Memory Testing Failures

Models, Algorithms, Current Test Analysis, Competitive Analysis

White Paper

« *Windows memory testing is recommended for assisted support, but is not recommended for factory or repair use.* »

As part of our ongoing efforts to maintain our position as the world leader in diagnostics and system health we conduct studies of various hardware components. This paper is a brief summary of our most recent study of 802 DIMMs (Dual Inline Memory Modules) that were returned from notebook PCs deployed in the field. It includes sections on the testing methodology and their results of this phase of the study. The results include competitive analysis of memory testing packages as well as the time to detect the memory failures.

The PC-Doctor memory study is a work in progress as there is more data that can be collected and more conclusions drawn from that data.

Purpose of this Study

The purpose of this study was to:

1. Measure the coverage of the PC-Doctor memory tests and compare to our competitors.
2. Determine the optimal test time for memory.
3. Evaluate new test algorithms.
4. Improve the PC-Doctor memory tests.
5. Determine how memory failures manifest themselves in the Windows operating environment.

Conclusions

1. PC-Doctor had the highest fault detection rate
2. PC-Doctor was the only diagnostic to catch 100% of failures
3. For factory testing when the memory is assumed good and there is a very low fallout rate we recommend running the Memory Fault + Address Fault test which take 12 s/GB.
4. For call center or assisted support we recommend running Advanced Pattern 4x4 which takes 83 s/GB
5. For depot repair of field repair we recommend running Advanced Pattern 8x4 or when the memory is suspect which takes 4:19 min/GB to run.
6. Many memory faults only showed on a subset of the systems. When possible memory should be tested in the system where it is used.
7. Windows memory testing is recommended for assisted support, but is not recommended for factory or repair use.

Background and Methodology

Memory errors can be classified into two groups:

1. **Soft errors** randomly corrupt memory data due to external factors such as electrical or magnetic interference.
2. **Hard errors** are repeatable errors due to physical defects such as memory cell defects, data path corruption, address path corruption, or control logic faults.

This paper deals with detection of hard errors.

Definitions

- **Failure:** A failure was defined as a memory that failed any of the following tests: PC-Doctor Protected Mode memory test, Memtest 86+ version 2.11, Microsoft Memory Test version 0.4, Ultra X RST version 2.66.
- **Coverage:** Coverage was defined as the number of memories in our failed sample set that failed a test or test suite divided by the number of memories in the failed sample set.
- **CNC:** Could Not Continue. Some of the memories would cause system lock up during testing. These memories were referred to as CNC and considered failures in determining coverage; however they were removed from the test time sample.
- **NDF:** No Defect Found. Memory that did not have a failure.

Methodology

The memory was divided into three groups:

1. No boot
2. Failure
3. No failure

All the memories were installed in systems that matched the speed rating of the memory being tested. For initial failure determination, multiple memories were installed in systems. If a failure was detected, the memories were tested individually to determine which memory was failing.

> *« With respect to the PC-Doctor tests, about 95% of the failures can be detected in 7 minutes 30 seconds per gig of testing…»*

There were three memories that appeared to have temperature-related failures. One of these memories would only fail when it was cold; the other two would only fail when the system was hot. These three memories were removed from the sample and set aside for further study.

After the failing memories were identified, testing was focused on each of these memories. Once a memory was determined defective, the memory was always tested in the same system to remove system differences.

- Each defective memory was installed in a system and tested.
- If a memory test suite offered the ability to run individual tests, the individual tests were run and the result and test time was collected.
- The Dell MpMemory test version 0463 was added to the data collection.
- New algorithms developed by PC-Doctor engineers were tested.
- PC-Doctor purchased a hardware memory tester and ran the failing DIMMS on the hardware memory tester.
- Different settings for the PC-Doctor advanced pattern test were tested.
- Patterns used by the PC-Doctor memory tests were modified to determine the optimal set of patterns.
- The seeds for the PC-Doctor random pattern tests were studied to determine if certain seeds would be effective for determining memory failures.
- Windows was started on the systems with defective memory.
- On the systems that would start Windows and Windows would idle without a lockup or blue screen, the PC-Doctor for Windows memory test was run.
- Finally samples of both good and bad memories were used in systems running Windows and the systems were stressed using Prime95.

Once the data was collected, the results were analyzed and coverage numbers were calculated. Test times were collected and normalized to 1 Gig memories. Failures were analyzed to look for patterns in failing addresses and failing bits.

Test Systems

The following systems and memory types were used:

System	Recommended Memory Type
1	DDR-333
2	DDR-333
3	DDR-333
4	DDR2-533
5	DDR2-533
6	DDR2-533
7	DDR2-667
8	DDR2-667
9	DDR2-667
10	DDR2-667
11	DDR2-667
12	DDR2-667
13	DDR2-667
14	DDR2-667
15	DDR2-667
16	DDR2-533
17	DDR2-667
18	DDR2-533
19	DDR3-1066
20	DDR2-667
22	DDR2-667
23	DDR3-1066
24	DDR3-1066

Observations

NDF Rate

The 802 memories included in this study, all returned from the field as defective, fell into these categories:

Total SODIMMS tested	802
No Boot	120
Test Failure	113
Total failures	233
No Defect Found	569
NDF Percentage	70.95%

There are several possible reasons (in no particular order) for this apparently high NDF rate:

« Memory that fails in one system may pass in another system.»

- Repair personnel are replacing multiple items at the same time. For example, simultaneous replacement of the motherboard, CPU, and memory is a common method of fixing systems that do not boot.
- Repair technicians may be replacing memory without testing. Troubleshooters may associate certain symptoms of computer problems with memory problems and may replace the memory when these symptoms are seen without testing to ensure that the memory is faulty.
- Memory that fails in one system may pass in another system. This was observed during the initial failure determination. If memory was moved to a system with a slower front side bus speed some of the memories would no longer fail. Also, some memories were observed to fail in one system and pass in another system with the same front side bus speed. Finally one memory was observed to fail only when paired with another memory and operating in dual channel mode. To control these outliers, once a failure was determined, the failing memory was always tested in the same system in which the failure was determined. The hardware memory tester, discussed later in this paper, also gives evidence that system variability affects memory test results.

Diagnostic Comparison

This table summarizes the study's findings concerning test coverage and average test times for the diagnostic software tested.

Test	Coverage	Avg Test Time / Gig
PC-Doctor Protected Mode Tests	100.00%	1:20:16
MemTest 86+	98.23%	0:21:12
Ultra X	94.74%	1:11:44
Dell MpMemory tests	94.50%	0:25:13
Microsoft (Extended)	89.47%	1:27:29
PC-Doctor For Windows	61.11%	0:03:00

Test Times

The PC-Doctor advanced pattern test has two configurable parameters. The test time and coverage can be adjusted by changing the number of phases and the number of patterns. The following table shows the test times and coverage of each of the PC-Doctor memory tests and adjustments of the advanced pattern test phases and patterns between 4, 8, and 16. The nomenclature of the settings is [phase]x[pattern]. Thus the 8x16 settings would use 8 phases and 16 patterns. The number of passes through memory for the Random Pattern test is also configurable and was set to 40 for this data collection.

Test	Coverage	Avg Test Time / Gig
Memory Fault	73.83%	0:00:04
Address Fault	79.44%	0:00:08
MF + AF	80.37%	0:00:12
Random Pattern	87.85%	0:13:01
4x4	87.62%	0:01:23
4x8	91.43%	0:03:09
4x16	93.33%	0:06:41
7x4	92.52%	0:04:37
8x4	93.33%	0:04:19
8x8	93.33%	0:09:55
8x16	96.15%	0:21:09
16x4	93.27%	0:13:36
16x8	94.23%	0:31:27
16x16	98.13%	1:07:03
Random Pattern + 16x16	100.00%	1:20:06

The above table shows that most of the errors (80.37%) are caught in a very short test time (12 seconds / gig). The time to catch the remaining errors is a logarithmic distribution.

The two outliers on the below graph are the random pattern test and the 16x8 advanced pattern test, which take more time than the others.

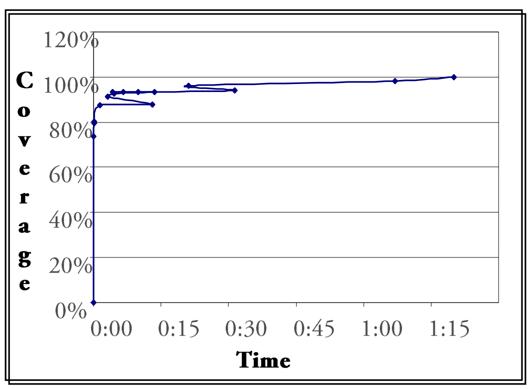

> « *Approximately 20% of the 113 memories that had test failures took more than 15 seconds/gig to detect.*»

Based on the time and coverage data shown above, the table below shows the recommended use of PC-Doctor tests during each stage of the PC lifecycle.

Test Script	Cover -age	Use	Tests	Test Time Per GB
Quick	80%	Factory	Address Fault Memory Fault	12sec
Thorough	88%	Support	Address Fault Memory Fault Advanced Pattern 4x4	91sec
Extended	93%	Repair	Address Fault Memory Fault Advanced Pattern 8x4	4.5min
Super	100%	Design	Address Fault Memory Fault Random Pattern Advanced Pattern 16x16	80 min

Detection Rates and Times

Despite the large sample size of 803 returned memories, the number of interesting failures (those requiring more than 15 seconds/gig to detect) is quite small because of the high NDF rate and the fact that most memory errors are easily and quickly detected.. Approximately 20% of the 113 memories that had test failures took more than 15 seconds/gig to detect. PC-Doctor tests detected a maximum 10 defects that were missed by Competitors. The difference between PC-Doctor and Dell MpMemory test or Memtest 86+ was 4 memories.

New Algorithms and Test Patterns

PC-Doctor evaluated over 50 algorithms for coverage and time effectiveness. As a result of this evaluation, PC-Doctor added the random pattern test to its test suite. Two specific memory defects were detected *only* by this random pattern test. For this data collection, the random pattern test was set to 40 passes through the memory. The time to failure data shows that the random pattern test fails at less than half of the time to pass. Thus the number of passes in the random pattern test could be reduced to 20, which would reduce the test time by half.

MATS	SOA-March C	March M	Algorithm A
MATS +	SOA-MarchC -	March PS	Algorithm B
SOA-MATS +	SOA-CFst	March SR	March 1/0
MATS ++	March G	March SS	March 6N
March X	March G+	March U	March 9N
March A	March GS	March U-	IFA-9N
March B	March LA	March UD	March 12
March B+	March LA-	March UD-	IFA-13N
SOA-March B-	March LAD	March Y	March 17
March AB1	March LADD-	March RAW	March 18
March AB	March LR	March RAW1	March SL
March C	March LRD	MOVI	
March C-	March LRDD	Marine scu A	

Compared to the existing PC-Doctor memory fault and address fault tests, none of the algorithms proved more effective for coverage vs. time than the existing tests.

Many patterns with varying numbers of zeros were evaluated for effectiveness for the advanced pattern test. These patterns, such as 0000, AAAA, 9249, 8888, 8421, and so on, did not improve test coverage or time.

« 53% of the defective memories caused catastrophic errors in Windows (blue screens, lockups, and so on).»

Hardware Memory Tester

PC-Doctor purchased a RAMCHECK DDR2 tester and a DDR2 SODIMM converter for this study. The average test time with this tester was 1 minute 40 seconds. The coverage was only 15.4%. The log files of the modules tested showed many modules had retests. After contacting the vendor of this memory tester, application note INN-8668-APN33 was given as an explanation of retest entries in the log. In part, the application note said:

It should be emphasized again that having some retests does not necessarily mean that the module is bad. It is only a comparison tool used by advanced users to compare modules.

The following is a brief description of the retest mechanism. When the RAMCHECK / RAMCHECK LX encounter an error while testing a block of memory, it tries to verify the error by repetition of the test of the same block several times (depending on the actual test phase). Every repeated test increments the retest counter. If the memory block passes the test during these repetitions, the test proceeds to the next memory block. If the error persists, the tester determines whether this is a fatal error or if the frequency is too tight and needs to be reduced. Often, retesting the same block at the reduced frequency is successful and the full RAMCHECK test completes successfully.

If modules with retests are considered failures (against the advice of the app note), then the coverage of the hardware tester is raised to 85%, which is on par with the coverage vs. time exponential curve presented above. This provides evidence that many of the memory failures are sensitive to frequency and timing changes. This evidence is consistent with the observations that some memory failures were not repeatable when the memory was moved across systems.

Memory Testing in Windows

This study determined two significant facts about how memory failures manifest themselves in the Windows operating environment:

1. 53% of the defective memories caused catastrophic errors in Windows (blue screens, lockups, and so on)
2. Of the remaining defective memories, 70% allowed prime95 torture test to run for 24 hours with no errors.

These facts illustrate how difficult it can be to detect defective memory in the Windows environment.

In those systems that would start Windows and allow the operating system to idle with no Blue Screen or lockups, the PC-Doctor for Windows memory test detected 61% of the defective memory modules. Including the defective memories that would not allow Windows to run, the PC-Doctor failure detection rate goes up to 82%.

Windows operating systems, including Win PE, are not recommended for testing memory in the factory, since the memory tests in non-Windows operating systems provide higher coverage and take less test time.

Windows environments are, however, recommended for troubleshooting system problems on customers' systems. Based on our sample data, systems running windows that pass the PC-Doctor for Windows memory test will have no defect 82% of the time.

« Although the modules failed repeatedly, many of the modules failed at different addresses and/or bits.»

Failure Variability

The random pattern test was run multiple times using different seeds for the random number generator. Although the modules failed repeatedly, many of the modules failed at different addresses and/or bits. In our sample, 23% of the memories always failed at the same address, 31% of the memories always failed the same bit, and only 6% always failed the same address and bit. The failing addresses were always near one another and the bits were usually shifted by one bit. Two example memories with varying address and failure bits are shown below:

Failure Address	Expected	Actual	Expected - Actual
0424D 4008	98CF F9D9	9847F 9D9	00880000
0424D 4008	A3CA 3213	A34A 3213	00800000
0424D 4008	250C A577	2504A 577	00080000
0424D 4028	B0FA FE8E	B07A FE8E	00800000
0424D 4008	13CC 3DB5	13C4 3DB5	00080000
0424D 4068	9D69 679E	9D60 679E	00090000
0424D 4008	A7EA 87D7	A7E2 87D7	00080000
0424D 4048	0B87 FBEB	0B07 FBEB	00800000
0424D 4008	C65E D14C	C656 D14C	00080000
0424D 4008	95842 BAD	95042 BAD	00800000
01C4E FF68	EE9E 36A1	EE9E 36A5	00000004
01C4E FF48	00512 4911	00512 4951	00000040
01C4E FF48	D6A6 09AB	D6A6 09EB	00000040
01C4E FF48	A123 E948	A123 E94E	00000006
01C4E FF28	E698 E63A	E698 E67A	00000040
01C4E FF28	2A7E C414	2A7E C450	0000003C
01C4E FF28	F66B A88F	F66B A8CF	00000040
01C4E FF48	3423 DB8F	3423 DBEF	00000060
01C4E FF68	2A56 4802	2A564 806	00000004
01C4E FF48	0710 A8DE	0710A 8FE	00000020

These failures indicate that the memory failures are not strictly caused by cell or control logic errors, but rather that a particular area of cells is more susceptible to noise or influence from neighboring cells.

Related Work

Test time and coverage are clearly related. Infinite test time would always yield 100% test coverage, as there would be no escaping defects. However, it is important to balance test coverage and test time. PC-Doctor partnered with the University of Reno, Nevada to analyze the resulting data from the memory study. Dr. Ilya Zaliapin, Associate Professor in the Department of Mathematics and Statistics of UNR published a paper titled *Detection of Memory Failures: Analysis of Test Duration and Choosing the Optimal Test Strategy*. This paper develops a model to determine optimal test time. The model includes the cost of test setup, the cost of test time, and the cost of an escaping defect. The most cost-effective PC-Doctor test time can be selected using these three parameters. The publication is available by request from PC-Doctor.

DRAM Errors in the Wild: A Large-Scale Field Study, is a paper investigating ECC corrected memory errors and uncorrected memory errors on Google servers. This paper draws the following conclusions:

- DRAM errors are orders of magnitude higher than previously reported with 25,000 to 70,000 errors per billion device hours per Mbit.
- More than 8% of DIMMS were affected by errors per year.
- Strong evidence is presented that the memory errors are dominated by hard errors rather than soft errors.
- A small number of systems are responsible for most of the memory errors. 20% of the machines with errors make up more than 90% of all observed errors.

PC-Doctor, Inc.
The global leader in system health

5485 Reno Corporate Drive
Reno, Nevada 89511

Phone: 1.775.336.4000
Fax: 1.775.336.4099
E-mail: sales@pc-doctor.com
www.pc-doctor.com