

# Guideline for Microprocessor Testing

## 1. Introduction

### A. Overview

This document is intended to be a guideline for radiation tests of microprocessors, which have been the subject of several studies during the last 20 years [1-14]. The main emphasis is on single-event upset testing, first because microprocessors are highly sensitive to single-event upset effects; and second, because there are many technical challenges in performing such tests on modern microprocessors. Total dose testing is addressed only briefly, noting that most microprocessors are relatively immune to total dose damage because of the inherent effects of scaling on device design.

The goal of this work is to develop a guideline that is applicable to processors that are potentially useful in space. Thus, the guideline does not consider very high performance processors that are intended for server or high-performance applications where very large amounts of power are tolerated to gain performance because it would be impractical to use such high-power devices in typical space applications.

Although the results are applicable to hardened microprocessors, the main focus is on high-performance *commercial* microprocessors. Those devices are evolving very rapidly because of performance pressure in the high-volume commercial marketplace. Feature sizes of commercial microprocessors are now at the 90 nm node, and processors are available that operate at clock frequencies of several GHz, providing much higher performance compared to hardened processors.

There is considerable interest in evaluating single-event upset effects at high frequency. Some of the initial work on frequency effects and radiation testing will be discussed. However, the document will not make specific recommendations on testing devices at very high frequency because of the difficulties associated with board design and dealing with the very high power dissipation at high frequency (CMOS power dissipation is essentially proportional to frequency). Test fixture difficulties and power dissipation both act as interferences when tests are done at very high frequencies.

The first part of the document discusses the main technical issues that need to be considered for planning, executing and interpreting microprocessor radiation tests. A relatively sophisticated understanding of those points is essential before the details that are relevant for the guideline can be discussed. Specific recommendations and approaches for radiation tests are included at the end of the document, along with unresolved issues.

### B. Early Test Approaches for Microprocessors

Microprocessor testing was first done more than 20 years ago. The 8-bit devices that were available at that time were very elementary compared to the complex devices that are now available, with very simplified instructions and interface requirements. Because of the simplicity of older processors it was possible to develop machine language instructions for radiation testing, using fully custom hardware. Early microprocessors had 48 to 64 pins, compared to more than 400 for modern microprocessors, and did not contain on-board cache. Three basic approaches were used in the earlier studies:

- The “Golden Chip” approach, which implements parallel, fully synchronous operation of two microprocessors. Only one processor is placed in the radiation source while the other was not exposed. Errors are detected with byte-level comparison of the two processors, running a machine-language program.
- Dedicated machine-language programs that operated specific regions of the processor with individual bit-level visibility. For this type of test, the output of the processor was

compared with expected results using logical comparisons. An automated test system or machine-language emulator can be used to generate the bit patterns. When an error occurs the specific bit pattern and clock cycle are recorded and used to determine which internal functions had likely contributed to the error.

- Development board tests, using an extension cable to connect the processor under test. This method eliminates hardware development costs, but provides more restricted error visibility as well as limited control of I/O and timing.

Table 1 provides an overview of these test approaches. Note the differences in error resolution of the three test methods. Although it was not discussed above, a specific sequence of instructions must be selected and implemented in an operational sequence before any of the three approaches can be used. Results for the earlier processors showed that the SEU response was dominated by register upsets (including upsets in the program counter) [1]. Incorporating additional instructions had little effect on the results.

Table 1. Three Approaches Used in Older Radiation Tests of Microprocessors

Test Approach	Diagnostic Approach	Error Resolution	Comments
<b>Dedicated Machine Language Tester</b>	All pins by clock cycle	Individual bit level	Maximum flexibility of logic levels and timing
<b>“Golden Chip”</b>	Address and data bus comparison	Instruction level	Bit-level diagnostics can be added with logic analyzer
<b>Development Board</b>	Instruction groups	Block level	Relies on I/O features of development board

In principle these approaches could also be used for modern processors. However, the extremely complex interface requirements and the difficulty of designing an operational system exact a very high price for custom hardware development, making the first approach impractical. Although the “Golden Chip” approach can still be used, it is difficult to operate high-frequency processors in lockstep. Thus, most tests of complex microprocessors use the development board approach.

## 2. Key Properties of High-Performance Microprocessors

### A. Performance and Processing Evolution

Microprocessors have changed radically during the last 20 years. The earliest devices used 4-bits, with very primitive capability, but quickly evolved to 8 bits. Table 2 summarizes the properties of several types of microprocessors starting with the early 8-bit devices [1-5,15]. The key points are the drastic reduction in feature size, and the development of SOI processors during the last three years. Although not shown in the table, the number of register bits is still relatively small. However, new microprocessors contain large amounts of internal cache memory, increasing the total cross section for upsets if the cache is used.

Until recently most processor development concentrated on increasing clock frequency and adding architectural improvements such as advanced pipelining, out-of-order instruction sequencing, and increasing the size of on-board cache to increase throughput. At the present time there are several distinct branches in processor development because of the extremely high power dissipation that occurs in microprocessors that are intended for maximum clock frequency and throughput. That

branch, driven heavily by performance, is intended for server applications where the high power dissipation can be accounted for in overall system design; such devices are nearly impossible to use in space because of the extreme difficulty of cooling. A second branch of microprocessor design is intended for mainstream desktop computer applications. Those devices can also dissipate relatively large amounts of power, as much as 100 W. Although it is conceivable that such devices could be used in space, the high power dissipation is a major drawback. The third branch of microprocessor design decreases power dissipation to develop intermediate performance levels with power dissipation below 20 W.

These distinctions are important because (1) high-performance processors use complex packages with massive heat sinks that make it nearly impossible to perform radiation tests; and (2) the predictions of the Semiconductor Industry Roadmap are very different for high-performance, desktop and reduced power microprocessors, which can lead to erroneous conclusions about the performance and features of microprocessors.

Table 2. Comparison of Several Types of Microprocessors

Device	Manuf.	Year	Feature Size (μm)	Core Voltage (V)	Comments
Z-80	Zilog	1986	3.0	5.0	8-bit NMOS
8086	Intel	1986	1.5	5.0	16-bit
80386	Intel	1991	0.8		16-bit
68020	Mot.	1992	1.2		
LS64811	LSI	1993	1.2		
80386	Intel	1996	0.6		
PC603E	Mot.	1997	0.4		
Pentium	Intel	1997	0.35		
PC750	MOT/IBM	2000		2.4	
Pentium	Intel	2002			
PC7455	Mot.	2002	0.18	1.6	SOI process
IBM750FX	IBM	2002	0.13	1.3	SOI process
PC7457	Mot.	2003	0.13	1.3	SOI process

### B. Configuration Requirements

Microprocessors are not stand-alone devices. An extensive amount of supporting electronics is required in order to place a microprocessor in a working configuration. Fig. 1 shows a typical block diagram of a contemporary microprocessor. In this example the L1 cache is contained on-chip, with a direct chip interface to external cache (L2). Such processors typically contain 350 to 450 pins. The processor is designed to interface with a special bridge chip. Random access and non-volatile memory interfaces are done through the bridge chip, along with various input/output functions. A 32- or 64-bit interface is needed between the bridge chip and processor.

Designing a board to allow the processor to operate is difficult, particularly for processors with clock frequencies above 100 MHz. The interface logic levels have been reduced from 5 V for older processors to 2.5 V for more advanced processors. Terminated connections or differential line driver/line receiver pairs must be used at all interfaces. Errors or oversights in board design can lead to sporadic operation that will interfere with radiation tests. In most cases the development boards that

are available from mainstream microprocessor manufacturers have been carefully designed and checked out for operation under worst-case conditions.




Fig. 1 Block diagram of the operational blocks required for a modern microprocessor.

### C. Packaging

Although packaging is usually considered to be of secondary importance for radiation testing, the specific package type used for microprocessors has a large impact on radiation testing because of the difficulty of transporting heavy ions through the package. Modern microprocessors typically use “inverted” packaging with a ball-grid array. A diagram of such an inverted package is shown in Fig. 1. Contacts on the active surface of the die are made with a ceramic substrate. Pins (or direct connection to a circuit board) are attached to the ceramic substrate; they are not shown on the diagram. Because of this inverted structure and the large number of pins, it is not practical to remove the die from this inverted configuration and repackage it so that the active surface is at the top. Consequently for radiation testing with heavy ions it is necessary to maintain the inverted configuration, irradiating the device through the back of the package. The range of ions used for testing must have sufficient range to pass through the surface of the die, or the die thickness must be reduced. Typical die thicknesses are approximately 750  $\mu\text{m}$ . Relatively few ions are available with a range of this magnitude, severely limiting heavy-ion tests. Ion range will be discussed in more detail in a later section.




Fig. 2 Diagram of the inverted package structure typically used for high-performance microprocessors.

Various mechanical methods can be used to reduce the thickness of the back of a microprocessor die, allowing particles with less range to be used for testing. One method uses a high-speed diamond abrasive tool. Fig. 3 shows an example where this was done to reduce the die thickness to approximately 200  $\mu\text{m}$ .



Fig. 3 An example of an advanced microprocessor with an inverted package where the back of the die has been mechanically thinned, leaving approximately 25% of the original die thickness.

### 3. Software Requirements

#### A. Operating Systems

The response of a microprocessor to radiation depends on software as well as hardware. Although it is possible to operate a processor with dedicated machine-language instructions and avoid the need for an operating system, this is generally impractical for the complex processors that are used today, partly because of the need for a bridge chip (or emulated equivalent) to perform most of the I/O and memory interface functions.

Nearly all microprocessor testing is done with some form of operating system. Development boards typically contain very basic operating systems. It is also possible to use more sophisticated operating systems. For example, Hiemstra and Baril conducted proton tests of Pentium processors using a board with the Windows NT operating system [7]. That choice was made because Windows NT was to be used in the manned missions that incorporated the processor. During those tests malfunctions in the operating system frequently occurred, and actually interfered with attempts to examine register upsets, resulting in the “blue screen” that indicates an operating system crash. The hang rate was so high that it was not possible to determine the error rate for registers in those tests. Fig. 4 shows the cross section observed for “hangs” during the tests by Hiemstra and Baril (“hangs” disrupt operation, typically requiring a cold re-boot to restore operation). The “hang” rate for proton tests done with a more primitive operating system for the Power PC processor are 3-4 orders of magnitude lower than for the tests done with Windows NT. These results provide dramatic evidence of the importance of the operating system on single-event upset testing. If a complex operating system is used, it will heavily influence the results and may interfere with attempts to characterize the response of the processor. Thus, very primitive operating systems are preferred for microprocessor

testing. Note however that tests with complex operating systems may be the preferred approach if they are actually used in the application.

The significance of hangs and crashes will be discussed in more detail later. Note that while some of these events may be caused by failures in the operating system, others may occur because a critical part of the processor (such as the instruction register) was affected.




Fig. 4 Cross-section for “hangs” during proton tests of two different types of processors [7]. The large difference in the cross section is almost certainly due to the complex Window NT operating system used for tests of the Intel processors. The Power PC tests used a very primitive operating system provided with a development board which was far less sensitive to processor errors.

## B. Software: Issues and Options

### Register-Level Tests

As discussed earlier, tests of first- and second-generation microprocessors showed that nearly all of the responses to heavy ions were directly related to state changes in registers. Nearly the same results occurred for tests that used a broad range of instructions compared to those with more restricted instruction sets, when the results were examined in the context of errors in internal registers. The transparent operation of 8-bit microprocessors provided direct visibility of the program counter as well as interrupts, facilitating the interpretation of test results. This led to the development of *register-level test software*. This typically consists of a compact program that initially loads specific values into a large number of registers and continually examines the status of the registers during the time that the device is irradiated. The register status test usually takes very little time to execute, and is run periodically at intervals from 1-200 milliseconds, providing a nearly continuous evaluation of register status. A flow chart showing this approach is shown in Fig. 5. If an error is detected during the sequence, the contents of the register and the register status are logged, along with the elapsed time during the irradiation when the error occurred. The error is corrected and the active register test loop continues until the test is stopped. Reasons for stopping the test include (a) reaching a set (pre-planned) fluence; (b) detecting an appropriate number of error counts; and (c) abnormal results, including a “crash” or “hang” that prevents the dedicated register test program sequence from being completed. Note that this test method assumes that the processor works properly nearly all of the time during the test. It will not work effectively unless the error rate is relatively low; and dominated by register errors. The results of this type of test can be reported either in upsets per bit (the preferred approach), or upsets per chip.




Fig. 5 Flow chart showing register-level test software. The test program runs continually during the irradiation, providing nearly continuous visibility of the status of internal registers.

Newer processors use more internal registers than older processors. The cross section of some types of registers may be different because of differences in the device geometry. Consequently it is necessary to refine the register-level test approach to evaluate the response of different types of internal registers (i.e., general-purpose registers, tags and flags). An added complication for newer processors is asymmetry in the sensitivity, resulting in a much larger cross section for upsets in a specific direction (i.e., 1 to 0), compared to upsets that cause the opposite transition to occur. In order to deal with the asymmetry problem the register test must be repeated with different internal register patterns that deliberately test upset symmetry.

### Tests of Internal Cache

Modern processors dedicate a significant part of the chip area to on-board cache memory. Cache memory cells are designed somewhat differently than register cells, typically using more compact geometry. Nearly all processors use 6-T memory cells for cache, although some papers have discussed the possibility of using other technologies, including DRAMs because of the much smaller cell size.

Cache memory needs to be evaluated separately, using dedicated software that is specifically designed to evaluate cache. Some processors are designed to allow internal error correction of the cache memory. This provides an additional degree of freedom for radiation testing. In principle cache tests can be done in the same way as register tests. Typically the number of bits in the cache is much larger than the total number of register bits, providing better counting statistics during an SEU test run because the number of errors is larger. The cross section of the cache bits is typically smaller than that of the registers. Cache bits may also exhibit asymmetric sensitivity to stored 1's and 0's.

### Operational Software Tests

Tests of specific software applications can also be done, although the interpretation of such results is less straightforward compared to register tests. Results of operational software can often only be measured using a “go/no-go” criterion, stopping the test whenever the output of the program differs from the expected result. In such cases a series of runs are made, stopping the beam after an error is detected. The processor program is reloaded after each test, restarting the test until another error occurs. The results can only be reported as a total cross section, not a per-bit cross section as for register or cache tests, and require a sequence of test runs.

Fig. 6 compares tests of the Power 603E for registers, cache (the larger number of bits), and a fast-Fourier transform program. The results are all reported as total chip cross section. The cross section for cache and registers scales with the number of bits. The cross section for the FFT program is much lower, which is the typical result when tests are done with operational software. The reasons for this are related to register usage and visibility. First, even a complex program may use only some of the internal registers; and second, many of the errors that occur in registers will only affect the results if they appear between the time that the register is loaded with information used for the calculation and the time that the step using that information is completed. The latter factor reduces the “visibility” of register errors by factors that are typically between 10 and 100.




Fig. 6 Comparison of the total chip cross section for cache (64Kbits), general purpose registers (5024 bits), and a fast-Fourier transform program. The processor was the Motorola PC603E [5].

The results in Fig. 6 show some key points that affect not only microprocessor testing, but also the way that test results are interpreted in the context of upsets or malfunctions in real applications. In principle it is possible to calculate the upset rate of a specific application program from the more fundamental results from register and related tests, but such calculations require a thorough knowledge of the processor design and architecture.

### Program “Hangs”

The other critical problem that must be dealt with is the incidence of “hangs.” Although “hangs” occurred relatively infrequently in older processors, they occur far more frequently in modern processors that have complex internal architectures. In nearly all cases recovery from a “hang”



condition cannot be done by applying a reset pulse, but requires removal of power and cold restart. Consequently “hangs” are extremely important for modern microprocessors.

The “hang” cross section for Intel and Power PC processors was shown earlier in Fig. 4 for proton tests. In that example the much higher cross section for Intel processors was attributed to crashes in the operating system (Windows NT in that example). Fig. 7 shows specific “hang” cross section results for two highly scaled PowerPC processors that were obtained with a far more primitive operating system [10]. Auxiliary tests (including JTAG) support the conclusion that the “hangs” are due to conditions in the processor, not failure of the operating system.

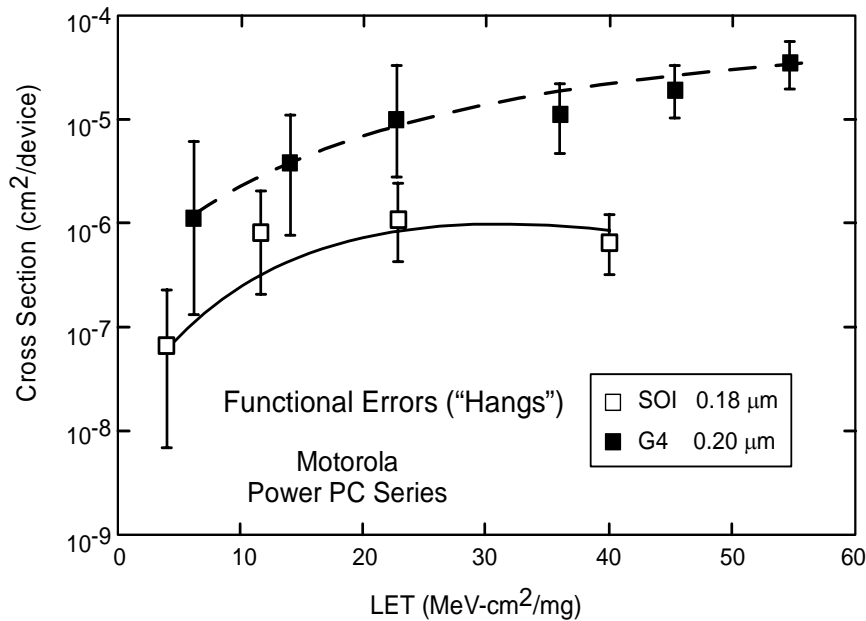


Fig. 7 “Hang” cross section for advanced versions of the PowerPC processor. The version using SOI technology has a lower cross section, but the threshold LET is approximately the same [10].

## 4. Accelerator Beam Requirements

### A. Heavy Ions

The active region of high-performance processors is very thin, typically about 2-3 μm. However, as discussed in Section 2, the inverted package used for microprocessors requires irradiation from the back of the package. The thickness of a typical microprocessor die is 700-800 μm. In order to test such a device, the range of the ion beam must exceed the die thickness. Note that there is significant energy loss by such beams as they travel through such thick regions, requiring correction for energy loss. Fig. 8 shows how the LET of four different types of ions available at the Texas A&M (TEM) cyclotron changes with distance traveled in silicon. For the ions shown, neon is the only ion with sufficient range to get through the normal die thickness. The LET varies rapidly with distance near the end of the range, and thus it is necessary to know the die thickness to within about 2%. Ar, Kr and Kr ions can be used on devices where the back of the substrate has been reduced by mechanical thinning.

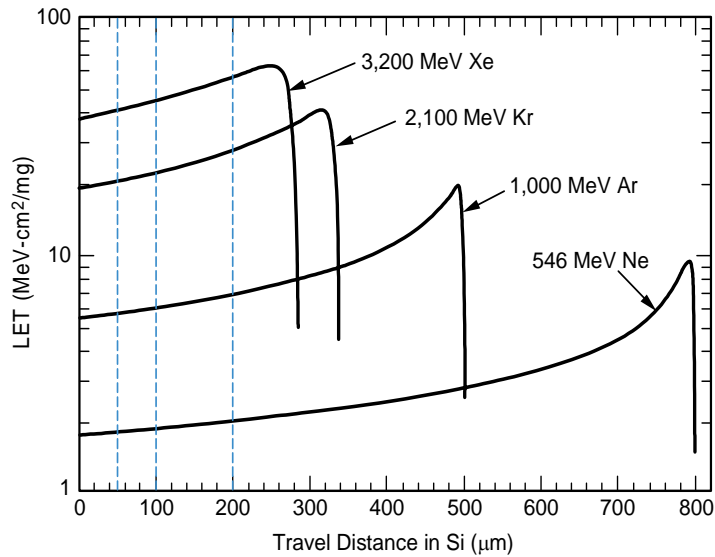


Fig. 8. Change in linear energy transfer with distance for four different ion species available at the Texas A&M cyclotron. When irradiation is done from the back of the die, the LET must be corrected to account for energy loss as the beam traverses the silicon to the (thin) active front surface.

*B. Beams Available at Various Accelerators*

The beams that are available for three commonly used accelerators are shown in Tables 3-6. Tables 3 and 4 show the range of high LET ions at Brookhaven and UC Berkeley. Although ions with lower LET have somewhat greater range, the ions that are available from those facilities have such a limited range that it is nearly impossible to test processors using irradiation from the back of the die because the die would have to be reduced so much in thickness that it would affect the packaging and lead integrity of the “flip-chip” bonding.

Table 3. Range of Ions with LET > 30 MeV-cm<sup>2</sup>/mg at Brookhaven

Ion	Energy (MeV)	LET at Normal Incidence (MeV-cm <sup>2</sup> /mg)	Range at Normal Incidence (µm)
Br	305	36.9	38.7
Ag	345	52.9	34.5
I	370	60.1	34.3
Au	390	84.1	29.1

Table 4. Range of Ions with LET > 30 MeV-cm<sup>2</sup>/mg at UC Berkeley

Ion	Energy (MeV)	LET at Normal Incidence (MeV-cm <sup>2</sup> /mg)	Range at Normal Incidence (µm)
Cu	290	30	45
Kr	380	41	46
Xe	600	63	50
Bi	950	90	50

The Texas A&M cyclotron produces ions with far greater range. Tables 5 and 6 show the energy and range of ions with energies of 25 MeV and 40 MeV per nucleon. The LET values are the LET at the surface. LET increases as the ion loses energy during its transition through the silicon, increasing the LET. The LET must be corrected to account for that energy loss. Fig. 9 shows how the LET changes with distance for the three types of ions with energies of 40-MeV per nucleon.

Table 5. Range of 25 MeV per Nucleon Ions at Texas A&M

Ion	Energy (MeV)	LET at Normal Incidence (MeV-cm <sup>2</sup> /mg)	Range at Normal Incidence (μm)
Ne	545	1.7	790
Ar	991	5.4	485
Kr	2081	19.3	332
Xe	3197	37.9	286

Table 6. Range of 40 MeV per Nucleon Ions at Texas A&M

Ion	Energy (MeV)	LET at Normal Incidence (MeV-cm <sup>2</sup> /mg)	Range at Normal Incidence (μm)
Ne	800	1.2	1655
Ar	1598	3.8	1079
Kr	3117	14.2	622




Fig. 9. Variation in LET of 40-MeV per nucleon ions with distance in silicon.

### C. Proton Testing

Protons with energies above 65 MeV have sufficient range to get through the thick substrate used in processors with inverted packages, although it may be necessary to correct for energy loss. The range of protons with energies between 15 and 100 MeV is shown in Table 7. Unless the substrate is thinned, it will be nearly impossible to determine the threshold proton energy for upset because of straggling and uncertainty in device thickness. Thus, although the proton cross section can be determined at high energies without modifying the device, it is necessary to use thinned devices in order to determine the energy threshold, typically less than 30 MeV.

Table 7. Range of High-Energy Protons in Silicon

Energy (MeV)	Range in Silicon ( $\mu\text{m}$ )	Range in Silicon (mils)
15	1,585	62.4
20	2,580	101.6
30	5,820	229.1
50	17,700	696.8
65	24,400	960.6
100	46,600	1,834

Proton testing can be done “in air,” with supporting equipment located close to the device. This makes proton testing inherently more straightforward compared to tests with heavy ions. However, relatively high fluences are required, which may damage the device during a series of test runs. Proton cross sections for highly scaled devices are on the order of  $10^{-13}$  to  $10^{-14}$   $\text{cm}^2$  per bit. As an example, if 5,000 registers are used by the software in a specific test, a fluence of  $5 \times 10^{11}$   $\text{p}/\text{cm}^2$  is required to measure (on average) ten upsets if the cross section is  $10^{-14}$   $\text{cm}^2$  per bit.

## 5. Examples of Test Results for High-Density Microprocessors

### A. Register and Cache Tests

#### Power PC Processors

Test results for several different PowerPC processors are shown in Fig.10, along with test results provided by the manufacturer of a radiation-hardened processor with the PowerPC architecture, the RAD6000. The feature size of the commercial processors is shown in parentheses. Tests of the commercial PowerPC devices were done with a development board, using an elementary operating system. Irradiations were done from the back of the die, using the long-range ions available at Texas A&M.




Fig. 10. Register test results for three different commercial PowerPC processors.

Cache test results for the PowerPC processors are similar, with a slightly lower saturation cross section. The lower cross section is due to the smaller size of the transistors that are used in the cache memory.

Cache test results for SOI processors with two different feature size are shown in Fig. 11. Despite the decrease in feature size, the threshold LET is essentially the same.




Fig. 11. Cache test results for two SOI PowerPC processors with different feature size.

Proton tests of registers are shown in Fig. 12. The proton cross section is about five orders of magnitude lower than the heavy ion cross section (Fig. 10), which is consistent with the lower interaction probability for protons because of the small nuclear cross section. The energy threshold is below 10 MeV.

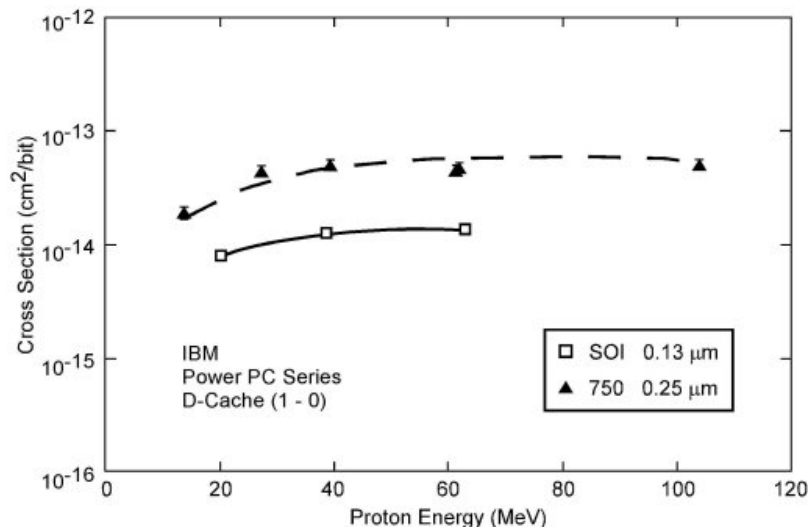


Fig 12. Proton cross section for cache memory in two different PowerPC processors.

### Intel Processors

Earlier work by Hiemstra [7] on proton tests of Intel processors was shown in Fig. 4. Heavy-ion tests of Intel processors were done by NASA Goddard [16]. They tested the following registers:

- 1- CPU registers (ebx, ecx, ebp, and sdi).
- 2- Math processor registers
- 3- MMX unit registers (pxor, por, pmul, pmulh, pads, addps, divps and mulps)

Figure 13 shows an example of their results. They are reported on the basis of total upsets at the chip level rather than “per bit”. The Intel processors are more difficult to test because of higher power dissipation.

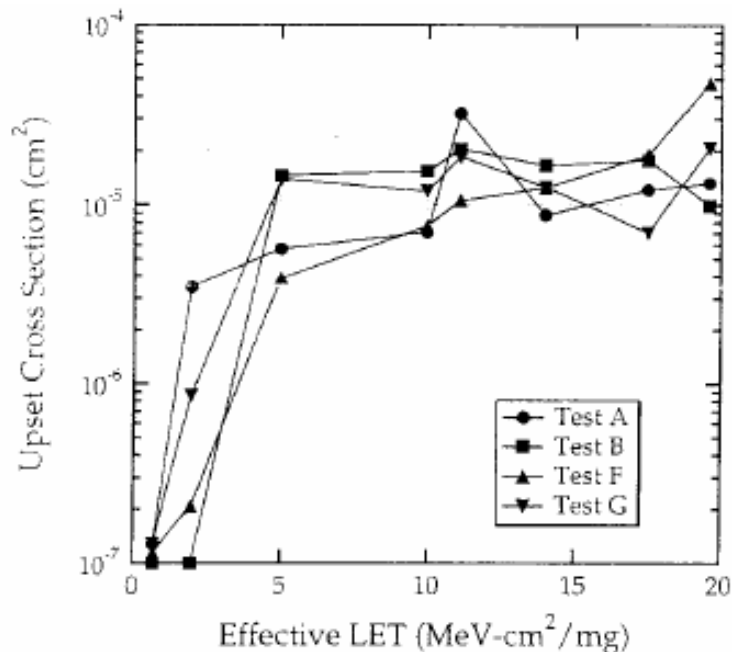


Fig. 13 Pentium III upset cross section as a function of the Effective LET, for various test cases [16].

The Goddard group also measured proton single-event upset for registers, cache, floating point and MMX units for Pentium III (P3) and AMD K7 processors [16]. They reported cross sections of about  $10^{-11}$  cm<sup>2</sup>. In their measurements very few single event upset were actually observed. There were problems in data collection, as the SEFI rate was sufficiently high as to impact the duration of runs. The data had to be collected with the Cache Off or the SEFI would have been too high to collect any significant data.

### B. Frequency Effects

One of the key questions regarding processor testing is the effect of clock frequency on test results. In principle, one would expect a higher upset rate when the clock frequency is near the maximum rated value. Such tests are difficult to do for several reasons. First, the noise margin and signal integrity of a test board that is modified for radiation testing may be sufficiently different from the conditions in a dedicated application to prevent the determination of the frequency dependence. Note that modern microprocessors operate at clock frequencies > 1 GHz. Second, processor generate a great deal of heat when operated at maximum speed. It is difficult to extract heat from a processor that has been modified to provide direct access to the chip (or the back of the chip) by an ion beam, and the increase in temperature that occurs at high frequency may further interfere with attempts to measure frequency dependence.

Fig. 14 shows recent measurements of frequency dependence for register upsets in an SOI version of the PowerPC. There is a slight increase in the cross section when the tests are done with a clock frequency of 1 GHz compared to tests at 350 MHz, but the effect is much smaller than implied by modeling studies for single-event transients. One reason for this is that although microprocessors operate at high switching speeds, the internal design has to be somewhat conservative in order to avoid yield problems.




Fig.14. Frequency dependence of register tests of SOI PowerPC processors.

The frequency dependence of Intel processors has also been examined, as shown in Fig. 15 [16]. The frequency dependence appears relatively weak for those results as well.

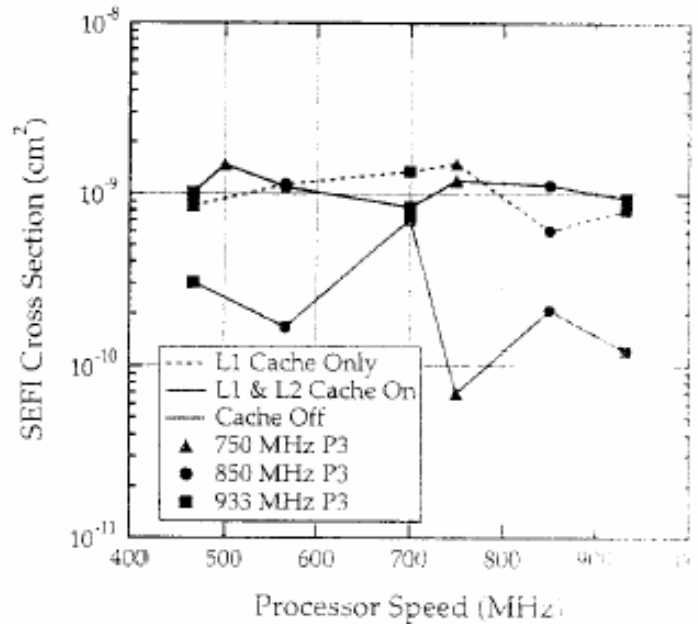


Fig. 15. Pentium III SEFI cross section as a function of the processor speed with various cache states.

## 6. Steps Required for Processor Testing

### A. Device Properties and Physical Preparation

The first step is to determine the basic properties of the device that is to be tested. One of the most critical features is the package type. If flip-chip bonding is used, then the thickness of the chip must be measured to determine if mechanical thinning is required in order to do the tests. Some processors incorporate heat sinks at the back of the die (top of package). The heat sinks may have to be removed or modified in order for the ions to reach the active part of the device.

The next step is to do the required modifications, and test the device afterwards to ensure that it still functions correctly.

### B. Hardware Requirements

Test boards must be fabricated (or, if commercial boards are used, adapted) that allow the device to be placed in front of the accelerator beam, with direct access to the device. Unless tests are done with an emulation system, other devices, such as memory, bridge chips and power control, are placed on the test board. The test board must be thoroughly checked out to ensure that the processor works satisfactorily at the frequency that will be used for the tests.

Special diagnostic methods, such as JTAG, require additional connections to the test board.

The hardware must include a temperature sensor to measure the operating temperature of the device during operation.

### C. Operating System and Software

An operating system must be selected. Primitive operating systems are recommended, as discussed earlier, but it is also possible to use more complex operating systems (i.e., Windows NT). However, it is far more difficult to distinguish functional errors in the processor from crashes in the operating system when a complex operating system is used.

Special software is usually developed for microprocessor testing unless the tests are intended to evaluate a specific software application. The general types of software that are required include:



1. Register tests, which load a predetermined pattern into several of the registers and continually evaluates the state of the registers, using a minimum of internal instructions in order to isolate register upsets
2. Cache tests, which are analogous to register tests, but specifically test the internal cache memory
3. Tests of specific instruction types or sequences

#### *D. Testing*

The first step is to select the facility that will be used and the properties of the ions that are required.

Once at the facility, the test hardware is placed in front of the beam, evaluating the performance of the hardware and software to ensure that it functions properly with the beam off. This is an essential step because the cable length, noise, and general interface issues may be somewhat different at the facility compared to conditions in a more conventional laboratory.

After the hardware and software are checked out, the device is temporarily removed from the beam. The accelerator is tuned for the specific ion energies required, using appropriate diagnostics to measure the particle flux (and energy, if required).

The next step is to place the device in the beam, turn on the accelerator and use the software and hardware to evaluate the microprocessor operation. Generally this is done at several levels, starting with tests of basic registers and progressing to tests involving more complex operations of the device.

That step is repeated for other types of ions, or for “degraded” ions where the energy has been reduced by inserting shields to lower the energy.

Test data are recorded during the tests, including measurements of the particle fluence for each test run. If functional errors occur, then the fluence at which the functional error occurred must be estimated using the diagnostic methods that were developed for the test.

## **7. Reporting and Interpretation of Results**

### *A. General Issues*

Test results must include a basic description of the approach used to test the devices, including the hardware and software that is used to evaluate the device, the operating frequency, and the operating system that is used (see Table 1, “Basic Test Approaches”). Because of the complexity of microprocessors and their associated test methods, a far more thorough description of testing details is required than for more conventional devices.

Commercial microprocessors tend to evolve rapidly, with a confusing array of part numbers and specifications. Thus, including the full part number in the report will not provide enough information for data interpretation. The details listed below must be included, in addition to the part number and date code.

- a. Package type
- b. Special treatment of the device, such as thinning or repackaging
- c. Maximum rated operating frequency of the device
- d. Core voltage
- e. Feature size

The properties of the ion beams used for testing must also be included in the report, including corrections for energy loss.

There may be specific features of a processor that affect the results. For example, some types of processors allow cache error correction to be turned on. Alternatively, some microprocessor tests are done with specific application programs, not with more basic tests that allow upsets in registers, cache and other regions of the processor to be determined. Although application-specific software results are

often the end goal of processor testing, the results tend to be of limited use unless the software is documented in a way that allows more general interpretation.

Finally, processors generate a great deal of heat, particularly when they are operated near maximum frequency. Device temperature should be monitored and reported.

### *B. Register and Cache Tests*

Register and cache tests can be treated in an analogous way to tests of static random-access memories. Just as for memories, it is essential to include error bars for counting statistics in the results. The type of pattern loaded into registers and cache should be included as well. The usual practice is to report such data as upsets per bit.

### *C. Functional Errors, Hangs and Crashes*

Functional operation is the most critical problem for microprocessor tests, and it also the most difficult feature to evaluate during testing. As discussed earlier, functional operation is somewhat dependent on the operating system that is used during testing. Unlike tests of registers or cache, functional errors are reported on an error per device basis. The basic features related to functional errors, hangs and crashes that need to be reported include:

- a. Diagnostic results that partially isolate the operating system from the results, including JTAG
- b. Fraction of the test runs that result in functional errors, and an estimate of the cross section for functional errors compared to the cross section for registers and cache
- c. Categorizing functional errors by the type of malfunction that occurs
- d. Steps required to restore operation, e.g., application of a RESET pulse, or power removal and complete restart.

## **8. References**

1. W. E. Will, *et al.*, "Total Dose Response of the Z80A and Z8002 Microprocessors," IEEE Trans. Nucl. Sci., **28**(6), pp. 4046-4051 (1981).
2. J. H. Elder, *et al.*, "A Method for Characterizing a Microprocessor's Vulnerability to SEU," IEEE Trans. Nucl. Sci., **35**(6), pp. 1678-1681 (1988).
3. R. Velazco, *et al.*, "Heavy Ion Results for the 68020 Microprocessor and the 68882 Coprocessor," 1991 RADECS Proceeding, pp. 436-440.
4. J. M. Kimbrough, *et al.*, "Single vent Effects and Performance Predictions for Space Applications of RISC Processors," IEEE Trans. Nucl. Sci., **41**(6), pp. 2706-2714 (1994).
5. F. Bezerra, *et al.*, "Commercial Processors Single Event Test," 1997 RADECS Data Workshop, pp. 41-46.
6. V. Asnek, *et al.*, "SEU Induced Errors Observed in Microprocessor Systems," IEEE Trans. Nucl. Sci., **45**(6), pp. 2876-2883 (1998).
7. D. M. Hiemstra and A. Baril, "Single Event Upset Characterization of the Pentium MMX and Pentium II Microprocessors Using Proton Irradiation," IEEE Trans. Nucl. Sci., **46**(6), pp. 1453-1460 (1999).
8. G. M. Swift, *et al.*, "Single-Event Upset in the Power PC750 Microprocessor," IEEE Trans. Nucl. Sci., **48**(6), pp. 1822-1827 (2001).
9. N. Seifert, *et al.*, "Impact of Scaling on Soft-Error Rates in Commercial Microprocessors," IEEE Trans. Nucl. Sci., **49**(6), pp. 3100-3106 (2002).
10. F. Irom, *et al.*, "Single-Event Upset in Commercial Silicon-on-Insulator PowerPC Microprocessors," IEEE Trans. Nucl. Sci., **49**(6), pp. 3148-3155 (2002).
11. F. Faure, *et al.*, "Impact of Data Cache Memory on the Single-Event Upset-Induced Error Rate of Microprocessors," IEEE Trans. Nucl. Sci., **50**(6), pp. 2101-2106 (2003).
12. F. Irom, *et al.*, "Single-Event Upset in Evolving Commercial Silicon-on-Insulator Microprocessor Technologies," IEEE Trans. Nucl. Sci., **50**(6), pp. 2107-2112 (2003).

13. F. Irom, *et al.*, "Frequency Dependence of Single-Event Upset in Advanced Commercial PowerPC Microprocessors," IEEE Trans. Nucl. Sci., **51**(6), pp. 3505-3509 (2004).
14. N. Renaud, "ATMEL TSC695F Radiation Hardened 32-bit SPARC Processor Single-Event Test Results," unpublished. ATMEL, Nantes, France 44306; [nicolas.renaud@nto.atmel.com](mailto:nicolas.renaud@nto.atmel.com).
15. A.H. Johnston, "Radiation Effects on Advanced Microelectronics Technologies," IEEE Trans. Nucl. Sci., 45(3), pp. 1339 (1998).
- 16.. J. W. Howard Jr., *et al.*, "Total Dose and Single Event Effects Testing of the Intel Pentium III (P3) and AMD K7 Microprocessors," 2001 IEEE Radiation Effects Data Workshop, pp.38-47.