
Micron	MT47H128M8CF-25:H [4]		160	50 nm
Samsung	K4T1G084QF [5]		160	5x nm
Hynix	H5PS1G83EFR-S6C [6]		160	5x nm

3.3 Parametric Studies

Parametric measurements on DDR2 devices are important for assessment of reliability. Datasheets show a very large number of parameters that can be measured. This includes everything from input capacitance to the structure of the clock. However, as indicated earlier, the majority of these parameters cannot be

measured with the resources available to this task in the quantity or detail required. We have determined that the most appropriate parametric studies that can be performed on DIMMs are to measure the standard datasheet IDD values, verify functionality across different data patterns, measure the time-dependent nature of the storage cells, and attempt to correlate initial outliers with reduced overall life performance.

In a DIMM, IDD values are combined from multiple devices. The IDD values will be extracted using the Eureka 2 tester. The measurement descriptions listed in Table 3.3-1 are those extracted by the Eureka 2 tester. Values in Table 3.3-1 represent the manufacturer’s specification for individual devices.

Table 3.3-1: IDD values measurable by Eureka 2 system and their specification for individual devices in DIMMs [4-6].

IDD Item	Description	Specification (mA, at 800 MT/s, CL=6)		
		Micron	Samsung	Hynix
IDD0	Operating One Bank Active-Precharge Current	65	45	75
IDD1	Operating One Bank Active-Read-Precharge Current	75	51	85
IDD2P	Precharge Power-Down Current	7	10	10
IDD2Q	Precharge Quiet Standby Current	24	20	32
IDD2N	Precharge Standby Current	28	25	45
IDD3P	Active Power-Down Current	20	23	25
IDD3N	Active Standby Current	33	37	55
IDD4W	Operating Burst Write Current	125	72	170
IDD4R	Operating Burst Read Current	120	80	160
IDD5	Burst Refresh Current	145	105	170
IDD6	Self Refresh Current	7	10	10
IDD7	Operating Bank Interleave Read Current	210	160	230

We also expect to use the Eureka 2 tester to provide information about the voltage and frequency space in which devices function nominally. This is extracted by obtaining shmoo plots of the voltage and frequency space with a given device functionality test, which determines if the device performs successfully.

Additional parametrics can be measured by the Credence D10. These are standard operating voltages and currents: leakage currents on all pins, output driver strength, logic high and low values, edge timing, and other items that can be measured with frequency below 200 MHz (which is too low for many measurements). Note, however, we have determined that for general reliability studies the Credence D10 would be a significant bottleneck to our planned sample test structure and has been eliminated from the test planning.

The items discussed here are used as part of the characterization normally performed on DIMMs during initial characterization. Periodic characterization of devices selected for life testing enables us to monitor changes to these parameters, which could help isolate outlying devices and correlate deviations in parametric response to early failures or devices with high soft error rates.

3.4 Cell Data Storage

The data cells are often expected to show large variation in storage capability with life and stress exposure, although for a given technology the cells may show minimal storage degradation with life exposure. The cell array and its addressing structures cover the majority of the device area. The storage cells also are more exotic in structure than standard transistors, utilizing needle-like capacitors, and

driving the overall density. Thus analyzing the performance of the storage cells is very important, and may reveal reliability behavior significantly different from the other device structures.

Cell data retention results from FY11 were inconclusive, showing improved cell retention after life testing, and problems with elevated temperature data retention due to test system operation. This clearly argues for the two changes implemented in the current characterization test regime. First, the cell retention measurements must be performed with multiple data patterns. And second, better control over refresh is required. These changes and others including transfer of device images are discussed in Section 5.3.2.

3.5 Limited Life Testing

Because of the issues with the number of devices required, life testing in accelerated environments is very resource intensive and largely outside of the scope of work that can be done under this NEPP task. For planning purposes, however, limited life testing is expected to be useful on current and future test devices.

Initial characterization is expected to identify outlier devices. Outliers are identified by how they perform during characterization testing of new devices (i.e., the parametric and cell characterization discussed earlier). With more than 100 of each DDR2 device of interest it should be possible to identify standard samples and separate the outliers. The outliers and a few standard samples can then be used as the basis for a limited life test to observe if there is a correlation between outlier behavior and reduced reliability.

It should be noted that DIMMs are unlikely to have more than one outlier along with the other seven or more devices on the DIMM. These other devices would be sufficient for providing a control group, except that they all will share common mode problems. Thus at least one additional in-family DIMM is required to provide an adequate control sample.

In terms of test planning, life testing is principally driven by whether or not the test samples show interesting (outlier) properties that warrant extended life testing. If no devices show outlier behavior, then no basis for improved up-screening is useful. This is also an important potential observation.

4.0 HARDWARE DEVELOPMENT

A major part of the effort this year was switching to DIMMs as the vehicle of DUT testing. In order to support DIMM testing we enlisted an industrial tester and upgraded the D2RT to enable functional testing of DIMMs. In this section we discuss the hardware upgrades to the overall effort, including firmware and software upgrades where appropriate.

4.1 DIMMs

For this year we redirected the DUT approach to enable testing of DIMMs that reduce the cost per device to approximately \$1 (DIMMs with 16 DDR2 devices readily sell for about \$20). However we also identified that the rate for devices with a relatively high soft error rate, or hard error, is on the order of 0.1 to 1%. Thus, we need between a few hundred and a few thousand devices to ensure a test population with a handful of devices of interest for this study. For work here we decided to start with hundreds (~\$200–\$500), understanding that future work may require thousands (\$2,000–\$5,000).

DIMMs provide a common interface, and meet a specification that makes running them somewhat straightforward. But this specification does not allow individual examination of devices, making DIMMs an excellent source for observing soft errors in large numbers of devices, but a mediocre source for isolating device parameters.

A compromise on DDR2 DIMMs is that DIMM adapters for loose individual devices exist. In addition an adapter for connecting previous DUTs for this NEPP task was produced to enable testing in DIMM-based equipment. This adapter is shown in Figure 4.1-1.



Figure 4.1-1: Adapter for connection of NEPP DDR2 individual DUT daughter cards.

4.2 Eureka 2 Tester

The Eureka 2 tester is a standard acceptance tester that can test DDR and DDR2 devices with a wide variety of standard tests, including the IDD and March tests. It can also perform shmoo testing across operating voltage, frequency, and any particular test desired. The Eureka 2 test system is shown in Figure 4.2-1.



Figure 4.2-1: The Eureka 2 test system is an acceptance tester for DDR2 DIMMs.

For this work we have inserted the Eureka 2 test system into the standard DIMM characterization approach. It is ideal for testing DIMMs at speed, with standard test patterns. It alleviates the need to make the D2RT perform these standard tests (which is very difficult due to operating frequency requirements). Using the adapter discussed in Section 4.2, the Eureka 2 has also been used to perform tests on DUTs from the 2011 data set. Further, adapters that allow loose devices to be connected to a standard DIMM slot can be used to enable collection of DIMM-equivalent data from potential flight parts.

4.3 Functional Tester DIMM Upgrade

The D2RT was developed under this NEPP task in previous years and development continued this year. This tester is designed to enable device stress during life testing without tying up limited resources, and to provide measurement of device characteristics that require time-intensive testing (such as cell retention characterization). As a result, the tester is focused on functional capability rather than parametric measurements. The tester enables device bias, temperature, stored data pattern, and pattern alternation – enabling periodic electric field oscillations (though at slow speed).

Development efforts in FY12 include hardware design, firmware development (including multiple data pattern options, refresh, and full device data image collection), and software development to enable and automate the upgrades. This tester has been used to collect 35°C data on Hynix DIMMs and engineering data on Micron DIMMs.

4.3.1 Upgrades to Hardware

Hardware upgrades for the D2RT were focused on establishing a working DIMM adapter. This effort focused on an approach that is expected to translate well to DDR3 or 4. Bringing up the DIMM adapter was not trivial. It was observed that movement from individual SDRAM or DDR2 devices to the architecture of a complete DIMM requires significant engineering work due to a series of pitfalls, including: termination current, multiple DIMM architectures (buffered, unbuffered, and registered), and multiple DIMM parameters (ranks, signal multiplexing, etc.). For this work we determined the most effective approach was to use unbuffered and unregistered DIMMs, and to program the DIMM parameters into the firmware rather than dynamically detect the DIMM architecture.

The D2RT DDR2 DIMM adapter was built as a two-rank 72-data bit mezzanine card. The prototype design (four built) is shown in Figure 4.3-1. Some errors were found that have been folded into a redesign that is slated for fabrication in FY13 to enable us to bring up nine DIMMs in environmental chambers simultaneously.

At present we have shown that the prototype mezzanine cards can support three DUTs (connected to three motherboards) in an environmental chamber simultaneously, all controlled by a single operations computer. However, because of jumper wires and termination power problems, it is very difficult to swap DUTs, making this a difficult to use and sometimes unreliable test system.




Figure 4.3-1: The JPL functional tester has been improved by the addition of DIMM test capability. The DIMM adapter is connected to the FPGA board and provides separate power for the DIMM.

4.3.2 Upgrades to Firmware

Firmware upgrades cover four general areas. The majority of upgrades to support DIMM operations were general upgrades, such as increasing the data width. The other upgrades target increased ability to operate DUTs and are focused on refreshing the DUTs and improving pattern capabilities.

4.3.2.1 General Upgrades

Some general upgrades were made to the functional tester. The biggest change was to improve the reliability of the firmware by removing unnecessary speed requirements that provided little benefit when testing at speed is now performed with the Eureka 2 system. The DUT clock was slowed down to 33 MHz and changes were made to operations firmware developed to enable operation of each type of test device (Note that this slow test speed operates the DUTs in test mode, and is only used to measure cell retention—which is clearly dominated by processes that take minutes to hours and are easily separated from effects picked up when running at 33 MHz, instead of the specification minimum 125 MHz).

The firmware was also modified to handle 72-bit wide data words with a burst length of 8 data words. The D2RT now supports individual error counters on every one of the 72 DQ (data) lines in an ECC DIMM, as well as separate device-level error counters for every access cycle (burst read of 8 data words).

These counters make it easy to spot outlier devices and plot error rates on a device-by-device, or data line-by-data line basis.

4.3.2.2 Refresh Operations

The refresh method used in earlier testing was to access all rows in the device within the refresh interval of 64 ms. This method did not work during FY11 testing due to lack of datasheet support for this very old refresh method. This issue was handled under joint development of improved refresh operations that benefitted from collaboration with MSL testing in FY12. The test system is capable now of performing auto-refresh operations on a time-scale consistent with the datasheet requirements (i.e., 8192 auto refresh cycles in 64 ms). Some minor alterations of the MSL-developed system were required to enable the cell retention measurements that were not needed under MSL.

4.3.2.3 Pattern Capabilities

DUT characterization, as indicated earlier, requires improved examination of cell-level response, especially with multiple data patterns. The initial pattern capabilities of the JPL Functional Tester were limited to a simple address-based pattern and its inverse. In order to improve the data pattern options for DIMM testing we modified the DIMM firmware to support the following pattern generation algorithms.

1. Fixed pattern—a fixed 64-bit pattern is written on every 8-byte burst (this pattern is bad for identifying compromised device operation). An example of this pattern is “all 0s.”
2. Address-based—a pattern that uses the current address (8-byte boundary) to generate a data pattern (this pattern is highly regular and may mask some error modes). This pattern can be inverted with the inversion flag.
3. Pseudo-random—two 63-bit linear feedback shift registers (LFSR) with outer feedback and a generator polynomial of $x^{31}+x^5+1$ are used with the 64-bit fixed pattern values being used to seed each generator (bits 62 to 32 are the same as 30 to 0). This generator has a period of $2^{63}-1$. Each value fetched is 29 elements further in the LFSR cycle to limit the relationship between sequential values. (The LFSR shifts data bits to the next higher order, and in the case that bit 62 was a 1, it inverts the generator orders, bits 31, 5, and 0.) The pseudo-random generator output can be inverted with the inversion flag.

4.3.2.4 Device Image Transfer

The D2RT system was upgraded based on the leveraged capability built for MSL SDRAM SEFI testing. The entire device image can be transferred to the test computer in about 40 seconds under the current firmware. Figure 4.3-2 shows the error pattern for an entire 1 Gb DDR2 memory. This map was used to help troubleshoot problems with the firmware. This mapping will be valuable in the future for identifying error patterns of weak bits during retention scans (this is not part of the test planning at present, but can be leveraged if needed).

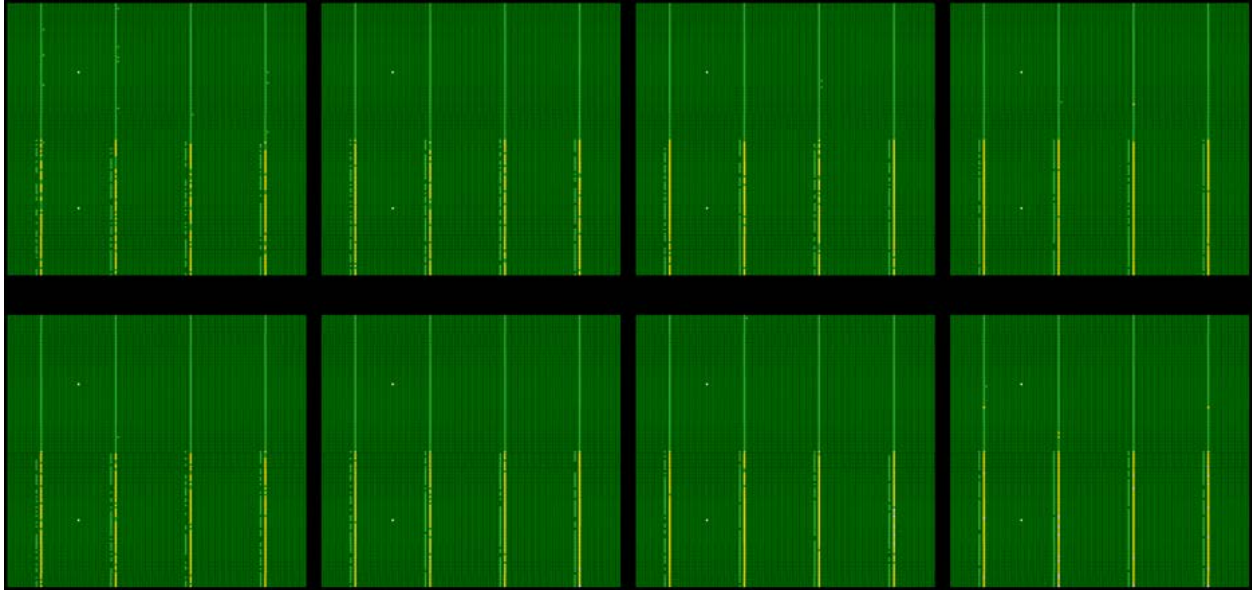


Figure 4.3-2: Graphic representation of the error map generated from a full device snapshot. The 8 regions correspond to 8 banks. The dark green dots (which are the majority of the image) are rows with no errors, while the other colors refer to various counts of errors.

4.3.3 Software Updates

The D2RT communicates with a software package on a PC through an Opal Kelly USB adapter. The PC software was upgraded in FY12 to enable automated retention measurements over multiple rank DIMMs with multiple data patterns. This system enables script-based reconfiguration which cycles through an arbitrary number of configurations, set by the test engineer, and then the system automatically samples data storage over a series of refresh intervals. Note that refresh intervals of interest run from 32 ms to 8 hours, and a single retention sweep can take a day or longer for a single rank and single data pattern. Thus testing a two-rank DIMM with ten data patterns can take multiple weeks. Automation can easily improve characterization times by a factor of two unless test engineers are on around-the-clock call for the entire duration of the scan.

4.4 Credence D10 System

The Credence D10 tester is a high pin count parametric tester with the ability to determine timing, voltage levels, and currents on many pins, but it is resource-intensive for DDR2 characterization and cannot provide very high operating frequency (i.e., above 200 MHz). This last limitation is not unexpected, as board design is very important for getting a circuit to work at such speeds.

Because of the specific benefit of the Credence system at performing detailed measurements, the idea of using it on a DIMM is inherently problematic because of multiplexed IO signals and shared power, ground, and control signals, amongst eight or more devices.

We still believe the Credence D10 system can be useful for mission-specific up-screening. However it is simply not useful for testing hundreds of devices, which is the low estimate of required devices to observe a device with compromised reliability. Thus the Credence D10 is only mentioned for reference and may be useful for flight lot parts, but cannot form the basis of a general reliability study under this NEPP task.

5.0 TESTING

Most of the FY12 efforts went into developing test plans that could make the best use of our resources, bringing new equipment on line, and obtaining 50-nm test devices. Part of the test planning included examining the capabilities of the various testers available to determine an appropriate test approach. The general results of the capabilities of our testers are discussed in this section as a basis for testing to be performed in FY13.

We performed baseline testing of the DDR2 operations that will provide the characterization data for the DIMMs. We used the Eureka 2 test system to perform a battery of standard measurements and verification of high-speed write and read operations. We used the D2RT to perform cell retention measurements on both Micron and Hynix devices (Samsung DIMMs have slightly different operating limitations and the D2RT is not yet configured correctly to communicate reliably with them).

5.1 Eureka 2

The Eureka 2 system was used to capture IDD measurements and perform March-# and random access testing to verify functionality. We also performed shmoo testing with the Eureka 2 system to obtain the voltage and frequency space for functionality. The IDD summary is given in Table 5.1-1 for Hynix devices.

Table 5.1-1: The currents observed during initial characterization of Hynix DIMMs, in mA.

Measurement	H1	H2	H3	H4	H5	H6
IDD0	378 mA	376 mA	369 mA	375 mA	378 mA	375 mA
IDD1	457	447	437	447	441	439
IDD2P	66	66	65	66	67	66
IDD2Q	176	175	172	202	178	175
IDD2N	174	172	170	172	175	172
IDD3P	62	64	60	64	64	62
IDD3N	544	541	533	535	546	533
IDD4W	533	425	517	414	539	531
IDD4R	1310	1281	1287	1173	1281	1146
IDD5	847	851	833	835	843	835
IDD6	37	37	37	37	37	36
IDD7	427	441	429	427	439	433

The shmoo plot of operating frequency versus voltage is provided in Figure 5.1-1. In order to be listed as a pass, the DIMM had to pass a March X test at the given voltage and frequency of each box. Note that the shape of the shmoo plot is not entirely expected, as the DIMMs do not appear to work in some frequency bands below the maximum operating frequency of 400 MHz. At 1.8 V the sample DIMM failed to pass the March X test at 360 and 370 MHz. It should be noted, however, that 400 MHz-rated DIMMs are expected to work at frequencies that are multiples of 66.7 MHz, and the DIMMs all worked correctly at 333 and 400 MHz. The dead spot around 366 MHz is not a useful spot in their functional envelope, thus this behavior seems acceptable.

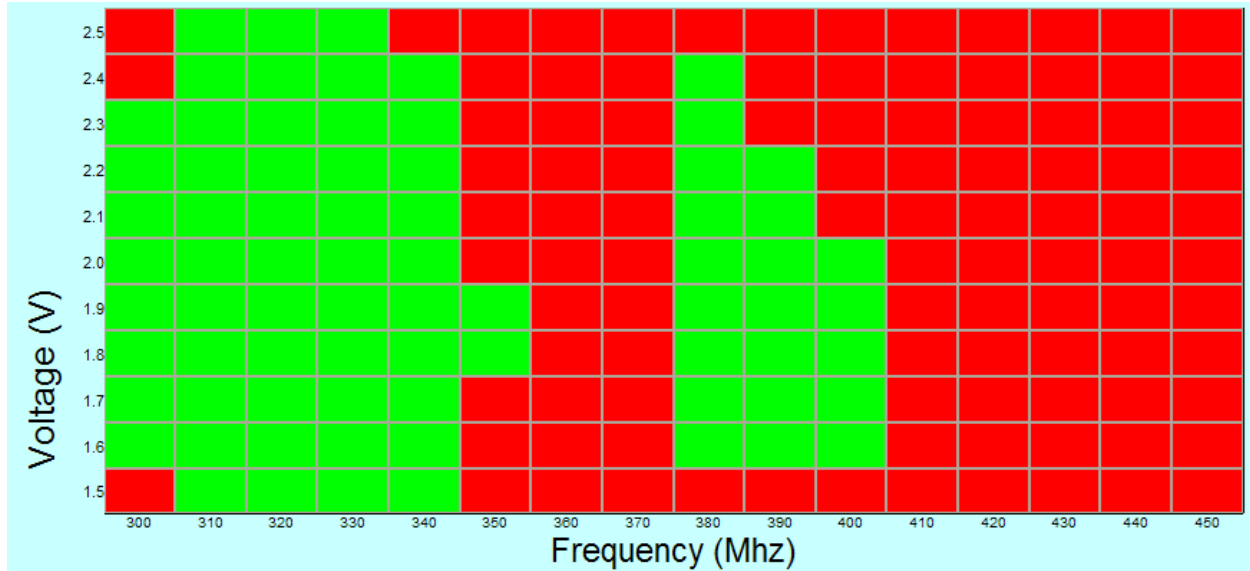


Figure 5.1-1: Shmoo plot of operating voltage versus frequency for Hynix DIMM H1. Note that this part does not fully work in a valley between 350 and 380 MHz, but there are no common operating frequencies in this range.

Similar measurements have been taken using the Eureka 2 test system with the following DUTs:

1. Hynix: 6 DIMMs (96 devices)
2. Micron: 9 DIMMs (144 devices)
3. Samsung: 3 DIMMs (48 devices)

5.2 Functional Tester

The DIMM modifications of the D2RT took significant time during FY12. We now have initial characterization on Hynix and Micron DIMMs that provide a component-by-component data set that can be used to identify outliers. Preliminary results for the Hynix DIMM labeled H2, tested with an address-based pattern (which is not considered to be an effective pattern for cell-level stress testing), are presented in Figure 5.2-1. The first bits start to fail at about 4 seconds retention time (at 35°C), and the weakest devices on the DIMM appear to be components 7 and 2. Figure 5.2-2 shows the summary of all Hynix DIMMs tested with the address-based pattern.

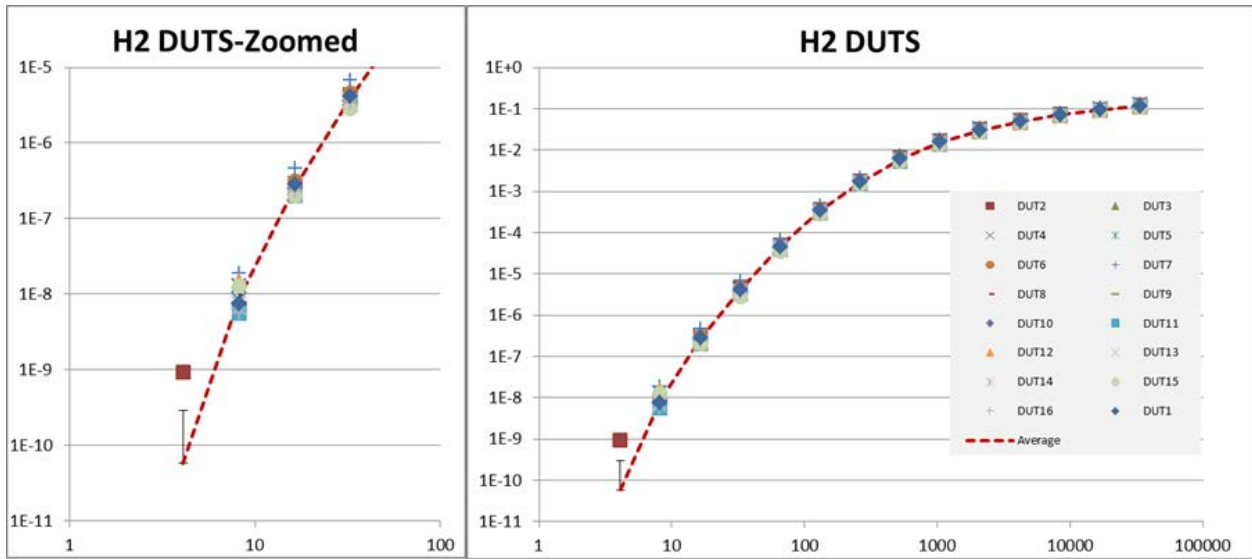


Figure 5.2-1: Cell retention time (using address-based pattern) for Hynix DIMM H2, where fraction of bits failing is plotted against retention time (in seconds). Note that component 2 (DQ8-15, rank 0) might be an outlier, while component 7 shows the worst overall weak-bit performance (DQ48-55, rank 0) (since it clearly has the most failed bits in the sub-100s refresh bins).

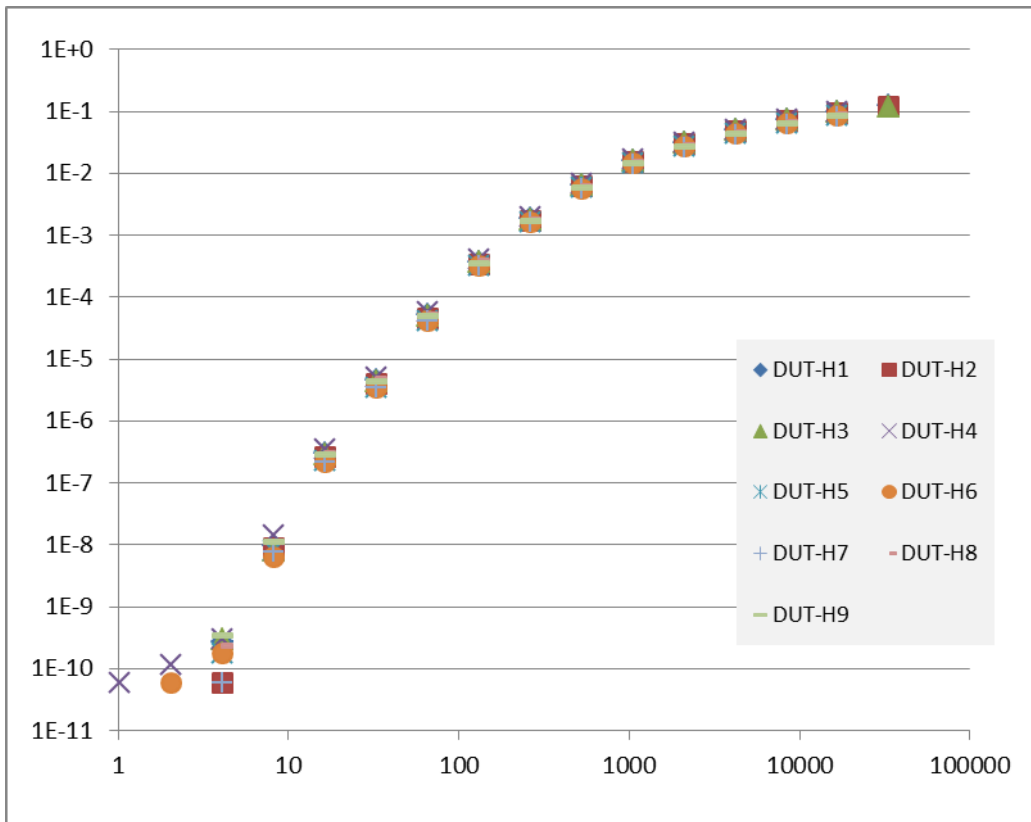


Figure 5.2-2: Summary curve showing data retention results from all 9 Hynix DIMMs using the address-based pattern.

Data taken on Micron devices was largely to validate the test system. Only five of the devices on one DIMM were characterized. We compare the five Micron devices to three Hynix devices analyzed with the system during this stage of development in Figure 5.2-3. These data show the number of bits failing as the refresh interval is increased (with measurements taken at room temperature). At about 9 seconds refresh

interval all of the Hynix devices outperform the Micron devices. The curves, however, look more similar than they do different. And based on findings from FY11, it is known that the test conditions here (room temperature) are not well enough controlled to draw conclusions from these plots. The test pattern was address-based.

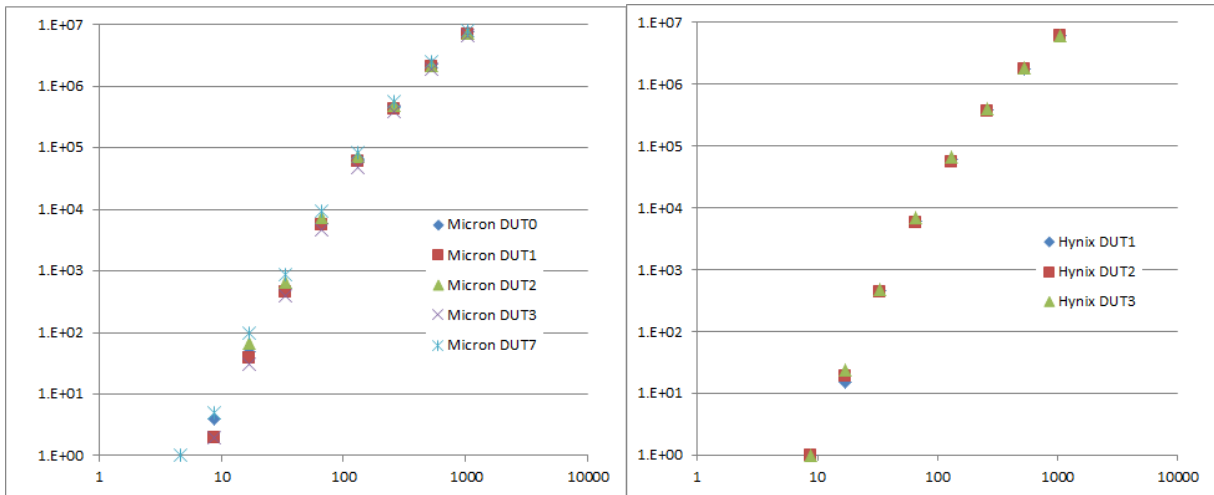


Figure 5.2-3: Micron (left) devices and Hynix (right) during initial evaluation of the functional tester for DIMMs. The plots show the number of failed bits versus the time between refresh cycles (in s). At this point some of the DUT ports on the DIMM were not reliable, so only subsets of the DUT data were analyzed. The behavior of the two device sets is fairly similar, but there is some indication of better performance at lower refresh rate with the Hynix devices. The test pattern was address-based.

Because the preliminary results were obtained using firmware with limited pattern capability (a single address-based pattern was available), we were not able to ascertain anything about the pattern sensitivity or the bit sensitivity. This is a key capability that is being added to the test system before the full characterization effort in FY13. Testing will be done with address based, multiple fixed patterns (all 0's, all 1's), and multiple pseudo-random patterns. It is expected that each of these types of patterns will be capable of providing insight into and examples of weak bits if they exist.

6.0 PARTNERING AND INJECTION

This section provides a brief overview of the actual collaborative work done under this task, and the potential collaborations being developed.

6.1 Leverage from MSL Effort

The D2RT is a robust platform for DRAM-type device testing. As such it was a natural choice to support SEE testing of MSL SDRAMs as a hardware platform. New capabilities were added to the test system that have been leveraged and discussed in the hardware development section, including refresh capability and device image transfer.

6.2 Use for Flight Screening

The characterization approach applied here is recommended as a screening approach for DDR2 devices for flight use. It is not expected that outlier devices will be observed during any up-screening effort. If outlier devices are observed, they should not be used for flight. More extensive characterization may still be warranted (possibly including life testing), but the characterization approach forwarded here will provide meaningful data sets with minimal schedule and budget impact on flight projects.

6.3 Community Partnering

The change of direction towards DIMMs this year was partially prompted by potential collaboration. At present the DDR2 capabilities available under this NEPP task are sufficient for a wide variety of testing. DDR2 devices are currently being examined by several aerospace organizations and we are actively pursuing collaborative options. However, it is likely that the best situation for collaboration in the near term is in DDR3, and as such this is a future direction for this task.

7.0 FUTURE WORK

This section briefly covers the future directions of interest under this NEPP task.

7.1 FY13 DDR2 Work

For FY13 we will do the following tasks. First we will perform a revision of the DIMM mezzanine card. We then plan to perform full characterization of Micron, Samsung and Hynix 1Gb DDR2 devices in 2 GB DIMMs (these two tasks can be pursued in parallel using the existing prototype boards). After identifying outliers we will also perform limited life testing.

Considerable FY12 effort was spent on firmware modifications to support DIMM testing and increased reliability of operation of the D2RT for functional testing. The majority of this is completed, but a few items remain. The modifications to support fixed and pseudo-random patterns have been accomplished. However, minor issues remain with the Samsung DIMMs due to low-speed operation (which is used to evaluate cell data retention only).

The DIMM mezzanine card developed in FY12 has been debugged, identifying design and implementation problems. Using the reworked prototype boards we have reached an almost 100% functional point on all boards. The updated card will feature all the fixes in the prototype boards, as well as a redesign of the termination power supply and the clock routing. This should allow operation at the minimum datasheet clock rate of 125 MHz, which will also require firmware modifications (but for cell retention measurements, we have moved to low-speed operation with plans to improve in the future with firmware revisions, so this clock rate is not required).

7.2 DDR3 Capability

Expansion of the DDR2 hardware to support DDR3 should be performed at the earliest possible time without risking the DDR2 hardware development and testing. This is not expected to be as major an effort as updating the single-DUT DDR2 system to test DIMMs. The majority of the debugging issues with bringing up the DDR2 DIMM system had to do with expanding the power capabilities, data bus size, and modifying the firmware to support more reliable data transfer with the DUT. All of these are similar issues in DDR3 DIMM implementation

The DDR3 hardware approach is essentially the same as the DDR2 approach. The functional tester provides stress during life testing and measurements of key functional parameters of the cell array. But it is largely unable to measure performance and high-speed data transfer reliability. An update to the Eureka 2 tester can be employed to test DDR3 memory, and we expect to obtain this update in the future of this task.

7.3 Migration to New Hardware Platform

The D2RT is built on the Modular Digital Test System (MDTS) prototype board 3b (MPB3b). This is a Xilinx Virtex 4 LX60-based board that is several years old. Its part number is DS-BD-V4LX60MB, from Memec Design, and it is no longer possible to order it. This means that we have a constellation of test equipment that is slowly decaying with no means to maintain or expand it as is. The most appropriate approach for the future of the task is to migrate to a newer development platform.

While new boards will alleviate resource concerns and improve overall designs by using the Virtex 5 with programmable output delay (Virtex 4 has only programmable input delay), there is no viable approach for implementing DDR3 at high speed (specification maximum) with any of these devices. Virtex 7 is able to support up to 1600MHz data rate [7], but the kits available today run more than \$3000, and 1600 MHz is not the specification maximum for many DDR3 devices. Thus, although we believe an upgrade is needed it is still clear that maximum datasheet parameters will not be measured with the new system. Instead, the new evaluation boards will function much as the current evaluation boards do—principally as functional

testers to investigate the cell-level storage capabilities. High-speed capabilities would still be handled by an industrial lot acceptance tester such as the Eureka 2.

8.0 REFERENCES

- [1] Schroeder, B.; Pinheiro, E.; Weber, W.-D.; “DRAM Errors in the Wild: A Large-Scale Field Study,” Sigmetrics/Performance, 2009
- [2] Guertin, S. M.; “FY11 End of Year Report for NEPP DDR2 Reliability,” JPL Publication 12-11, 2012
- [3] J.N. Bowles-Martinez and S.M. Guertin, “Interim Test Status Report for NEPP Scaled CMOS,” JPL publication 10-13, (2011)
- [4] Micron 1Gb DDR2 datasheet: “DDR2 SDRAM MT47H256M4, ...”, Rev. X, 10/11 (2007)
- [5] Samsung 1Gb DDR2 datasheet: “1 Gb F-die DDR2 SDRAM”, Rev. 1.2, Samsung Electronics (July, 2011)
- [6] Hynix 1Gb DDR2 datasheet: “1Gb DDR2 SDRAM H5PS1G43EFR, ...”, Rev. 0.4, (November, 2008)
- [7] Xilinx Virtex 7 data sheet 183 v1.4

APPENDIX A. ACRONYMS AND ABBREVIATIONS

ADC	address, data, and control
CMOS	complementary metal oxide semiconductor
DDD	displacement damage dose
DIMM	dual inline memory module
DQ	data line where Q is 0-7
DUT	device under test
FBGA	fine ball grid array
FPGA	field programmable gate array
FSM	finite-state machine
GSFC	Goddard Space Flight Center
IDD	total device current
IDD(q)	I _{dd} drawn by device while in operating mode q.
I/O	input/output
JPL	Jet Propulsion Laboratory
LCDT	low-cost digital tester
MCB	mezzanine card B
MCA	mezzanine card A
NEPP	NASA Electronic Parts and Packaging
SSTL	Stub Series Terminated Logic
TID	total ionizing dose
TBC	to be confirmed
TBD	to be determined

REPORT DOCUMENTATION PAGE

*Form Approved
OMB No. 0704-0188*

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) 01-01-2013	2. REPORT TYPE JPL Publication	3. DATES COVERED (From - To) N/A
--	--	--

4. TITLE AND SUBTITLE FY12 End of Year Report for NEPP DDR2 Reliability	5a. CONTRACT NUMBER NAS7-03001
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S) Guertin, Steven	5d. PROJECT NUMBER 103982
	5e. TASK NUMBER 03.02.02
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Jet Propulsion Laboratory California Institute of Technology 4800 Oak Grove Drive Pasadena, CA 91009	8. PERFORMING ORGANIZATION REPORT NUMBER JPL Publication 13-1
--	---

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) National Aeronautics and Space Administration Washington, DC 20546-0001	10. SPONSORING/MONITOR'S ACRONYM(S) NASA NEPP
	11. SPONSORING/MONITORING REPORT NUMBER

12. DISTRIBUTION/AVAILABILITY STATEMENT
Unclassified—Unlimited

Subject Category 60 Computer Operations and Hardware

Availability: NASA CASI (301) 621-0390 Distribution: Nonstandard

13. SUPPLEMENTARY NOTES

14. ABSTRACT
This document reports the status of the NEPP Double Data Rate 2 (DDR2) Reliability effort for FY2012. The task expanded the focus of evaluating reliability effects targeted for device examination. FY11 work highlighted the need to test many more parts and to examine more operating conditions, in order to provide useful recommendations for NASA users of these devices. In order to develop these approaches, it is necessary to develop test capability that can identify reliability outliers. To do this we must test many devices to ensure outliers are in the sample, and we must develop characterization capability to measure many different parameters. For FY12 we increased capability for reliability characterization and sample size. We increased sample size by moving from loose devices to DIMMs with an approximate reduction of 20 to 50 times in terms of per DUT cost. By increasing sample size we have improved our ability to characterize devices that may be considered reliability outliers.
This report provides an update on the effort to improve DDR2 testing capability. Although focused on DDR2, the methods being used can be extended to DDR and DDR3 with relative ease.

15. SUBJECT TERMS
DDR2, Reliability, Data Retention, Temperature Stress, Test System Evaluation

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 28	19a. NAME OF RESPONSIBLE PERSON STI Help Desk at help@sti.nasa.gov
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code) (301) 621-0390