Summary of the Radiation Testing of the Intel Pentium III (P3) Microprocessor

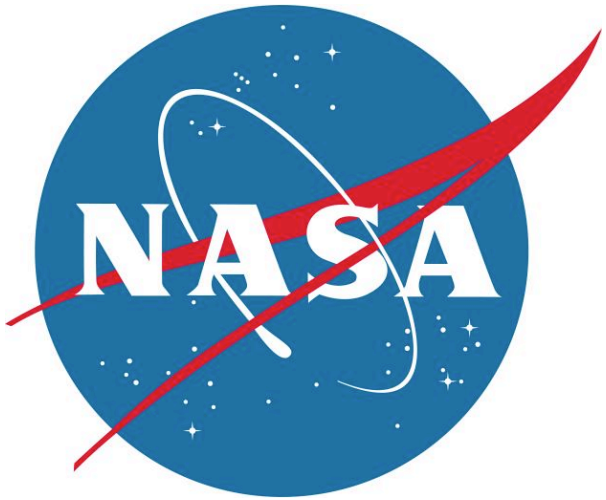

James W. Howard Jr.

Jackson and Tull
Chartered Engineers
Washington, D.C.

Martin A. Carts
Ronald Stattel
Charles E. Rogers

Raytheon/ITSS
Lanham, Maryland

Kenneth A. LaBel

NASA/GSFC Code 561
Greenbelt, Maryland

Timothy L. Irwin

QSS Group, Inc.
Lanham, Maryland

Outline

- Introduction
- Test Methodologies
- Hardware
- Software
- Test Issues
- Sample Data
- Summary

Introduction

- Many future NASA missions will require extensive on-board computation capability which raises the issue of availability, cost and capability of radiation hardened or radiation tolerant microprocessor systems. Radiation hardened computer systems are often costly and are actually two or three generations behind in computational capability, a significant shortfall for missions that may require state-of-the-art (SOTA) capability.
- To confront these issues, NASA instituted the Remote Exploration and Experimentation Project (REE) with the goal of transferring commercial supercomputer technology into space using SOTA, low power, non-radiation-hardened, commercial-off-the-shelf (COTS) hardware and software to the maximum extent possible. Testing of the Pentium III microprocessor was done with this effort in mind.

Test Methodologies

- Total Ionizing Dose/Displacement Damage Dose
  - Level of performance testing
    - Functional failure levels, timing errors, power draw
    - Biased vs. Unbiased, idle vs. operating
  - Changing technology testing
    - 0.25 μm vs. 0.18 μm
    - Design and operation speeds

Test Methodologies

- Single Event Effects
  - Architecture and technology implications
    - Test SOTA technology and exercise that technology
    - Exercise independent pieces of the architecture with maximum duty cycle
    - Investigate technology versus operational conditions (e.g., rated versus operation clock speeds)
  - System level impacts
    - Destructive events
    - Function interrupts vs. non-recoverable upsets vs. recoverable upsets

P3 Test Controller System

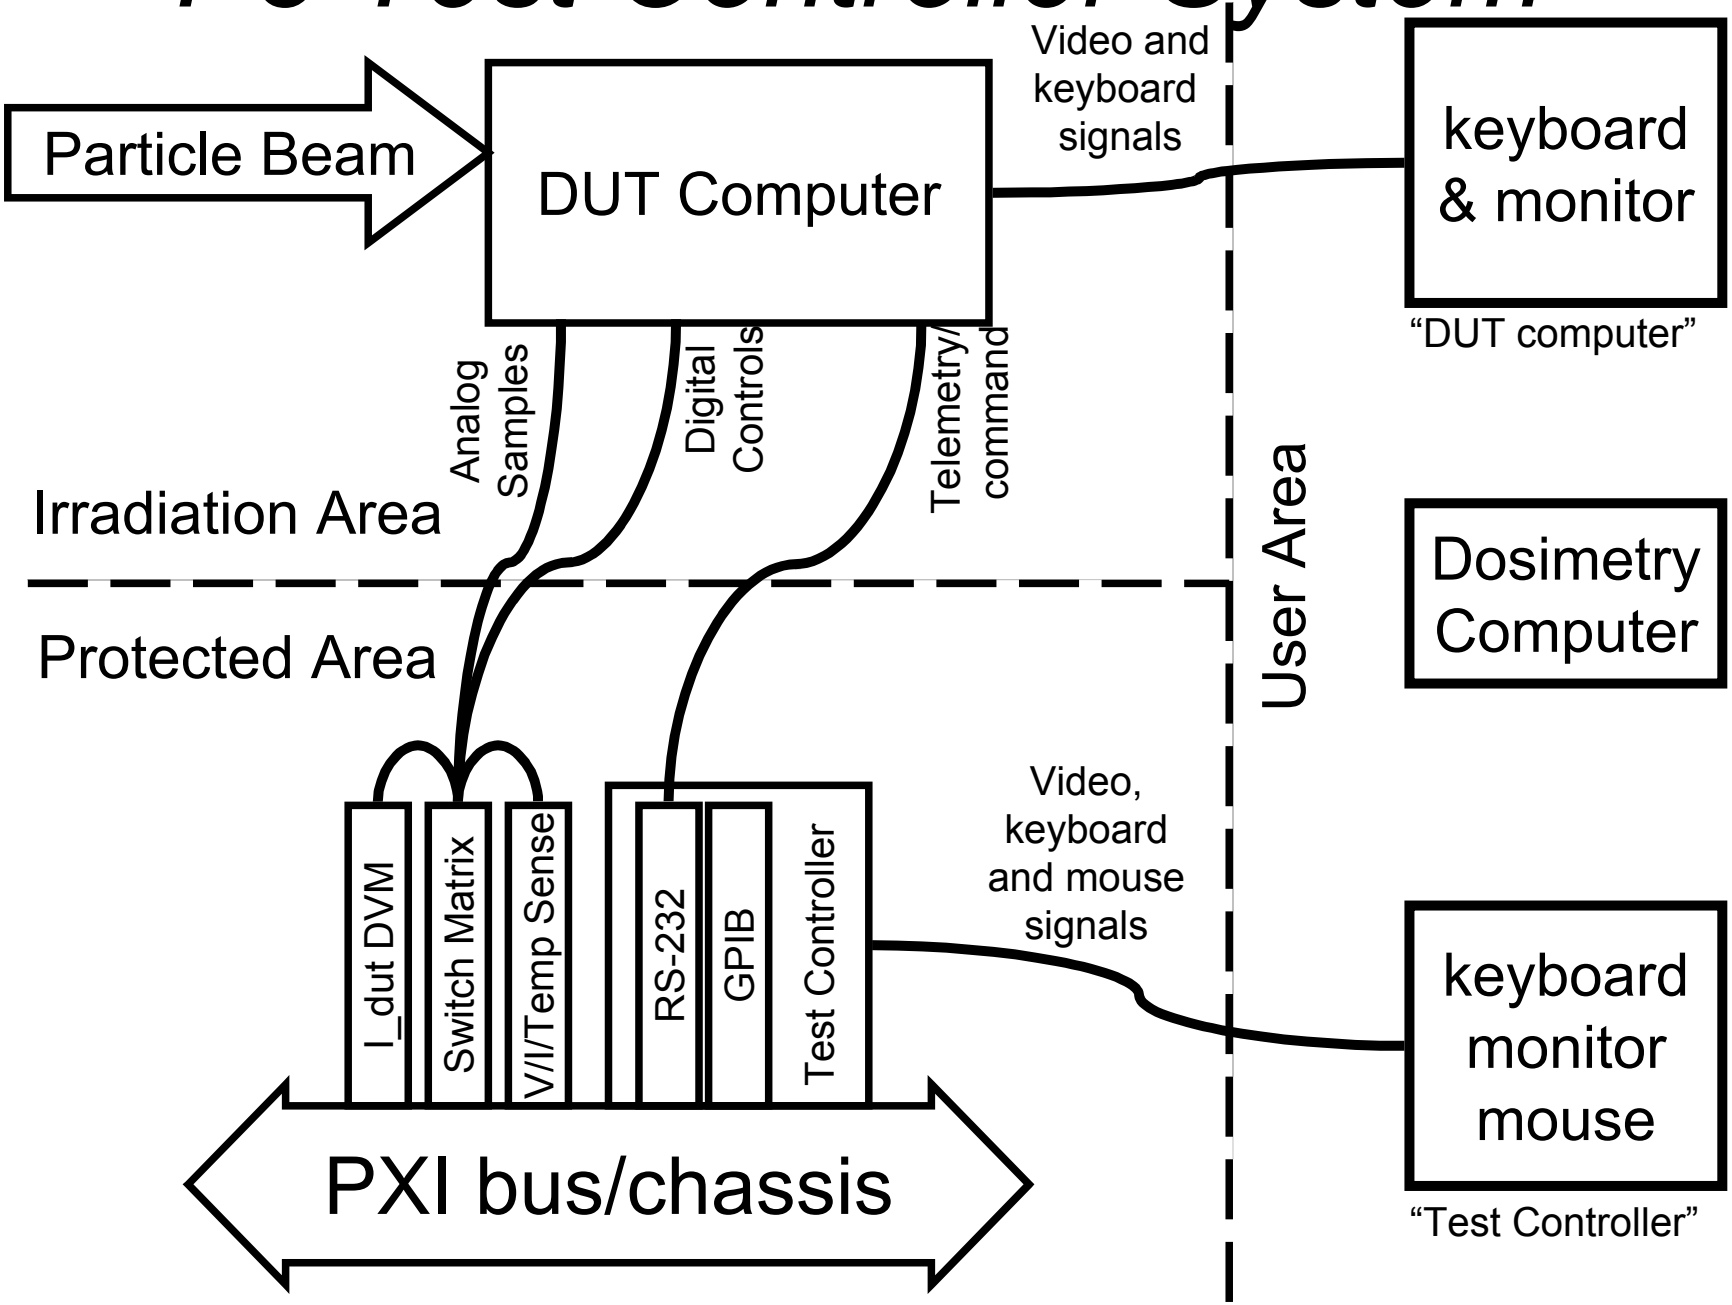

- The system hardware consists of two subsystems, the Test Controller and the DUT computer subsystems. The above figure illustrates the overall test configuration. The Test Controller hardware is based on the PXI specification. The PXI subsystem resides within the irradiation chamber but is removed from the DUT subsystem by more than 15 feet. It consists of the PXI components, the PXI Computer <-> DUT Computer cabling and the user interface.
- The PXI components include the PXI chassis, which contains an embedded controller (running WinNT, Labview (LV) environment, and a custom LV application), a signal switch matrix, and two digital multimeters (DMMs) in the voltage measurement mode. The switch matrix provides two functions: The multiplexing of analog signals to one of the DMMs, and contact closures (pulling signal levels to ground). The other DMM is dedicated to monitoring one specific analog value and measuring that without regular periodically switching, so that it may measure more frequently with less delay.
- The user interface for the PXI Computer, network connectivity (for data file access) and AC power feed are also components of the PXI Computer. An extended (via a CAT5 cable based extender from Cybex, Inc.) keyboard, monitor, mouse user interface provides user control of the PXI computer from the user facility. Most of the PXI Computer <-> DUT Computer cabling leaves the PXI subsystem from the switch matrix, exceptions are the AC power cable to power the DUT computer and a serial (RS-232) cable for telemetry/command of the DUT computer (telemetry originates within the DUT Computer, commands originate within the PXI Controller).

P3 DUT Computer Test System

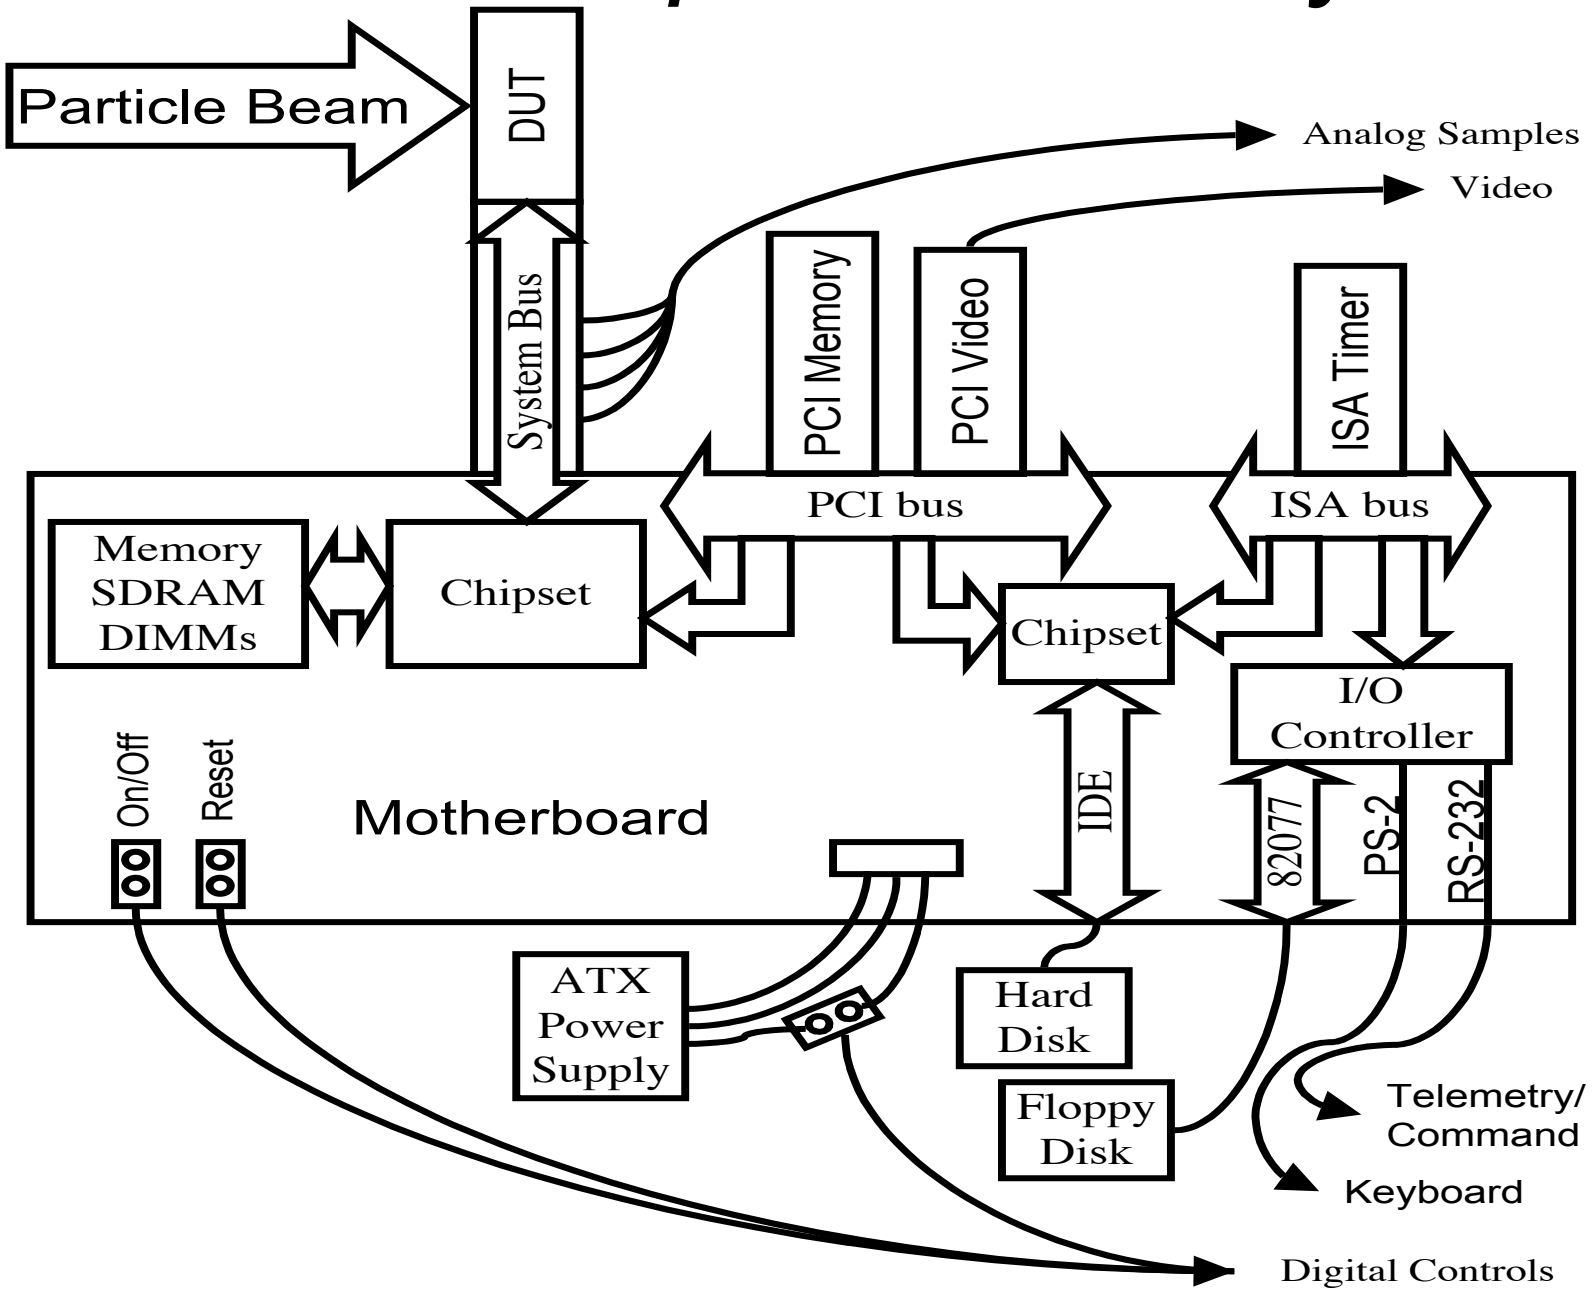

- The DUT computer subsystem consists of the components immediately connected with the operation of the DUT computer, including components mounted directly to the motherboard, components located nearby (e.g. disk drives) and connected via cables, and the user interface.
- Two different motherboard types were used (one for the SC-242 DUT and one for the FC-PGA370 and FC-PGA2 DUTs). Directly attached to one of these commercial motherboards are the modified P3 DUT processor (on top of a DUT processor extender card in the case of the SC-242 DUT), a RAM module (DIMM), the video, and a PCI-bus memory board (This board and the serial port are handed identical copies of the telemetry stream; the contents of the memory board will survive a soft reset of the DUT computer for later readout through the RS-232 port in case the RS-232 port crashes.).
- The DUT computer motherboard resides in the test chamber, positioned directly in front of the particle beam but so that only the DUT processor is irradiated. A metal base with standoffs holds the SC-242 motherboard horizontally so that the DUT sits vertically, high on its extender board. Only the DUT and its cooling system are irradiated. In contrast, the FC-PGA370 motherboard sits vertically, with its “front” edge in an (insulated) slot cut into the aforementioned metal base, with the top stabilized by a hinged support, so that the DUT, which is plugged directly into the motherboard socket, is also vertical. The proton beam penetrates the DUT, its cooling solution, and the motherboard but no active components are affected.

Programming Environment

- The DUT Software is written in the Microsoft Visual C++ environment with a Pharlap Add-in.
- Tests are written in a combination of C and Assembly Language.
- The software is executed on the DUT using the Pharlap Real-Time Operating System.
  - Pharlap was chosen for its low overhead, preemptive multithreading, short interrupt latency, and price.
  - The kernel has been stripped to its minimal functionality so that boot time is minimized.
  - Kernel interrupts have been disabled to allow the test running full attention of the processor.

DUT Tests

There are eight tests designed to exercise the various aspects of the CPU during SEE testing:

A: Register Test

B: Floating Point Register Test

C: Memory/Data Cache Test - Sequential

D: Task Switching Test

E: Instruction Cache Test

F: Floating Point Operation Test

G: MMX Test

I: Memory/Data Cache Test - Offset

- Test "A" checks the general-purpose registers of the CPU. There are eight general-purpose registers.
- Test "B" checks the Floating Point Unit (FPU) with a maximum of data transfer to the FPU. A buffer is loaded with the arguments and expected results for the five operations tested (fadd, fsub, fmul, fdiv, and fsqrt).
- Test C performs cache test. If the cache is on, then the cache is turned off, the memory is loaded with an incrementing pattern and the cache is turned back on before entering the test. The memory is checked word by word. After each word is checked, its value is changed to the bitwise complement of the baseline. If the value is not as expected then an error is reported and an attempt is made to reset the value to the baseline.
- Test "D" launches seven subthreads each with a counter that is reset to zero. Each thread increments its counter if the counter is less than 11 and then passes control to the next thread. The main thread then checks after 50 milliseconds to see if all of the counters have reached 11. If not an error is reported to the user.
- Test "E" runs through 16K of instructions repeatedly. The instruction sequence is to increment the eax register from 0 to 3 checking in between each increment to see if the value is as expected, then to decrement the eax register 3 times and check to make sure it returns to zero. Any errors are reported to the user and the cache is invalidated.
- Test "F" checks the Floating Point Unit (FPU) with a maximum of operations in the FPU. A buffer is loaded with the arguments and expected result.
- Test "G" checks the Matrix Math Extensions (MMX). A buffer is loaded with the arguments and expected results for the four operations tested (pxor, por, pmul, pmulh, padds, addps, divps, and mulps).
- Test “I” performs an alternate cache test. The difference between tests I and C is that bits are swapped in the location used to access the memory/cache. This has two effects. First, the area of cache used is different for the 50% and 25% cases. In test C 50% case, a contiguous half of the cache is used while in test I 50% every other tag (a tag marks 16 words) in the entire address range is used. Similarly for the 25% case.

Data Analysis Software

- GUI Interface
- Relational Database (3 Stages)
  - Setup data entered into database
    - Test configurations, software, dosimetry, etc.
  - Telemetry files analyzed and errors entered into database
    - Filter for allowed errors
    - Accuracy (Program shows possible errors and description)
  - SQL statements for filters to extract data

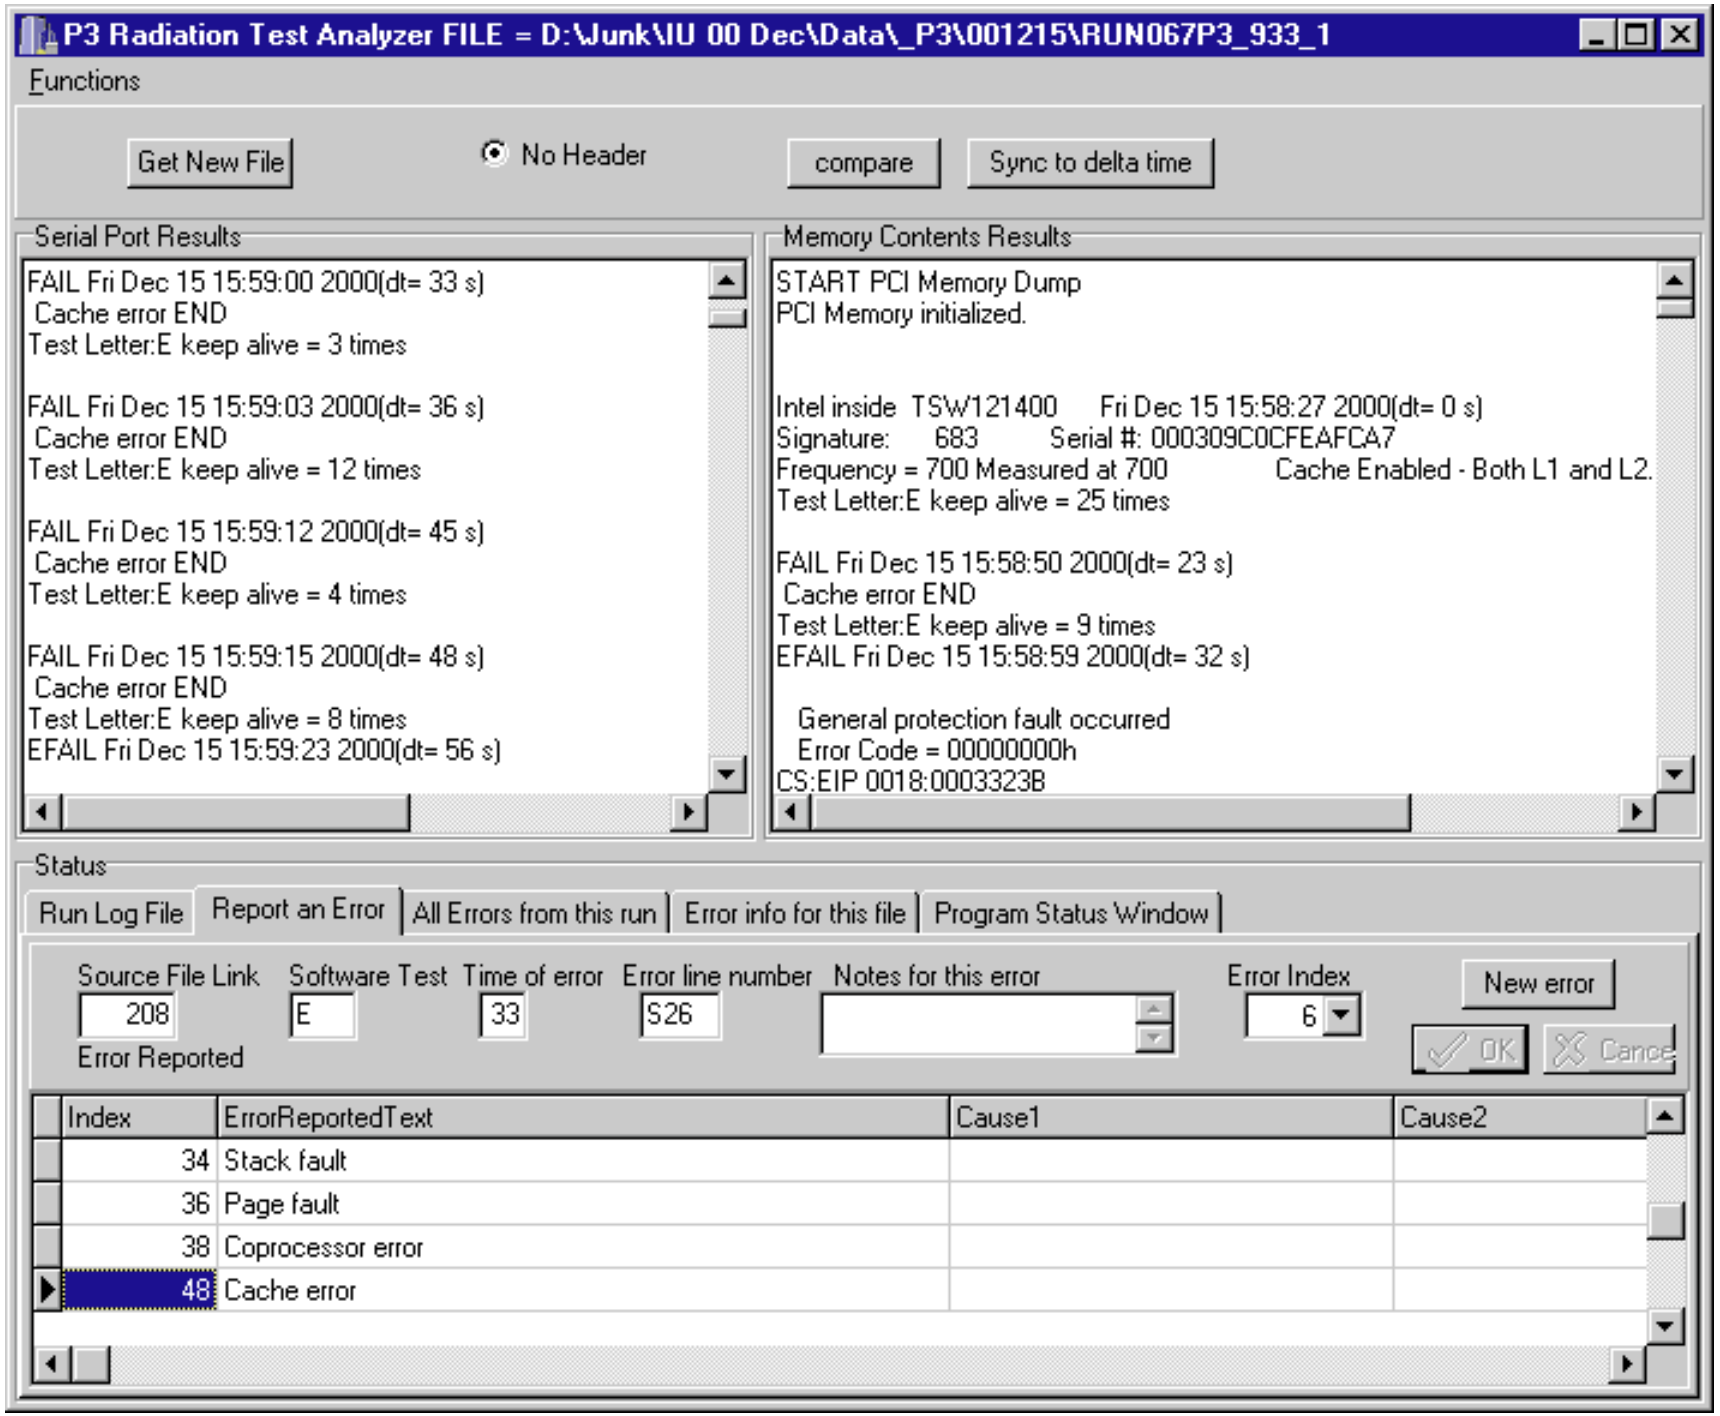
P3 Radiation Test Analyzer FILE = D:\Junk\IU 00 Dec\Data_P3\001215\RUN067P3_933_1
Functions
Get New File
No Header
compare
Sync to delta time
Serial Port Results
FAIL Fri Dec 15 15:59:00 2000(dt= 33 s)
Cache error END
Test Letter:E keep alive = 3 times
FAIL Fri Dec 15 15:59:03 2000(dt= 36 s)
Cache error END
Test Letter:E keep alive = 12 times
FAIL Fri Dec 15 15:59:12 2000(dt= 45 s)
Cache error END
Test Letter:E keep alive = 4 times
FAIL Fri Dec 15 15:59:15 2000(dt= 48 s)
Cache error END
Test Letter:E keep alive = 8 times
EFAIL Fri Dec 15 15:59:23 2000(dt= 56 s)
Memory Contents Results
START PCI Memory Dump
PCI Memory initialized.
Intel inside TSW121400 Fri Dec 15 15:58:27 2000(dt= 0 s)
Signature: 683 Serial #: 000309C0CFEAFCA7
Frequency = 700 Measured at 700 Cache Enabled - Both L1 and L2.
Test Letter:E keep alive = 25 times
FAIL Fri Dec 15 15:58:50 2000(dt= 23 s)
Cache error END
Test Letter:E keep alive = 9 times
EFAIL Fri Dec 15 15:58:59 2000(dt= 32 s)
General protection fault occurred
Error Code = 00000000h
CS:EIP 0018:0003323B
Status
Run Log File
Report an Error
All Errors from this run
Error info for this file
Program Status Window
Source File Link
208
Error Reported
Software Test
E
Time of error
33
Error line number
S26
Notes for this error
Error Index
6
New error
OK
Cancel
Index
ErrorReportedText
Cause1
Cause2
34 Stack fault
36 Page fault
38 Coprocessor error
48 Cache error

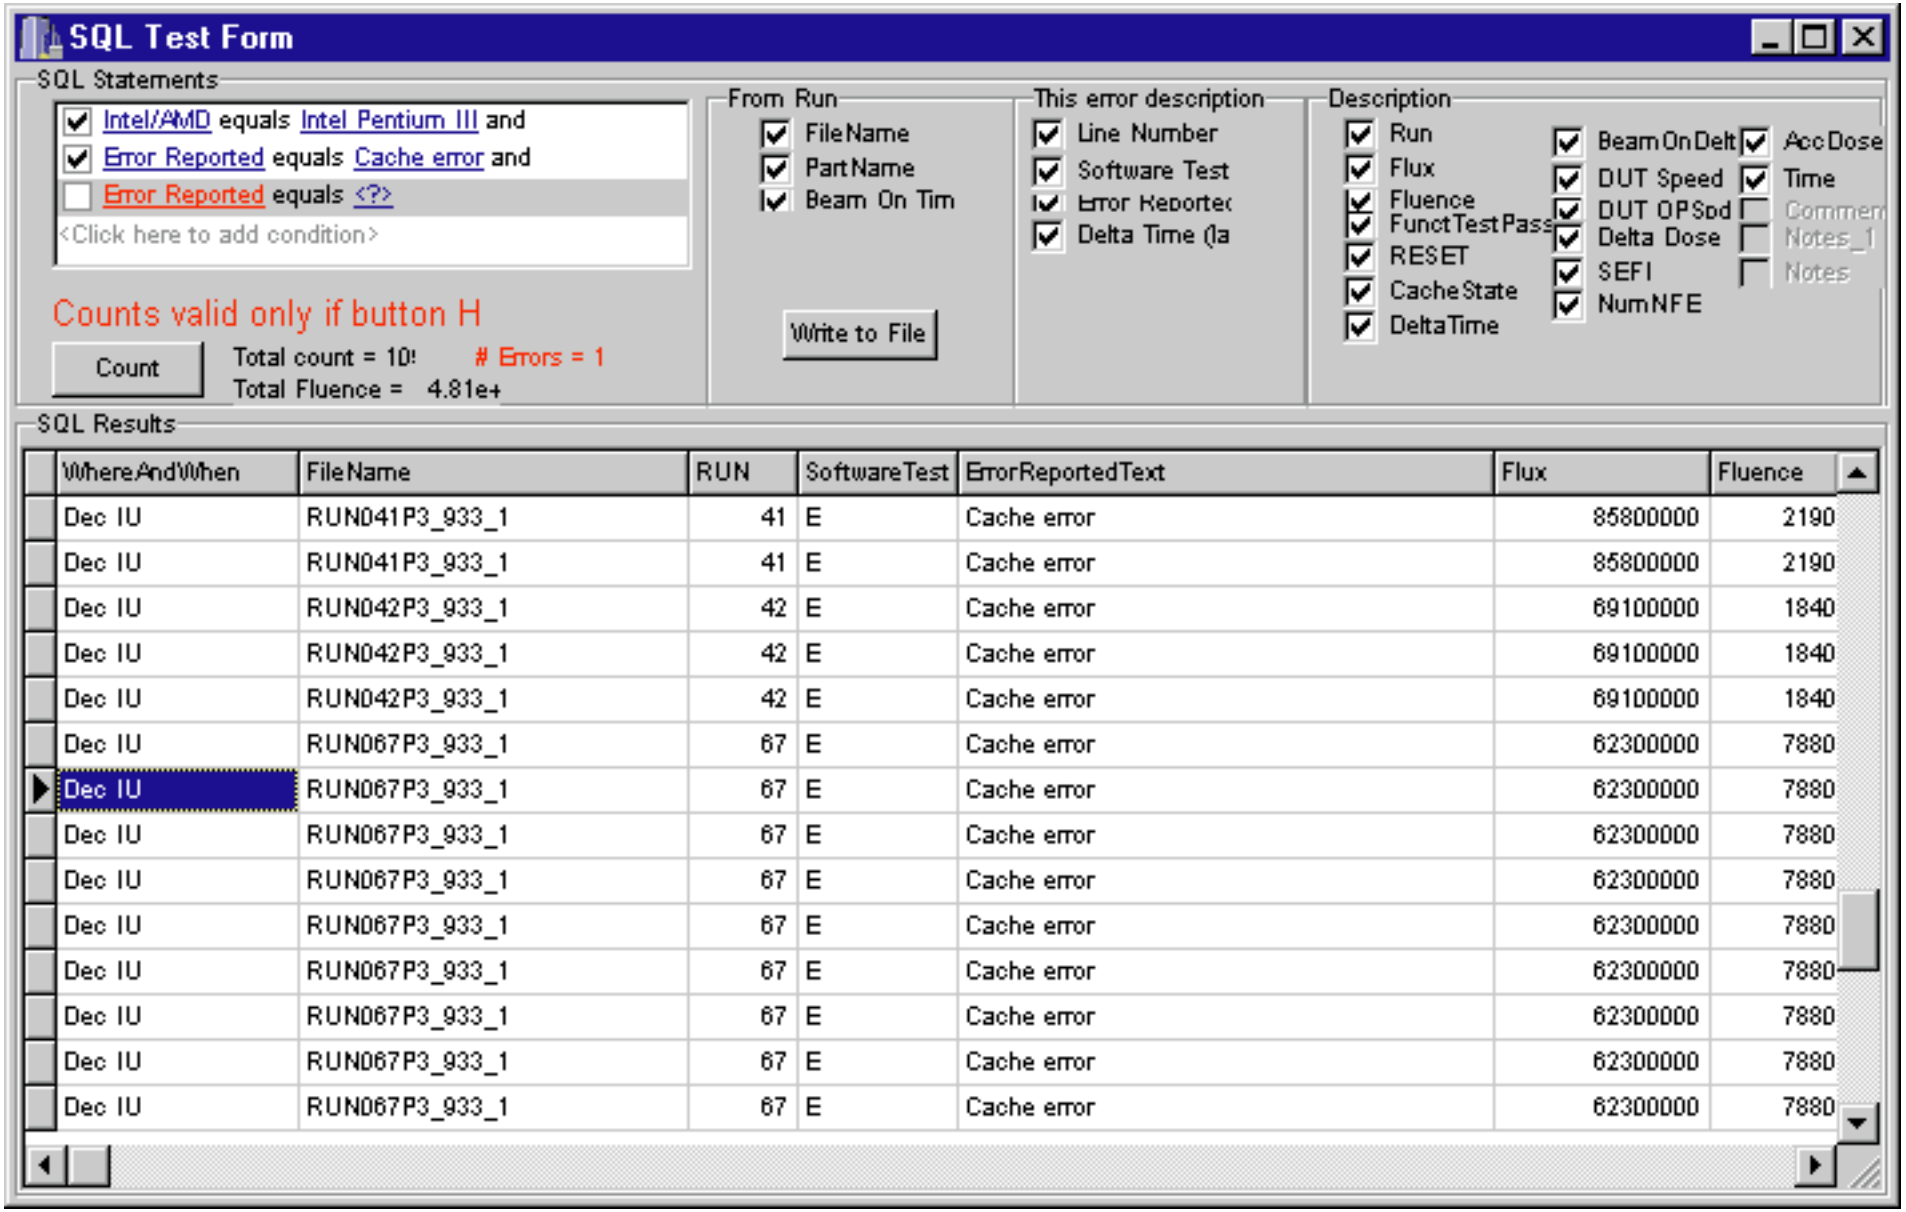
SQL Test Form
SQL Statements
Intel/AMD equals Intel Pentium III and
Error Reported equals Cache error and
Error Reported equals <?>
<Click here to add condition>
Counts valid only if button H
Count
Total count = 10!
Errors = 1
Total Fluence = 4.81e+
From Run
FileName
PartName
Beam On Tim
Write to File
This error description
Line Number
Software Test
Error Reporte
Delta Time (la
Description
Run
Flux
Fluence
FunctTestPass
RESET
CacheState
DeltaTime
BeamOnDelt
DUT Speed
DUT OPSpd
Delta Dose
SEFI
NumNFE
AccDose
Time
Commen
Notes_1
Notes
SQL Results
WhereAndWhen
FileName
RUN
SoftwareTest
ErrorReportedText
Flux
Fluence
Dec IU
RUN041P3_933_1
41
E
Cache error
85800000
2190
Dec IU
RUN042P3_933_1
42
E
Cache error
69100000
1840
Dec IU
RUN067P3_933_1
67
E
Cache error
62300000
7880

Testing Issues

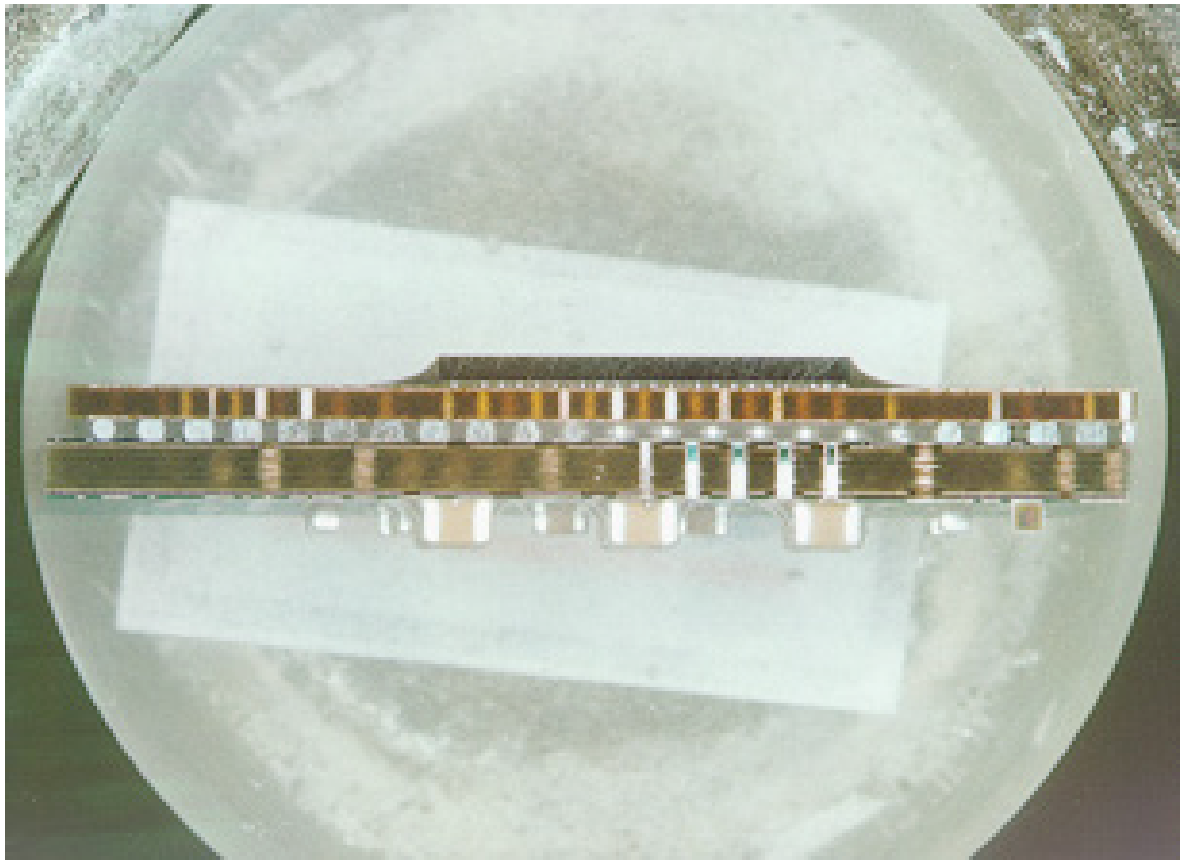

- Die Penetration
  - The Pentium III die is a flip chip solder bubble bonded die.
  - The sensitive regions of the processor are approximately 900 microns deep in the silicon die.
  - Thermal issues compound Heavy Ion Testing by requiring cooling material in the beam line, as well.

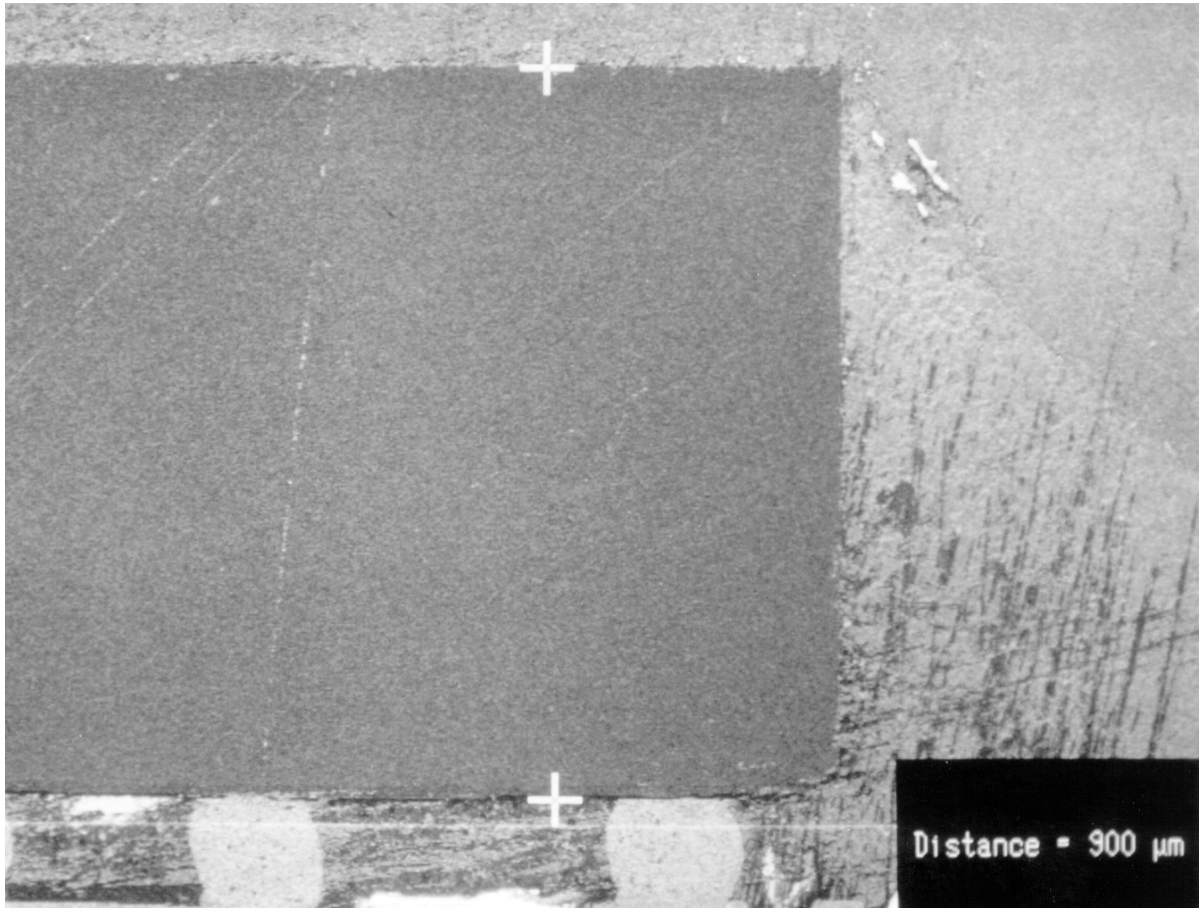

- Thermal
  - The Pentium III can draw in excess of 20 watts of power.
  - The packaged heat sink and cooling fan are removed and replaced with a water-cooled jacket, that is thinned to 10 mils over the die.
  - The large thermal issue is also the reason that the die cannot be thinned.

- The Pentium III, when operating under normal conditions with the caches enabled, will draw in excess of 20 watts of power. If left in that state with no cooling, the processor will not even boot. Also, the P3 die are flip chip solder bubble bonded die to the daughter cards. Since the beam must hit the die directly, the packaged heat sink and cooling fan must be removed. In its place a water-cooled jacket is used (thinned to 10 mils over the die for heavy ion testing).
- The large thermal source is also the reason that the die cannot be thinned, as has been done with other flip chip parts. The thick substrate is the main thermal path for removal of heat from the junctions. Thinning this would place excessive thermal stresses on the die and most likely lead to structural failures.
- With these parts being flip chip solder bubble bonded die, the sensitive regions of the processor approximately 900 microns deep in the silicon die with respect to the heavy ion incidence point. Thermal issues compound this by requiring cooling material in the beam line, as well. Therefore, only high energy and high Z beams are capable of penetrating and giving higher Linear Energy Transfer (LET) values in the sensitive regions.
- If the beam energy is such that normal incidence the ions do not have penetration issues into the sensitive region but when the beam is either degraded in energy or the part rotated (to give a larger effective LET), there are penetration issues. Roll-off of the cross section curves may then be observed (i.e., a higher effective LET data point having a lower cross section than the previous lower effective LET point). For the beam used in these tests, this roll-off point is expected to occur somewhere between an effective LET of 10-20 MeV-cm^2/mg.

What Have We Tested

- Intel Pentium III
  - Speed ranging from 550 through 1200 MHz
  - Represents 0.25, 0.18 and 0.13 μm technology

Where Have We Tested

- GSFC TID Facility
  - Biased and Unbiased Co-60 Testing
- Indiana University Cyclotron Facility
  - Proton Displacement Damage
  - Proton SEE
- Texas A&M University Cyclotron
  - 55 MeV/amu Argon and Neon
  - LET range from approximately 3 through 15 MeV-cm^2/mg

TID/DDD Data

Pentium III DEVICE UNDER TEST (DUT) TABLE

Device	Rated Speed	Test Condition	Source	Exposure Levels (krads)
P3	800 MHz	Biased	Co-60	*511
P3	933 MHz	Biased	Co-60	573
P3	550 MHz	Unbiased	Co-60	336
P3	650 MHz	Unbiased	Co-60	336
P3	650 MHz	Unbiased	Co-60	3700
P3	700 MHz	Unbiased	Co-60	336
P3	850 MHz	Unbiased	Co-60	697
P3	933 MHz	Unbiased	Co-60	2100

* Indicates part functionally failed

Dose Rate Biased - 2 rad(Si)/min

Dose Rate Unbiased - 7 rad(Si)/min

- Intel P3 parts were exposed to the total dose environment at the GSFC Radiation Effects Facility (Cobalt-60). The results of this testing are summarized in the above table. It should be noted that the one DUT rated at 550 MHz is 0.25 μm technology, while all other DUTs tested are 0.18 μm technology.
- Several P3 devices, in an unbiased condition, have been exposed to various doses, one in excess of 3700 krads(Si). They have shown little sign of degradation in either supply currents or timing and functionality testing. Biased testing of one Pentium III DUT did functionally fail after exposure to an approximate dose of 511 krads(Si). A replacement part was tested under bias, and had exceeded a dose of 573 krads(Si) when it was removed from testing.

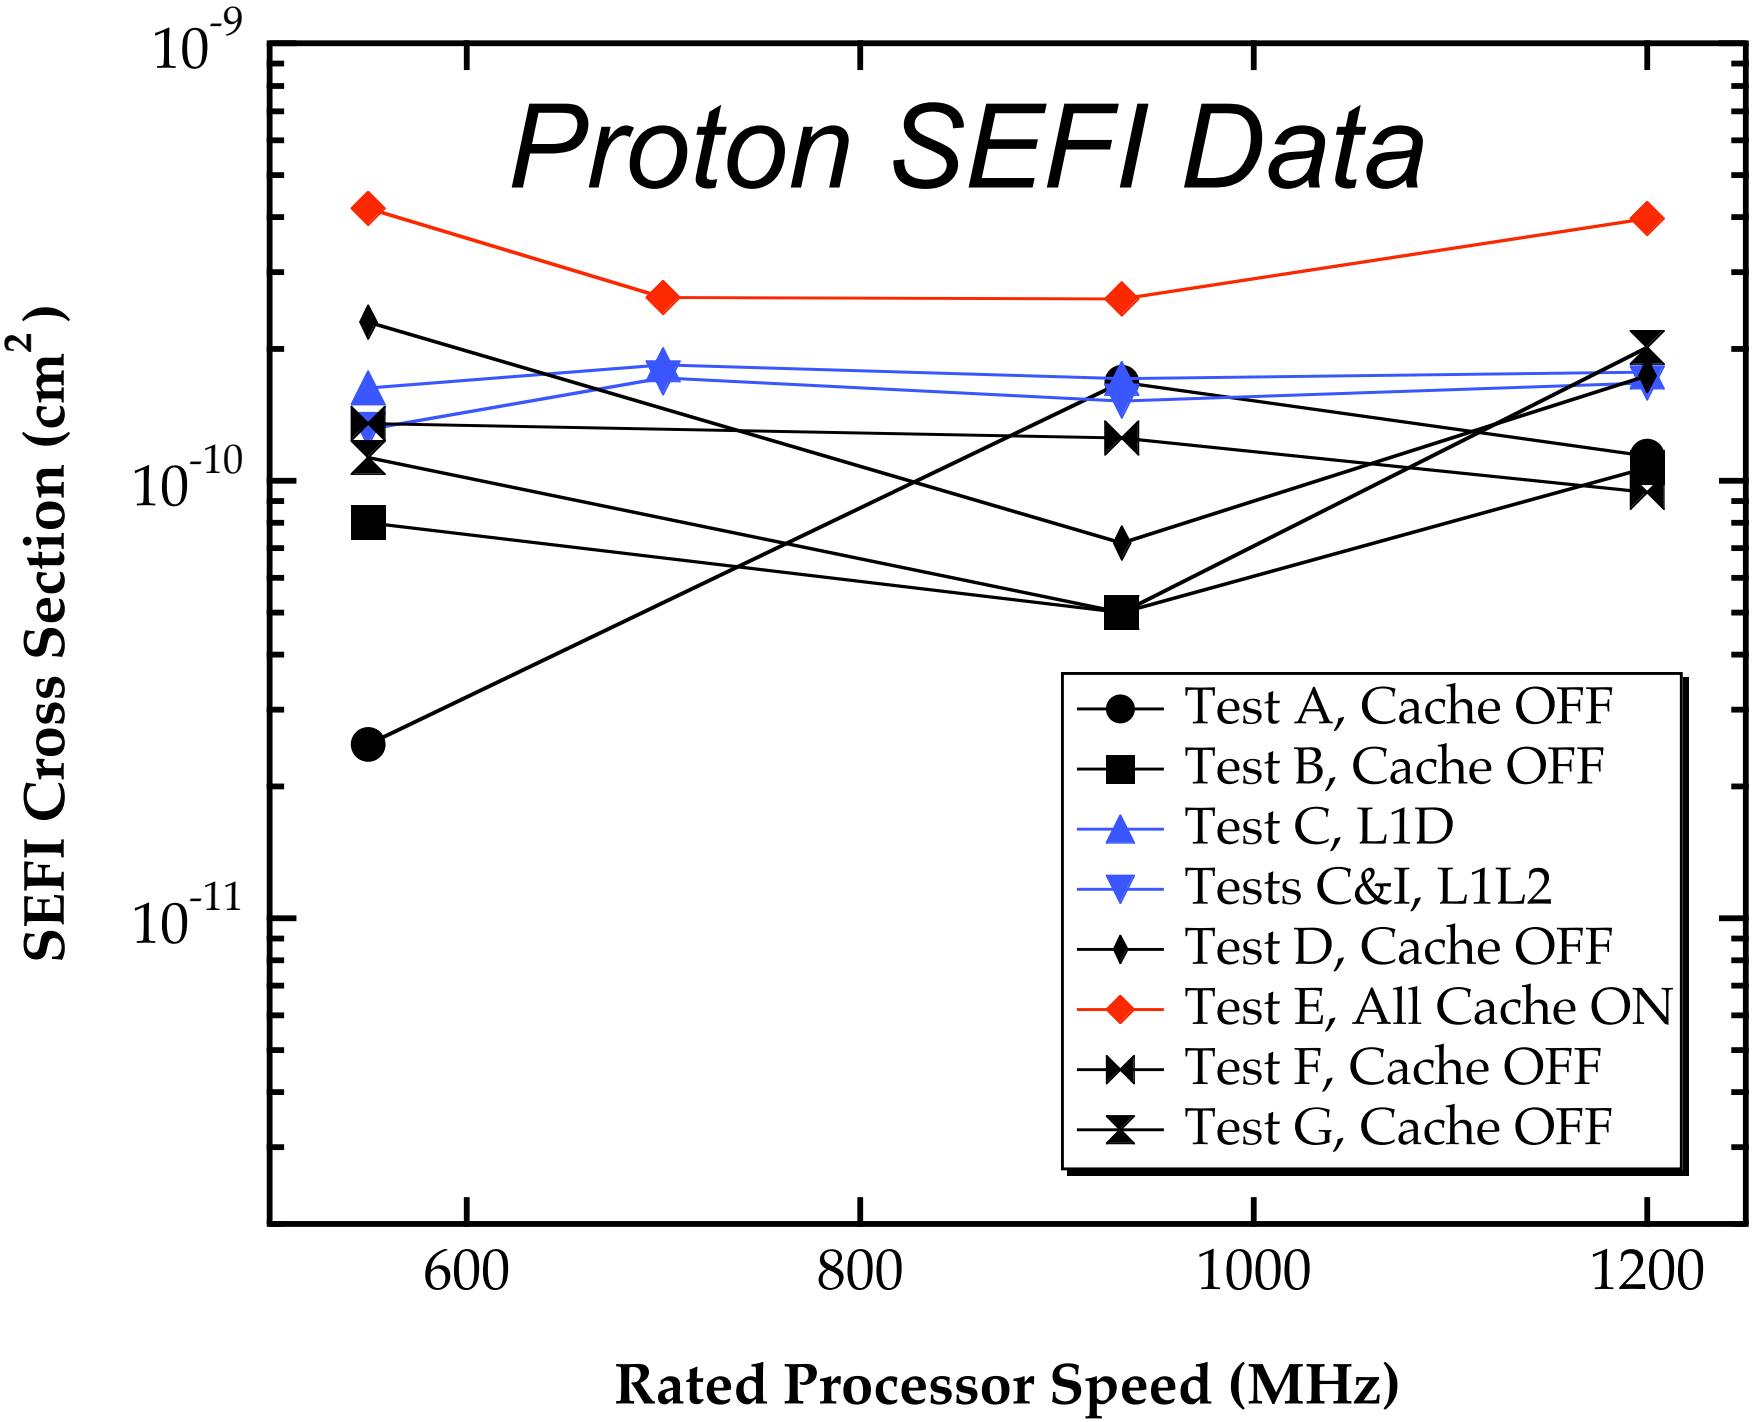
Proton SEFI Data
SEFI Cross Section (cm^2)
10^{-9}
10^{-10}
10^{-11}
600
800
1000
1200
Rated Processor Speed (MHz)
Test A, Cache OFF
Test B, Cache OFF
Test C, L1D
Tests C&I, L1L2
Test D, Cache OFF
Test E, All Cache ON
Test F, Cache OFF
Test G, Cache OFF

- The above figure shows the Pentium III (P3) SEFI cross sections as a function of the DUT rated speed for all test software and conditions. For test software being executed, there is little, if any, dependence as all data lies within approximately one order of magnitude. However, the curve that is the highest (Test E) is also the case where the L1 instruction cache is enabled. It could be argued that the difference observed is real and that the L1 instruction cache being enabled will lead to more SEFIs, but the data does not show enough difference to say it is statistically significant. Also, no speed dependence is observed.

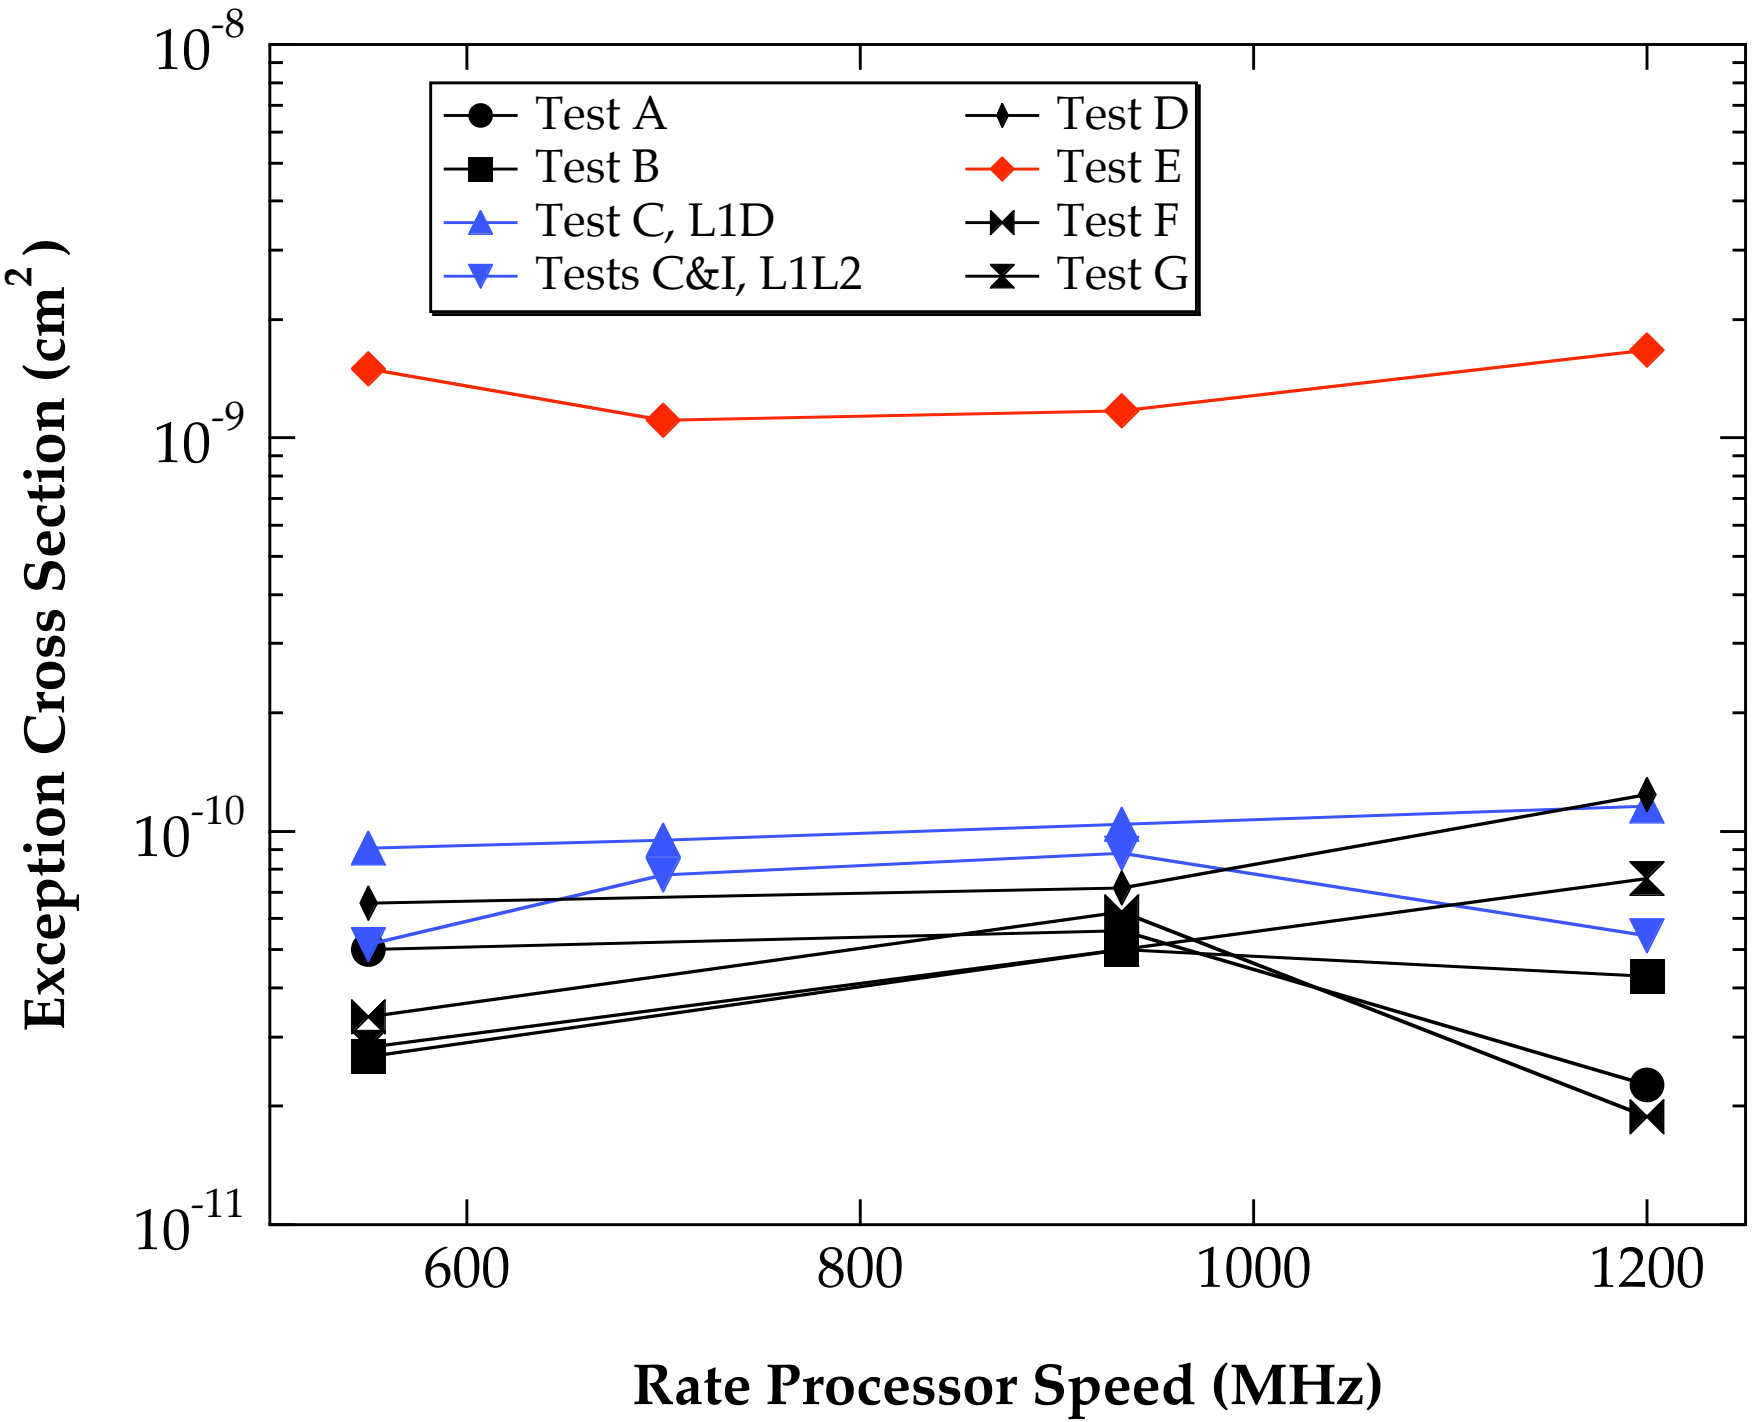
Test A
Test B
Test C, L1D
Tests C&I, L1L2
Test D
Test E
Test F
Test G
Exception Cross Section (cm^2)
10^{-8}
10^{-9}
10^{-10}
10^{-11}
600
800
1000
1200
Rate Processor Speed (MHz)

- The exception cross section (i.e., those events that if not handled would lead to a SEFI) is shown above as a function of the DUT rated speed for all test software and conditions. As with the SEFI data shown in , there is no discernable speed dependence is observed. Unlike the SEFI data, though, there is a significant difference between the "Test E" data point and the remainder of the test software cases. In this "Test E" case, the L1 instruction cache is enabled. With the instruction cache enabled, one would expect that the processor would be more prone to exceptions as any upset of program code in the instruction cache would lead to bad commands, address locations, etc., leading to an exception.
- What is interesting to note in the SEFI and Exception data sets is the relative relation between the curves. Every exception curve is approximately a factor of 3 to 4 higher than the same SEFI curve. Or stated another way, approximately one exception out of every 3 or 4 exceptions will lead to a SEFI event. This would seem to imply that the events that cause a SEFI event are related to the events that cause exceptions, independent of the cache state. The large difference in exception cross section for the "Test E" case (All cache ON) does lead to a larger SEFI cross section but no larger than any other test or cache condition. It is simply that there are many more exceptions occurring and this will lead to more SEFI events.

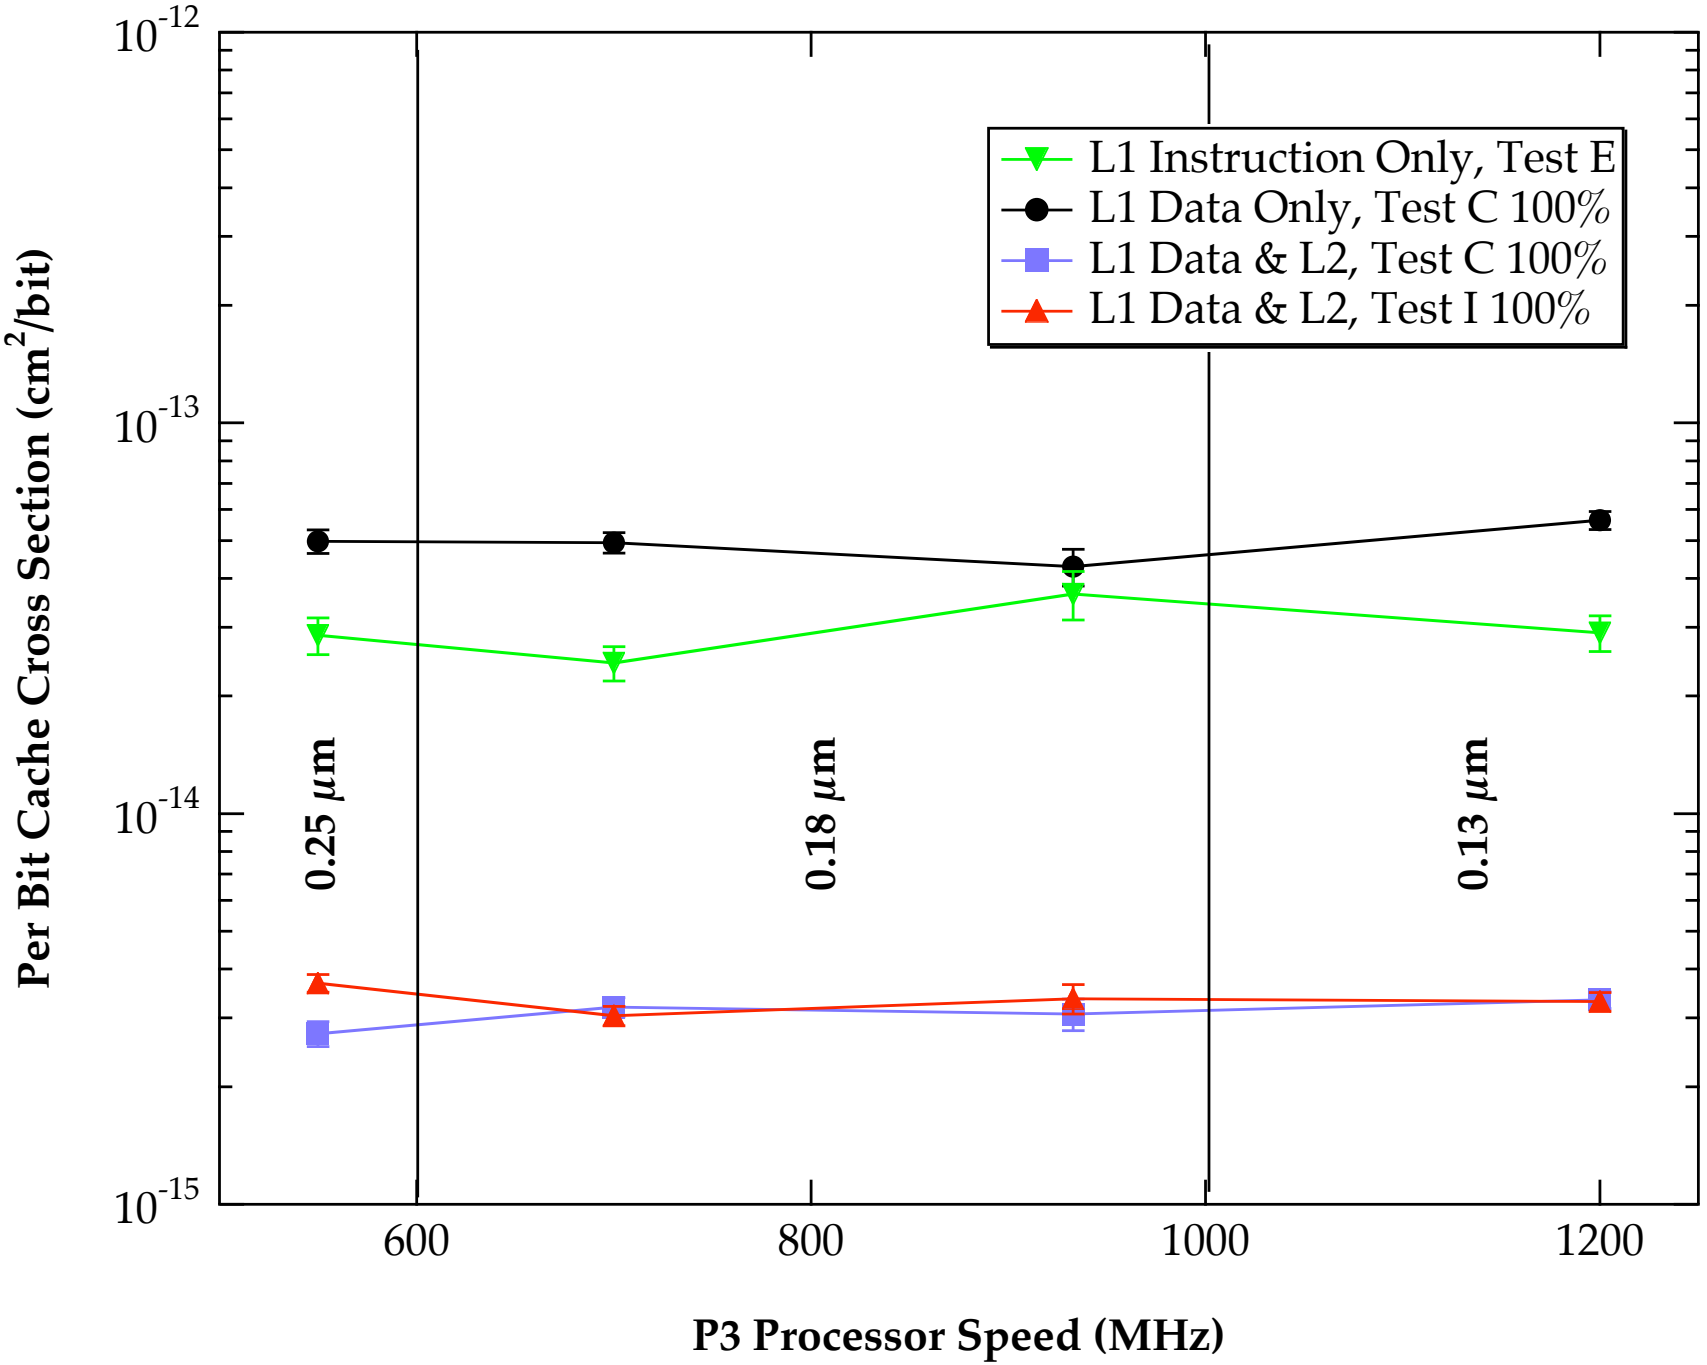
L1 Instruction Only, Test E
L1 Data Only, Test C 100%
L1 Data & L2, Test C 100%
L1 Data & L2, Test I 100%
0.25 µm
0.18 µm
0.13 µm
Per Bit Cache Cross Section (cm²/bit)
P3 Processor Speed (MHz)
10⁻¹²
10⁻¹³
10⁻¹⁴
10⁻¹⁵
600
800
1000
1200

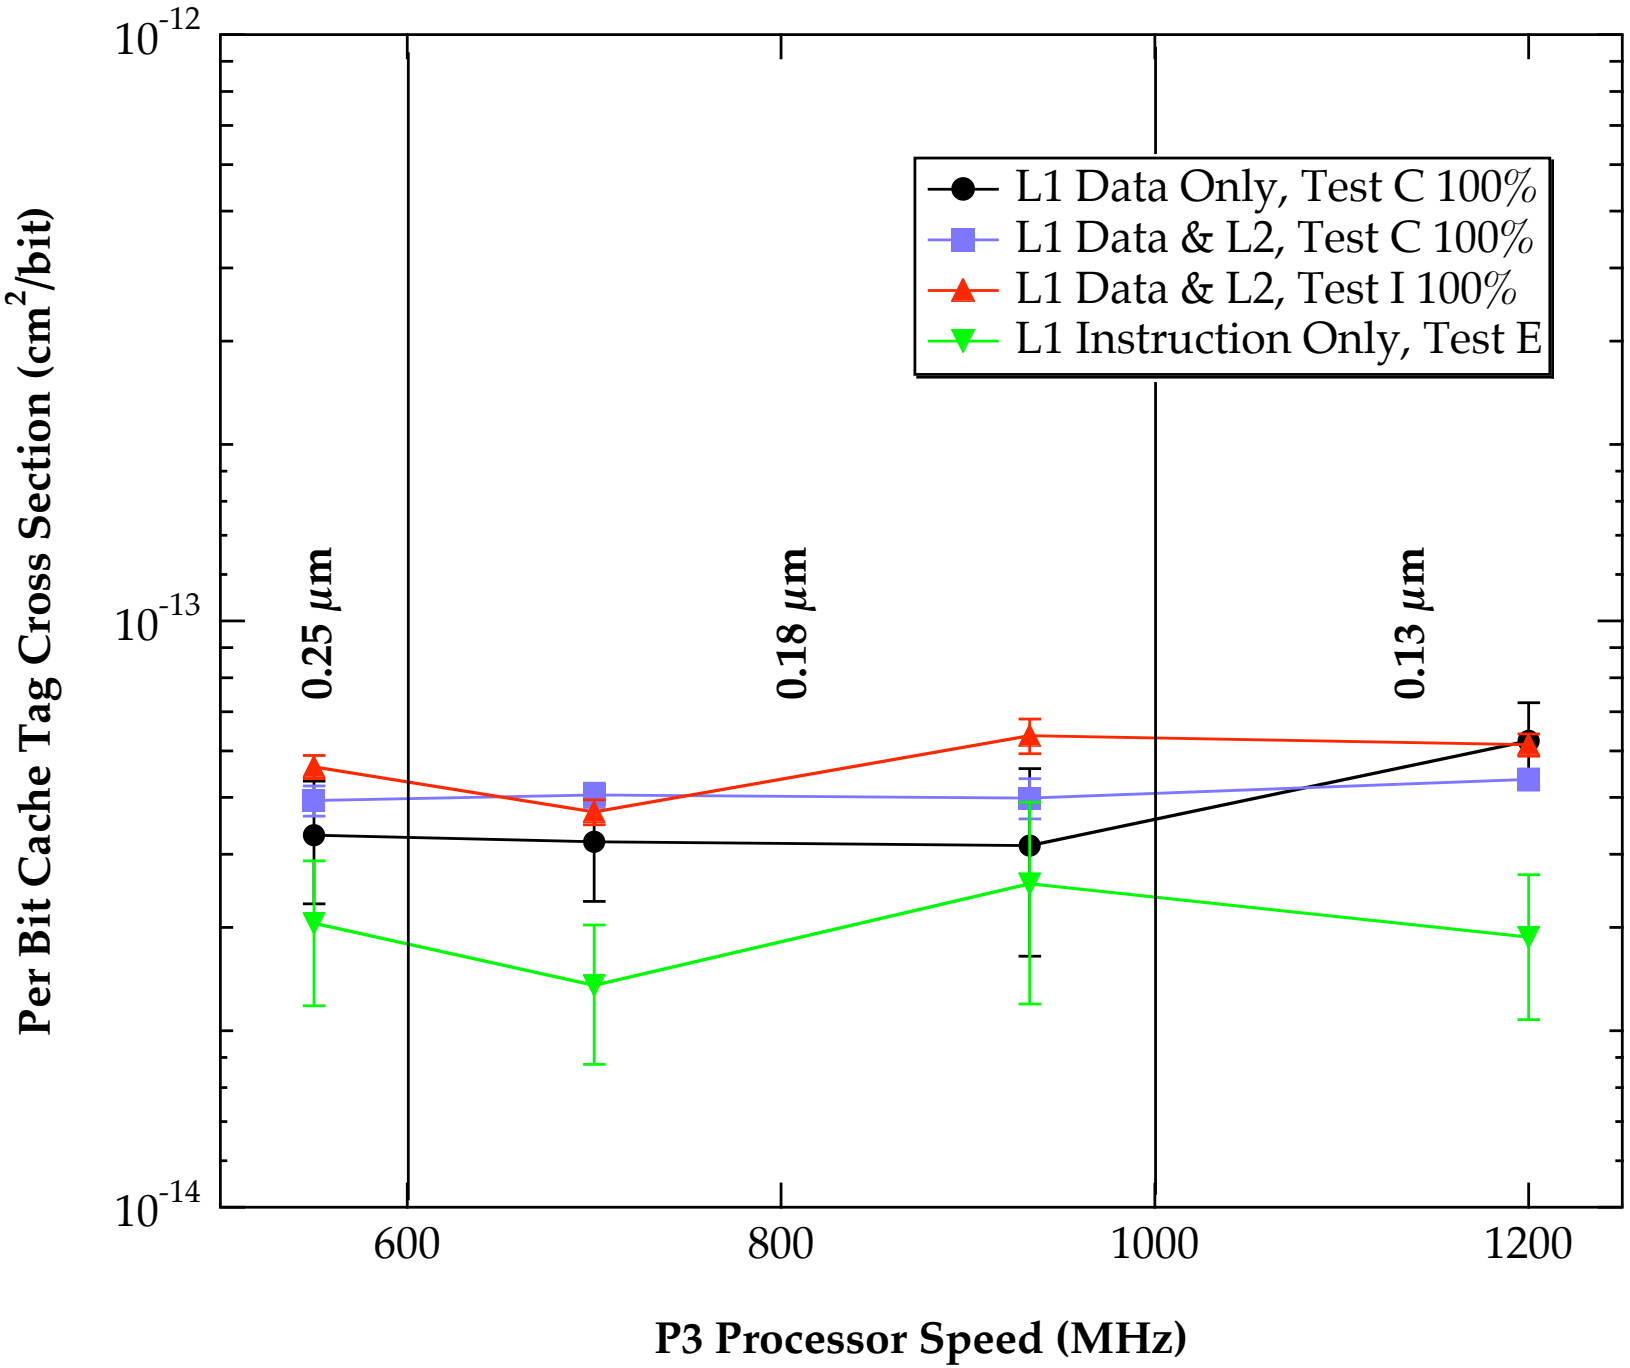
L1 Data Only, Test C 100%
L1 Data & L2, Test C 100%
L1 Data & L2, Test I 100%
L1 Instruction Only, Test E
0.25 µm
0.18 µm
0.13 µm
Per Bit Cache Tag Cross Section (cm²/bit)
P3 Processor Speed (MHz)
10⁻¹²
10⁻¹³
10⁻¹⁴
600
800
1000
1200

- One of the primary goals of this testing was to do a comparison across three technology generations, 0.25, 0.18 and 0.13 microns. To this end, all the data collected was grouped into these categories by using the knowledge that no processor clock speed variation (same processor clocked up or down) and that no significant part-to-part variation has ever been seen. The above graphs show this comparison for the cache bits and tag bits, respectively.
- It is obvious from these three figures that there is no difference in the upset cross sections for the caches (the first place these technology changes would expect to affect).
- The other interesting observation to make is the relative cross sections for all the cache and tag bits independent of which cache they are in/for. Both L1 cache bit and all cache (L1 and L2) tag bits show approximately the same per-bit cross section. The only area that does not show this cross section is the per-bit cross section for the L2 cache bits. This result does make sense as the L1 cache bits and the tag bits are small quantity storage arrays as compared the L2 cache bits and therefore are of the same geometrical structure that is difference from the L2 cache bits. It is also interesting to note that this difference exits across all three technology generations in the same manner, indicating that the same geometrical structures were used for all devices across the generations.

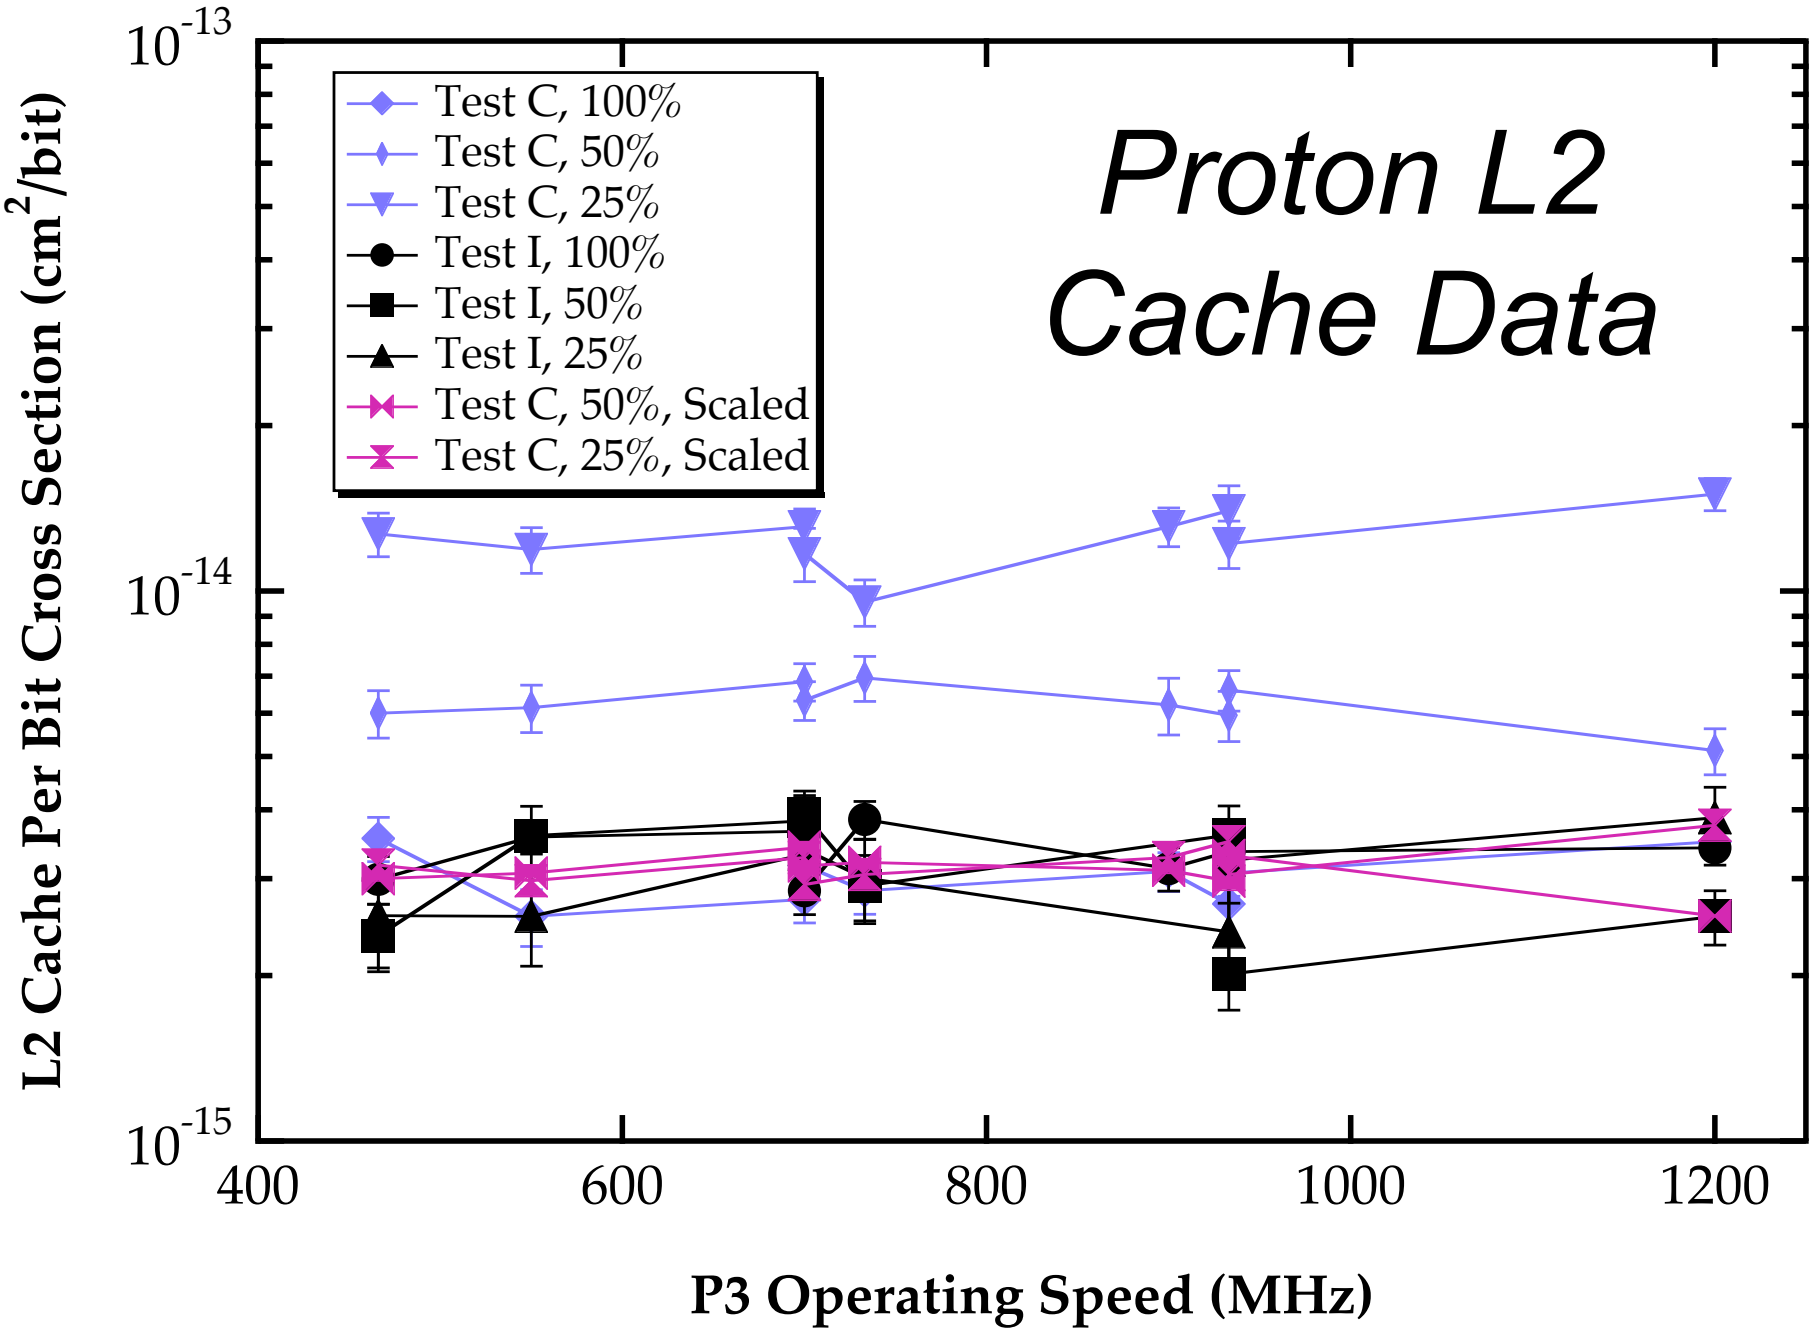
Proton L2
Cache Data
Test C, 100%
Test C, 50%
Test C, 25%
Test I, 100%
Test I, 50%
Test I, 25%
Test C, 50%, Scaled
Test C, 25%, Scaled
L2 Cache Per Bit Cross Section (cm²/bit)
10⁻¹³
10⁻¹⁴
10⁻¹⁵
400
600
800
1000
1200
P3 Operating Speed (MHz)

- The data, shown in above figure, points out a major testing issue that arose during the last heavy ion test. The issue was that when test C was run for the 50 and 25 percent cases, a different per-bit cross section was seen. This same effect was seen in this proton test (the one-sigma error bars do not allow for the possibility of statistics being the cause). However, for this proton test, test I software was written to provide an alternative method of testing the L2 cache that forced a non-sequential addressing scheme. As can be seen, the three percentage cases for test I do not show the same variation as test C. In fact the three test I curves fall on top of the 100% test C curve, giving credence that the 100% test C data gives the correct data.
- It is not clear what is causing this difference between test C and I. There is nothing in any Intel documentation that indicates any special data fetching, storages, etc., whenever data is continually fetched sequentially. Even is there was, that would not explain why the 100% case test C is unaffected by whatever mechanism is in place. Another observation that adds to the confusion is the data for all the per-bit cross sections for tag bits for all caches (L1 Data cache from test C, L2 cache from both test C and I, and L1 Instruction cache from test E) coincide. This would seem to indicate that the correct number of tags are used in each of the difference percentage cases in the sequential and non-sequential addressing tests, as their per-bit cross section are the same. With the correct number of tag bits being used, this definitively implies that the correct number of cache bits are being used.
- This points out a major testing issue when dealing with a “black box” part such as the Pentium III. A lesson-learned from this is that the parameter space for a test such as this, needs to be broadened to allow for unknown effects to arise, be seen, and hopefully understood.

Proton Cache SEU Data

Summary Per Bit Cache Cross Sections						
Cache State	**Single Tag Errors**	**Single TE Error**	**Multiple Tag Errors**	**Multiple TE Error**	**Cache Bit Errors**	**Cache BE Error**
L1 Data	4.62×10^{-14}	2.23×10^{-14}	9.55×10^{-16}	4.78×10^{-16}	4.96×10^{-14}	7.16×10^{-15}
L2	5.72×10^{-14}	6.15×10^{-15}	1.08×10^{-15}	1.08×10^{-15}	9.86×10^{-17}	4.30×10^{-16}
L1 Inst.*	2.77×10^{-14}	4.09×10^{-15}	6.03×10^{-16}	6.03×10^{-16}	3.19×10^{-14}	1.77×10^{-15}

** Calculated values based on percentage of tag versus cache bits.*

- To generate the numbers in the above table all DUTs tested, at all test speeds for the various test conditions were combined into the respective conditions (as all the data shown to date has shown no processor speed or part-to-part variation for the Pentium III devices). This table contains the data for tests C (ignoring the data for 50% and 25% L2 cache cases) and I. It shows the cache bits cross section, the cache tags cross section and a cache tag upset mode cross section that appears to be multiple bit upsets in the tags. Of interesting note is that this mechanism is only present for the L2 cache tags. It is possible that the mechanism exists for the L1 cache tags, as well, but due to the small size, did not manifest itself in this test. The table also contains the one-sigma Poisson error for each of these cross sections, as well.

Proton Other SEU Data

Total Cross Sections for Other Tests					
DUT	**Test**	**# of Upsets**	**Fluence (p/cm^2)**	**Cross Section (cm^2)**	**Cross Section Error**
P3	A	5	2.39×10^{11}	3.93×10^{-11}	1.76×10^{-11}
P3	B	0	2.56×10^{11}	$< 9.59 \times 10^{-12}$	9.59×10^{-12}
P3	D	2	3.96×10^{11}	2.36×10^{-11}	1.67×10^{-11}
P3	F	2	3.92×10^{11}	2.02×10^{-11}	1.43×10^{-11}
P3	G	1	2.62×10^{11}	1.05×10^{-11}	1.05×10^{-11}

- The main item to be seen in the above table is that very few events were actually observed for all the other test that were run. The cross sections calculated in this table are total cross section since, for all but test A, determining the number of effective bits is difficult. For test A, with five 32-bit registers being tested, a per-bit cross section is approximately 2×10^{-13} cm^2/bit, slightly higher than the cache or tag cross sections, but there is quite a bit of uncertainty in that number due to the small number of events. While it is not satisfying to have so few upsets and large variations in cross section, it does imply that the sensitivity of upsets in these areas of the processor is small and of little concern for space operations.

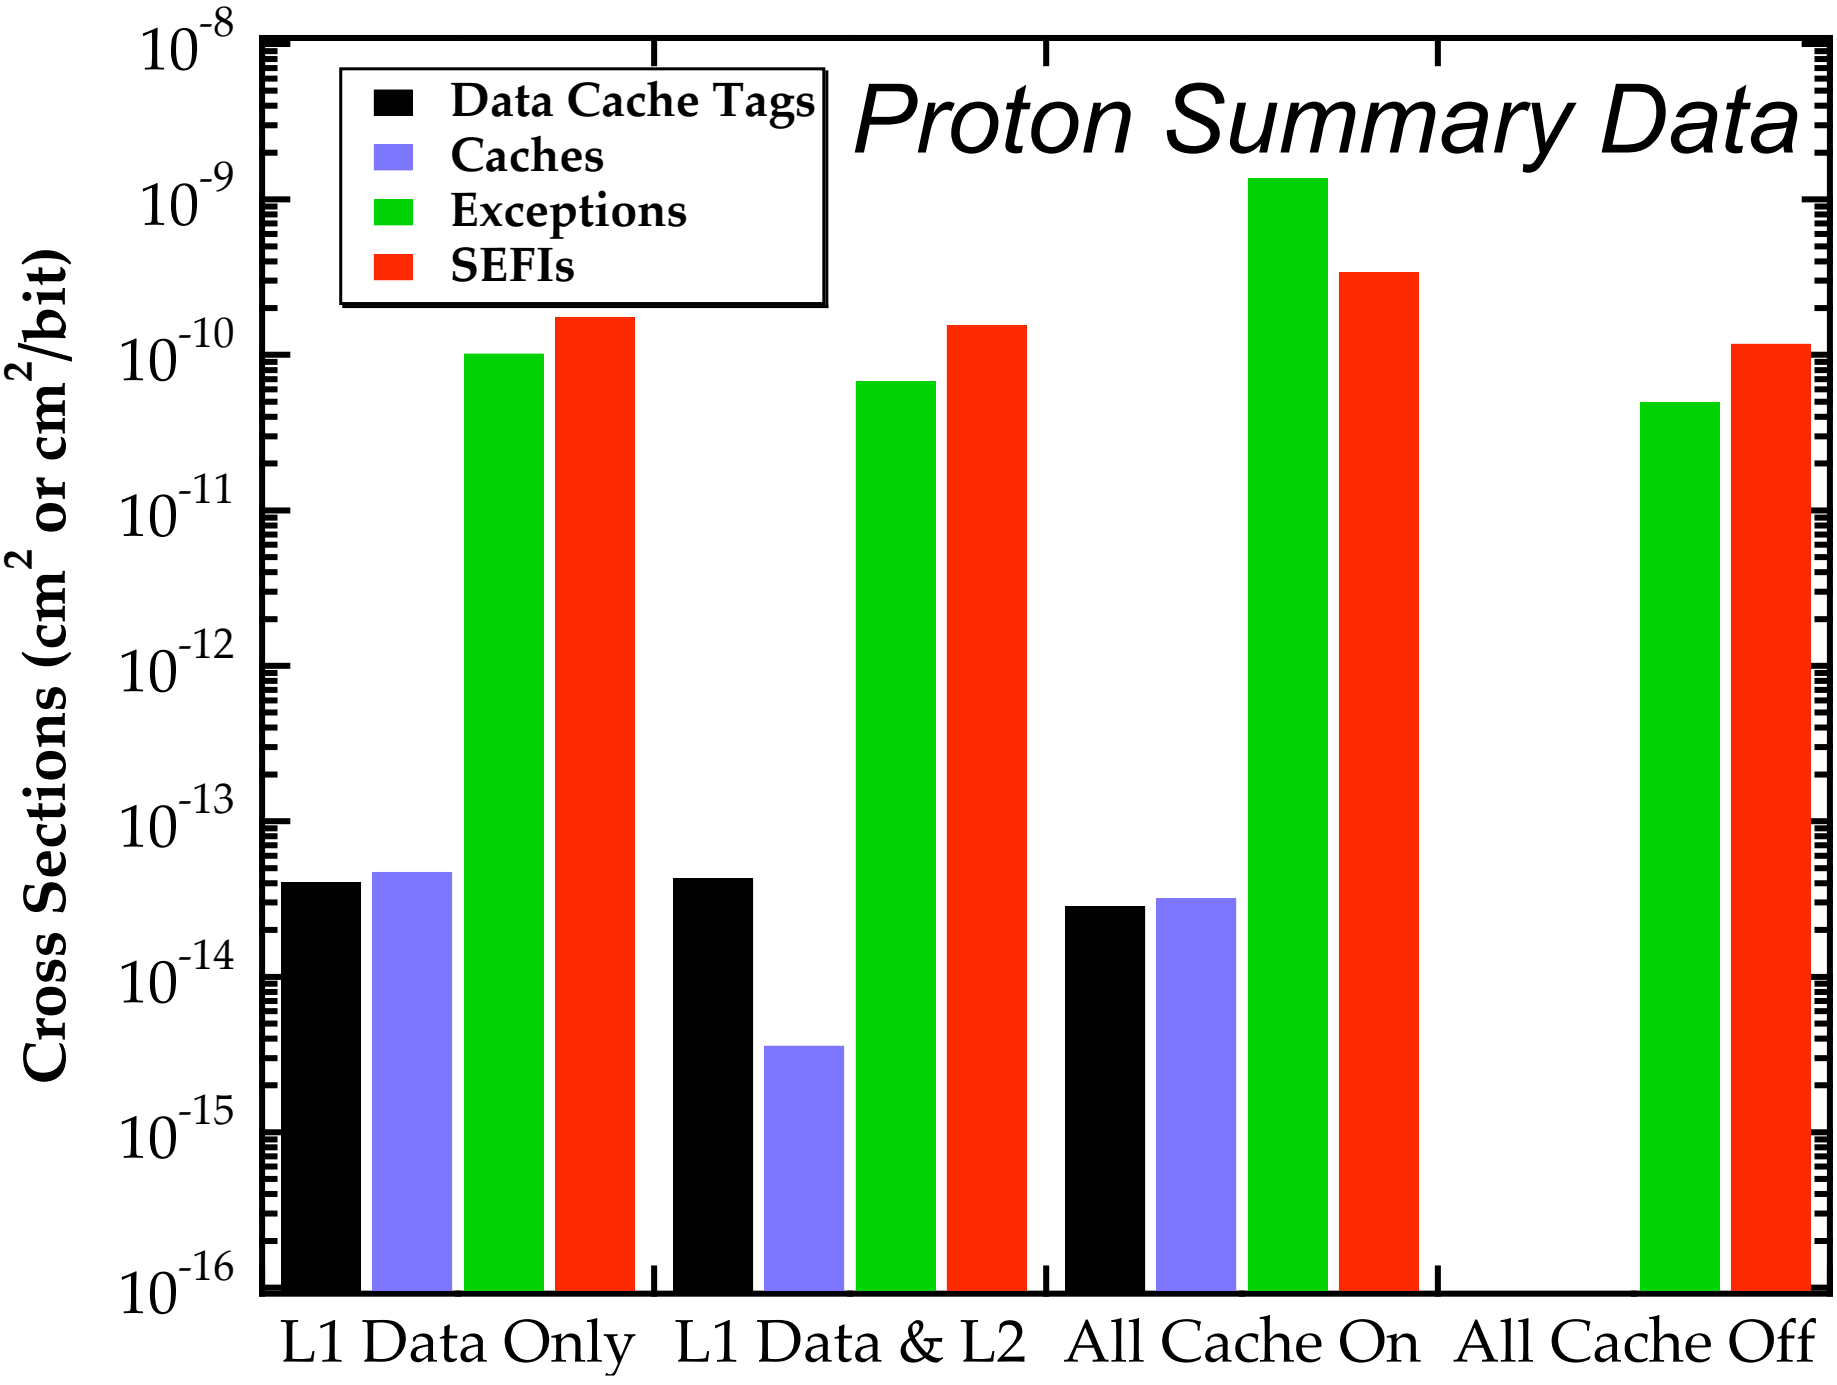
Proton Summary Data
Data Cache Tags
Caches
Exceptions
SEFIs
Cross Sections (cm2 or cm2/bit)
10-8
10-9
10-10
10-11
10-12
10-13
10-14
10-15
10-16
L1 Data Only
L1 Data & L2
All Cache On
All Cache Off

- The above figure shows a summary plot of the data taken for this testing as a function of the cache state. It is very obvious here that the exception rate is dominant for the case using the instruction cache but not for the other cases. This lends a strong impression that exceptions, or exception type events, are not leading to SEFI events. This impression is strengthened by the fact that the SEFI does not increase dramatically for the all caches case, even when the exception rate does.
- While it is not shown this way, if the instruction cache data were presented as a total cross section rather than a per-bit cross section, it's value would be approximately twice that of the exception cross section. Referring back to test E, approximately two-thirds of the instruction cache is tested and one-third is used for instruction space, or twice as much cache is tested as cache space that could lead to an exception. This implies that approximately every proton that produces an event in the instruction cache will lead to an exception. This makes sense, as the test conditions for test E would have to be considered heavy utilization of the instruction cache. For software that did not rely as heavily on the instruction cache, this ration would most likely be much less. The main point to take is the importance of exception handling when the instruction cache is utilized, since if these exceptions were not handled, the processor would halt, leading effectively to a SEFI event.

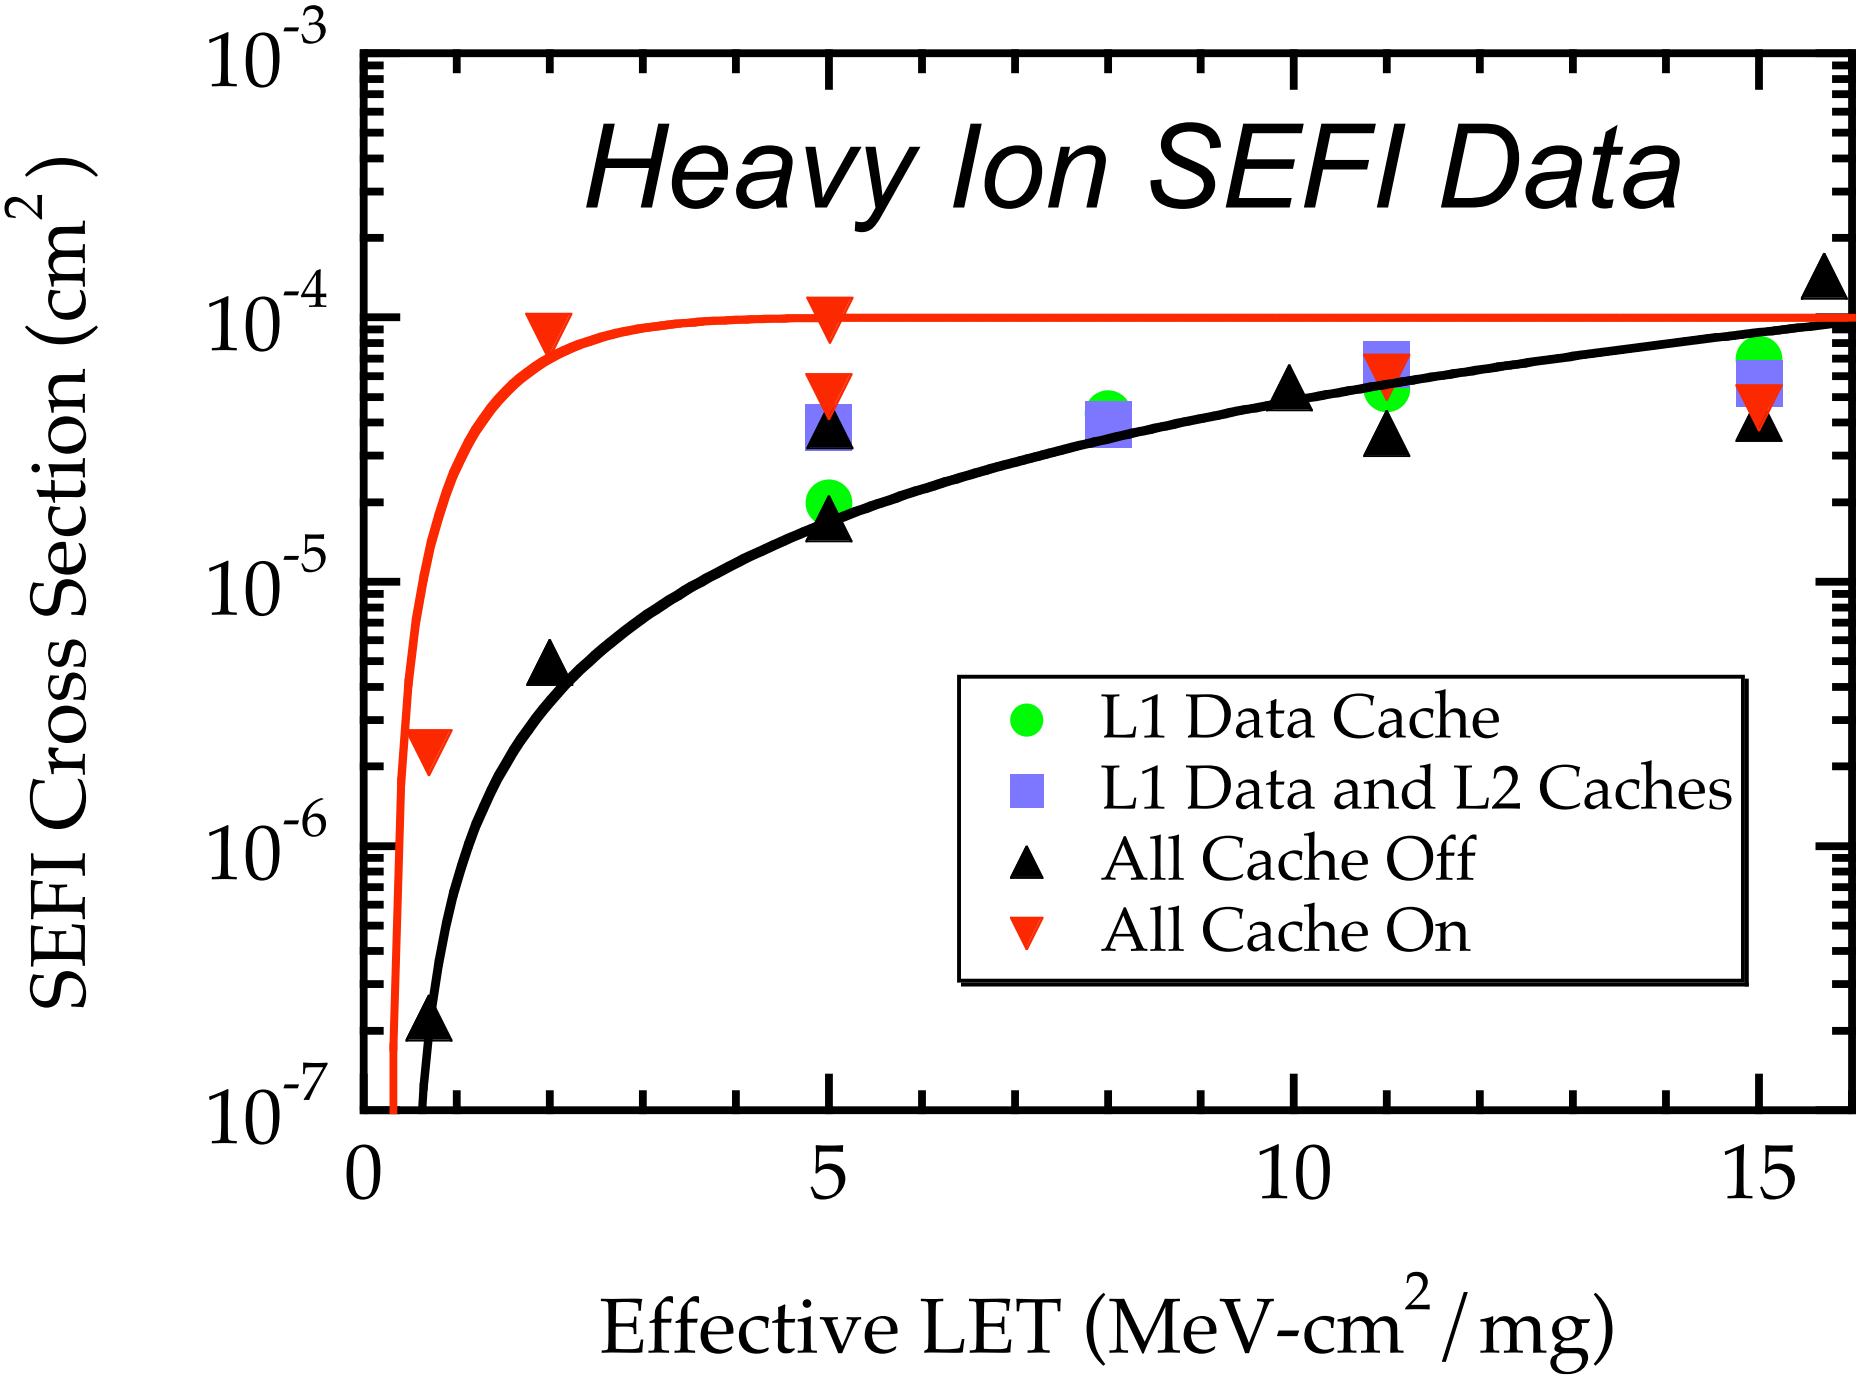
Heavy Ion SEFI Data
SEFI Cross Section (cm^2)
10^{-3}
10^{-4}
10^{-5}
10^{-6}
10^{-7}
0
5
10
15
Effective LET (MeV-cm^2/mg)
L1 Data Cache
L1 Data and L2 Caches
All Cache Off
All Cache On

- The above figure shows the SEFI cross section for all the parts tested as a function of the effective LET. Shown are four data sets representing the four cache conditions that were tested: All cache on, All cache off, L1 Data on only, and L1 and L2 Data cache on only. The first observation to make is that to first order the saturation cross sections for all cache conditions are in the same order of magnitude. While the “All Cache Off” case is consistently below the other cases, it is not significantly below, except possibly in the knee regions of the cross section curves. This data seems to indicate that the primary mechanism that leads to a SEFI event is not generated with an error event in the cache. The cache may play a role in a second order effect near threshold.

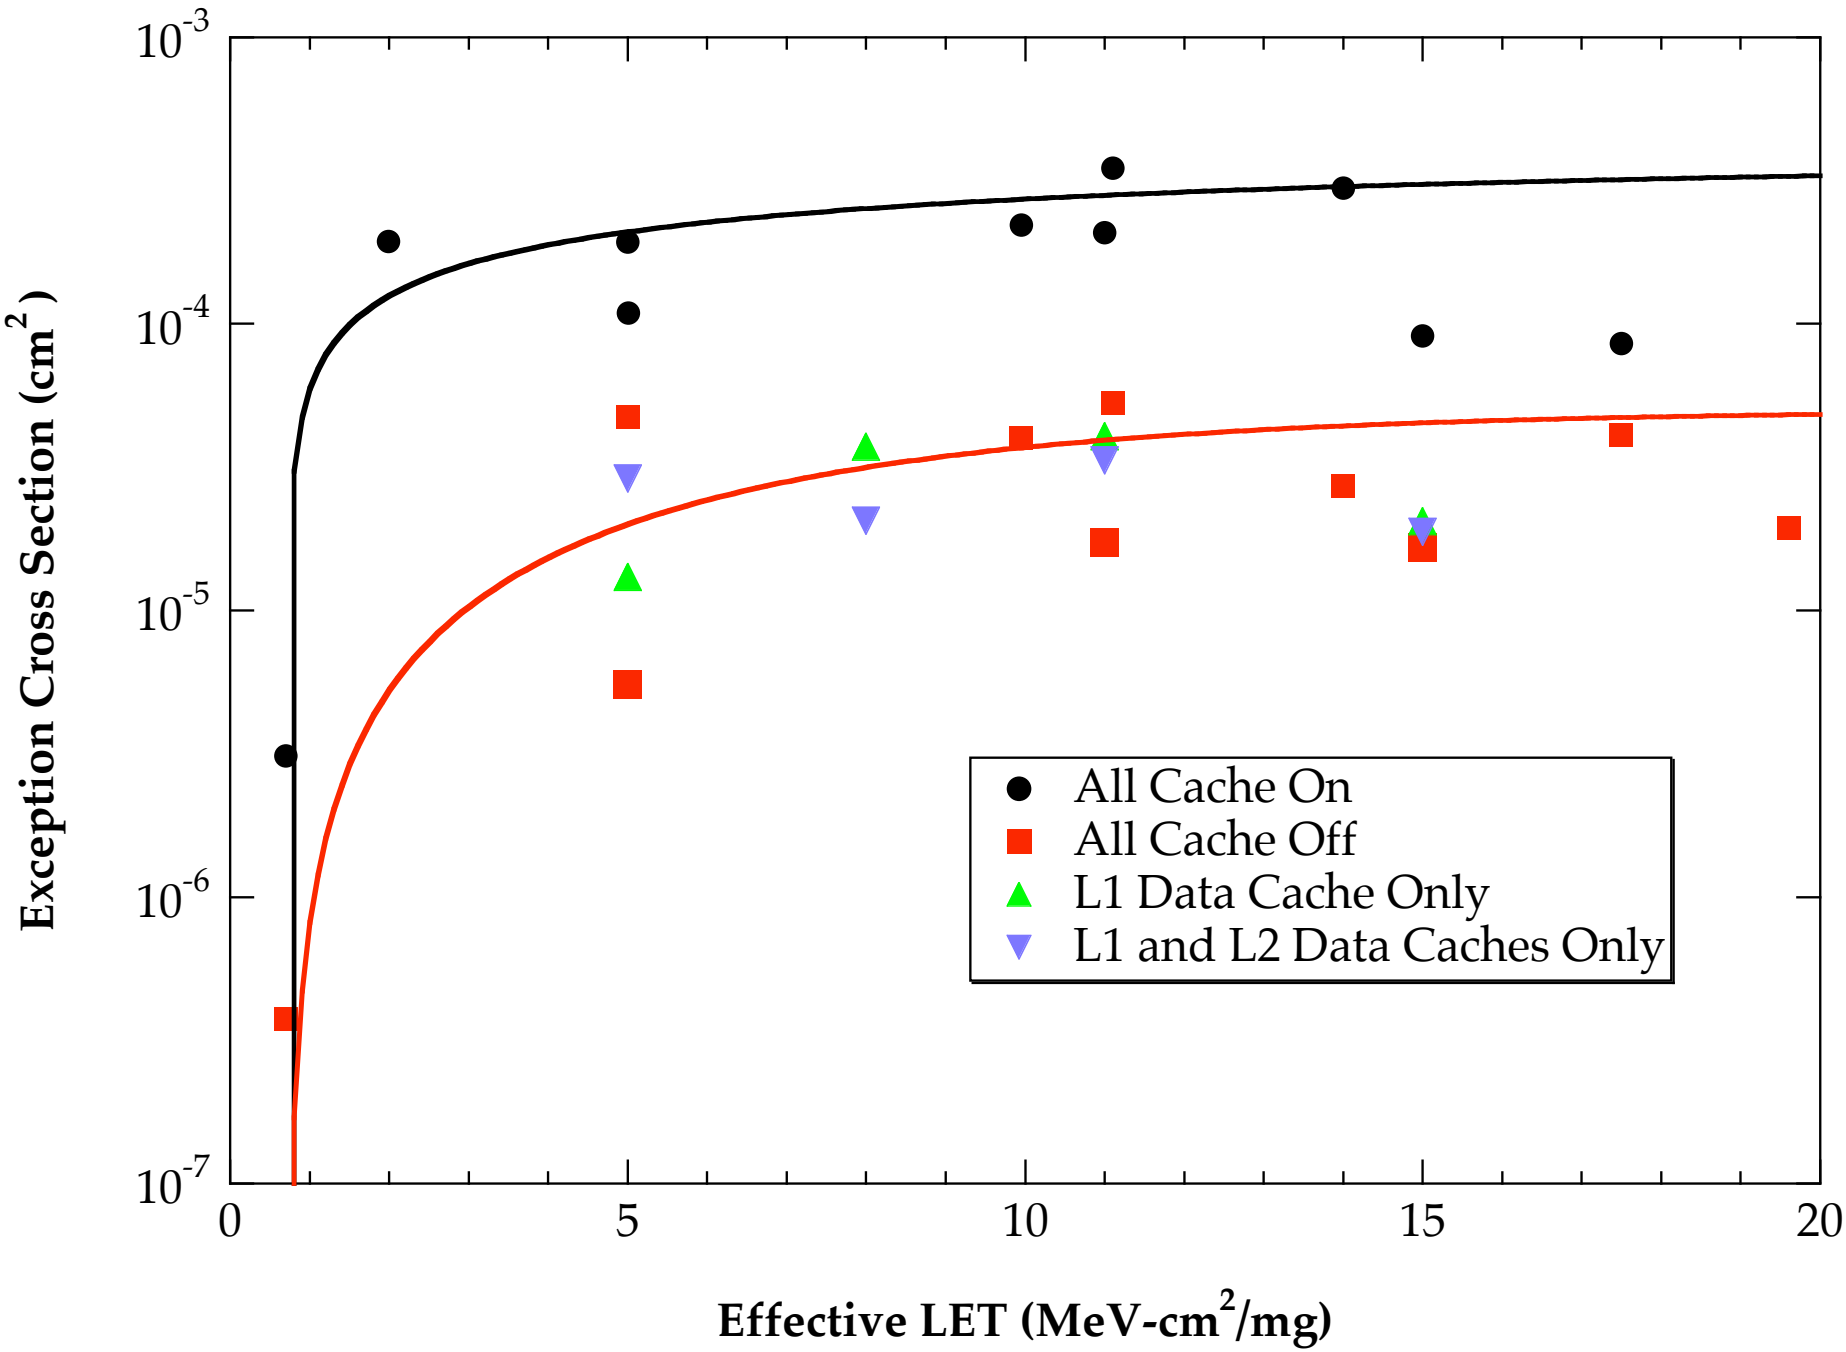
Exception Cross Section (cm^2)
10^{-3}
10^{-4}
10^{-5}
10^{-6}
10^{-7}
0
5
10
15
20
Effective LET (MeV-cm^2/mg)
All Cache On
All Cache Off
L1 Data Cache Only
L1 and L2 Data Caches Only

- The exception cross-section (i.e., those events, if not handled, would lead to a SEFI) is shown above as a function of LET and cache state.
- As with the SEFI, the LET threshold is less than 0.7 MeV-cm^2/mg. There does appear, however, to be a more significant roll-off with LET above an LET of 10 to 15.
- Unlike the SEFI cross-section, it can be seen in that there is a strong dependence on the cache state for the exception rate. There is about an order of magnitude difference between the cases of "cache off" and any of the "cache on" cases. There is a small difference when the data caches are enabled, but the primary effect is seen when the instruction cache is enabled.
- Other non-SEFI events are errors during the individual software tests (Tests A-G, see Software Section). Tests A, B, D, F, and G will be considered first as in all these cases very few to no errors were observed. Fluences during these tests were typically in the range of 10^5 to 10^6 ions/cm^2. Therefore, the registers, the floating-point unit registers and combinatorial logic and the MMX unit have very small cross-sections, on the order of a few time 10^{-6} cm^2 or less.

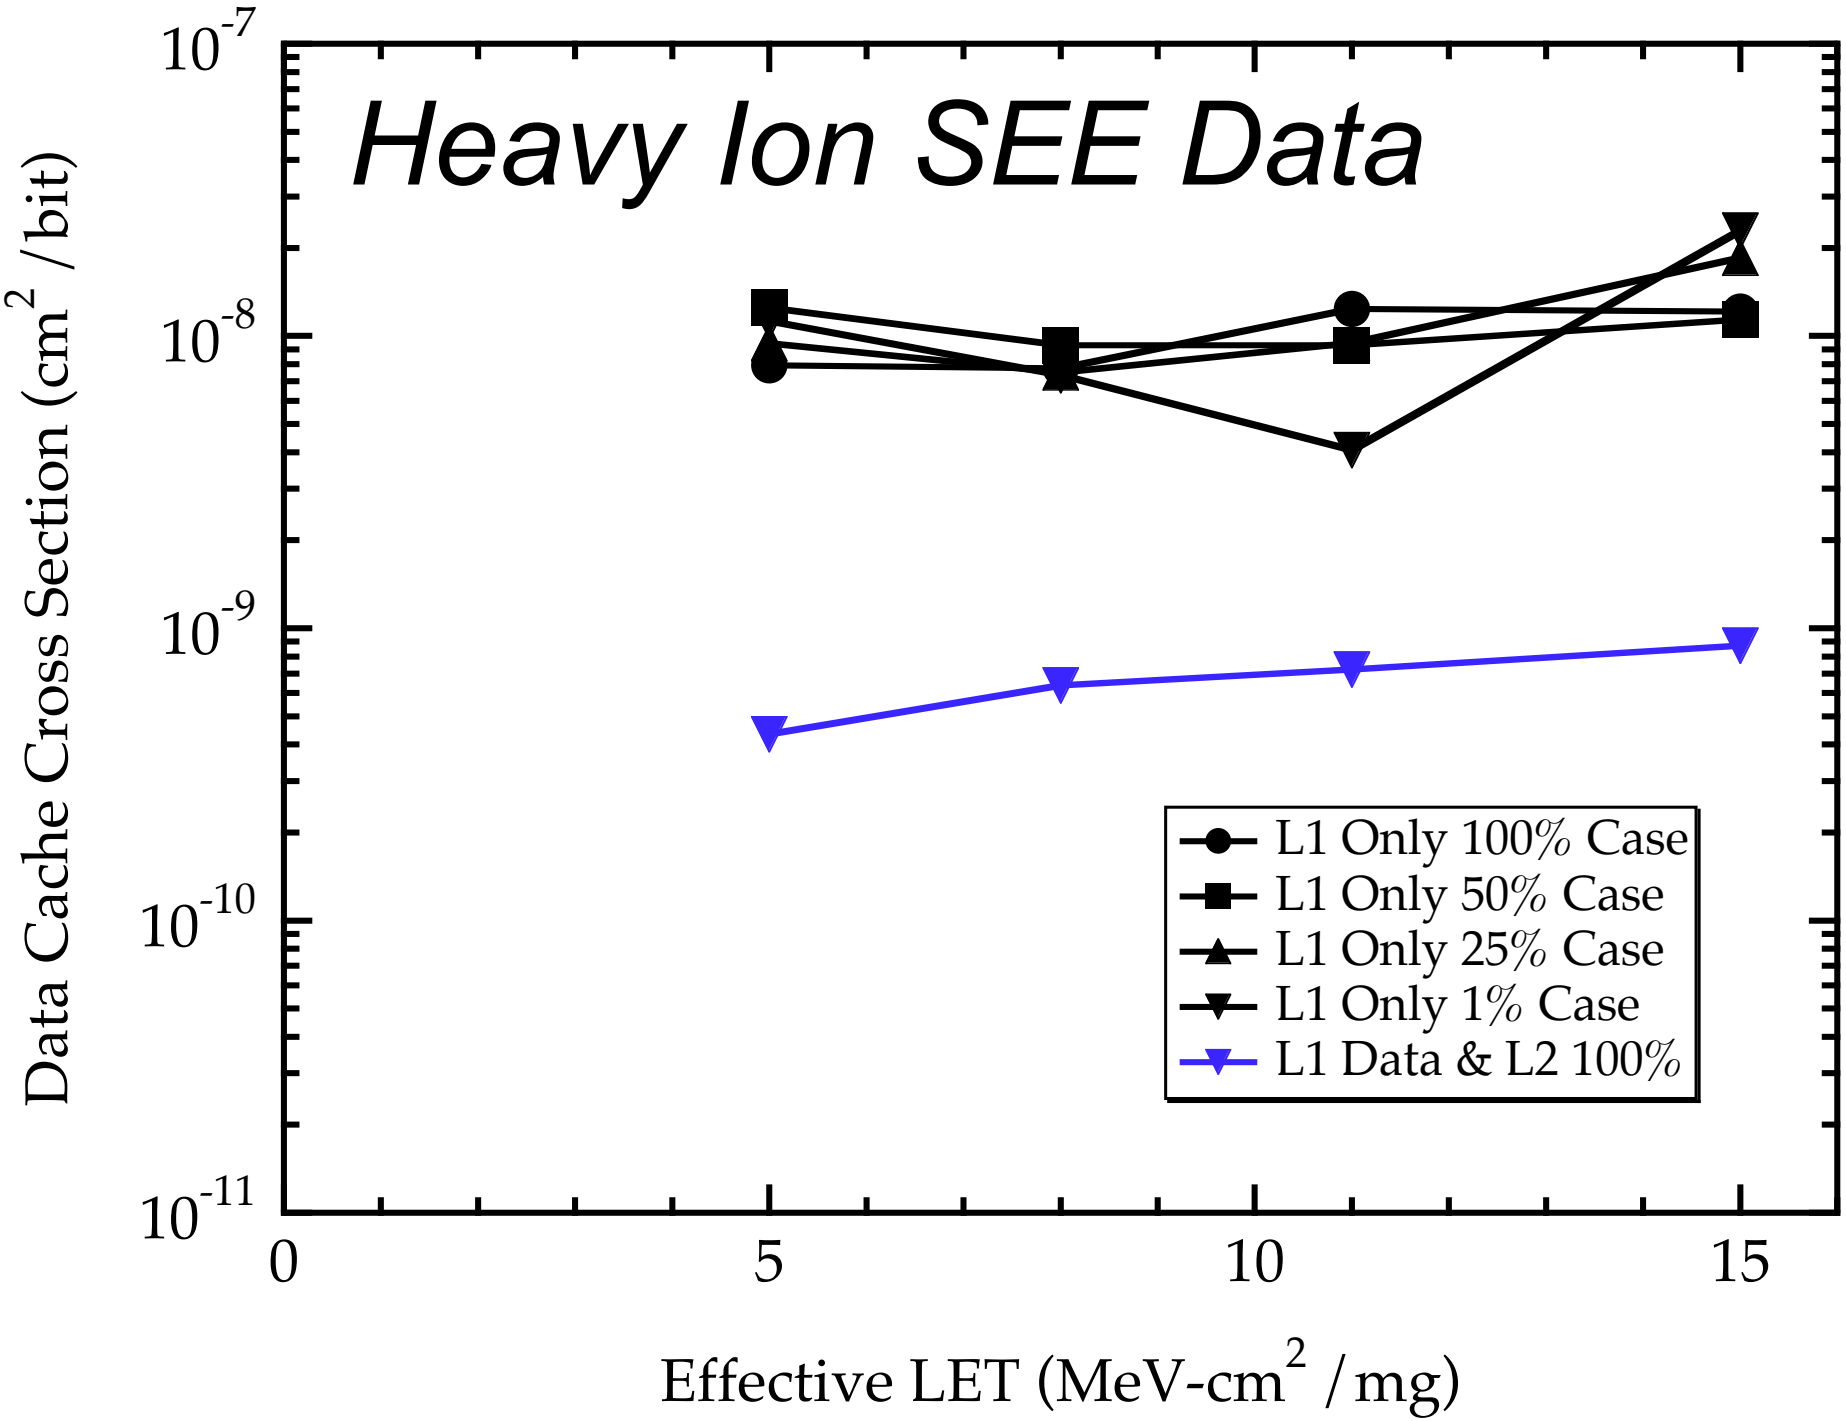
Heavy Ion SEE Data
Data Cache Cross Section (cm2 /bit)
10-7
10-8
10-9
10-10
10-11
0
5
10
15
Effective LET (MeV-cm2 /mg)
L1 Only 100% Case
L1 Only 50% Case
L1 Only 25% Case
L1 Only 1% Case
L1 Data & L2 100%

- The same effect that was pointed out for the proton data (L2 cache percentage issue) was also observed with heavy ions. Proton testing demonstrated that the 100% curve for the L2 cache case contains the correct data. Plotted above is that data as well as the four percentage cases for the L1 data cache cases. As with the proton data, all four L1 data cache case curves fall on top of each other. Also, similar to the proton data, the L1 data cache curves are approximately an order of magnitude larger in cross section that the 100% case for the L2 cache.

Summary

- Extensive data has been collected on the total ionizing dose and single event response of the Intel Pentium III microprocessors.
- The data indicates:
  - high tolerance to TID
  - no susceptibility to SEL from protons or heavy ions to an LET of 15 MeV-cm^2/mg
  - Single event upsets and functional interrupts are present
- Care must be taken in testing parts like the P3 where it must be treated like a “black box”. Our cache testing showed a dramatic difference in test results simply by how the cache is utilized.