National Aeronautics and Space Administration

FY11 End of Year Report
NEPP SOC Devices

Steven M. Guertin
Jet Propulsion Laboratory
Pasadena, California

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

JPL Publication 12-7 3/12

National Aeronautics and Space Administration

FY11 End of Year Report
NEPP SOC Devices

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

Steven M. Guertin
Jet Propulsion Laboratory
Pasadena, California

NASA WBS: 724297.40.49.11
JPL Project Number: 104593
Task Number: 40.49.03.04

Jet Propulsion Laboratory
4800 Oak Grove Drive
Pasadena, CA 91109

http://nepp.nasa.gov

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

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

Copyright 2012. California Institute of Technology. Government sponsorship acknowledged.

TABLE OF CONTENTS

Abstract 1
Notation 2
1.0 Background and Task Review 3
1.1 General Statement 3
1.2 Review of Last Year's Efforts 4
1.2.1 UT699 Testing 4
1.2.2 Maestro Efforts 4
1.3 Overview of This Year's Efforts 4
1.3.1 Why Move Off Maestro and UT699? 4
1.3.2 Review of SOCs of Interest to NASA 5
1.3.3 Test and Development Work on the Freescale P2020 5
1.3.4 P5020 Development and Test Effort 6
1.3.5 Continued Monitoring of Maestro 6
1.3.6 Collaboration 6
2.0 Seven-Step Approach 7
2.1 Collaboration with Manufacturers and Users 7
2.2 On-Chip Peripheral Approach and Prioritization 7
2.3 Fault Tolerant Device Test Approaches 8
2.4 RHBD Device Challenges to Test Development 8
2.5 Multicore Device Unique Techniques 8
2.6 General Test Methods 9
2.6.1 Test Hardware Selection 9
2.6.2 Test Hardware Operation 9
2.6.3 Test Device Preparation 10
2.6.4 Software for Basic DUT Sensitizing 10
2.6.5 Software for Application-Like DUT Sensitizing 10
2.6.6 Test Matrix 11
2.7 Collecting Results from Sample Testing 11
3.0 Key Types of SOCs 13
3.1 Background of SOC options 13
3.1.1 Key Input on SOC Survey 13
3.1.2 Key Types of Devices 13
3.2 Devices Examined for Potential Study 14
3.2.1 List of Potential SOCs for Study 14
3.3 Key Devices for this Task 14
3.3.1 Freescale P2020 14
3.3.2 Freescale P5020 16
3.3.3 Aeroflex UT699 18
3.3.4 Opera Maestro 19
4.0 Test Methods Under Development 20
4.1 Basic Methods 20
4.1.1 Low-Level Structure Testing 20
4.1.2 Circuit Elements/Blocks Testing 20
4.1.3 Application-Like Testing 21
4.2 Test Applications 21
4.2.1 DUT Tests Itself 21
4.2.2 Lock-Step Testing/Test Vectors 22
4.2.3 Debug Hardware and Software 22

4.2.4 Hardware Interaction and Development 23
4.3 Test Data Approach 23
5.0 Test Data 25
5.1 General Results this Year 25
5.2 Testing the P2020 25
5.2.1 Test Approach 25
5.2.2 Test Software Plan and Approach 25
5.2.3 Proton Testing 26
5.2.4 Heavy Ion Testing 27
5.3 Testing the P5020 28
5.3.1 Device Preparation 28
5.3.2 Test Results 28
6.0 Future Work 30
6.1 P5020 Work 30
6.1.1 Test Devices 30
6.1.2 Software and Hardware Development 30
6.1.3 P5020 Testing 30
6.2 P2020 Work 30
6.2.1 P2020 Test Upgrades 30
6.2.2 P2020 Testing 31
6.3 Future SOCs 31
6.3.1 Atmel AT7913 31
6.3.2 SiLabs C80501F310x 31
6.4 Maestro Efforts 31
6.4.1 JPL Testing 31
6.4.2 Other Testing 31
6.5 Collaborative Review 31
7.0 References 33
Appendix A. Acronyms and Abbreviations 34

ABSTRACT

This is the final report for fiscal year (FY) 2011 NEPP System-on-a-Chip (SOC) Devices task. This task seeks to provide improved methods for SOC testing to the NASA community, with the goal of providing data that is of interest to NASA and that meets the standards of the wider radiation-effects community. This report describes continued efforts to establish qualification and radiation testing methods appropriate to SOC devices.

This NEPP task seeks to understand the impact of radiation effects on devices of interest to NASA programs. Because SOC devices are combinations of several types of computer circuit elements, many types of test approaches are applicable for the various elements of the devices. The current work is part of a multi-year effort to establish test methods based on examination of existing methods, development of new methods, and in-the-field verification of methods by radiation-testing SOCs of interest.

The present report for FY11 discusses the efforts carried out under this NEPP task. This year, an examination of SOCs with NASA applicability was conducted, augmented by consideration of device types with minimal previous coverage as well as those that can expand the scope of the previous efforts. These were combined to provide the widest coverage of device types of interest to NASA, but were not so broad of scope as to limit the effectiveness of any test development or radiation-test validation of test methods.

This task is expected to continue, gathering information about test methods appropriate for NASA missions and developing test data appropriate for devices of interest. The longer term goal of the task is to develop a guideline describing a qualification approach and radiation testing methods. The methods discussed in this report, relating to multicore, RHBD, and multiple functional blocks on an SOC are expected to be developed further in the FY2012 follow-up effort under this task.

NOTATION

This document uses several specific notations briefly described here.

- 0xDDDD_DDDD (or 0xDD) is used to express numbers in hexadecimal format. The digits 0–9 and a–f refer to values from 0 to 15. '_' is used to aid the eye.
- 0bDDDD_DDDD (or 0bDD) is used to express numbers in binary format. '_' is used to aid the eye.
- %l, %o, %g, %i are prefixes for registers in the Sparc™ architecture.
- Endianness is always little endian, indicating that digit 0 is the least significant digit and is at the trailing (right) end of the number; i.e., 0×1234 has 4 as digit 0 and 1 as digit 3.

1.0 BACKGROUND AND TASK REVIEW

This NEPP System-on-a-Chip (SOC) Devices task started in fiscal year 2010 with the primary goals of (1) establishing appropriate test methods for SOC devices and related microprocessors for NASA, (2) investigating the appropriateness and methods for engaging device manufacturers in testing of their SOCs, (3) engaging the NASA community to determine the most desired test data to be collected from SOC and microprocessor devices, and (4) surveying available SOC devices (primarily microprocessor-based SOCs) to validate test methods and identify general strengths or weaknesses in various architectures.

1.1 General Statement

Here we briefly state the rationale for goals stated above. The general approach was to provide the most appropriate coverage of the SOC topic. Since there is significant overlap among technology-implementing and radiation-test organizations, engaging the community where possible was desired. However, the subject is vast, so it was also important to engage the community in order to streamline identification of test candidates and methods.

SOCs are the future of computer electronics and are likely to be injected into space missions soon. These devices present complex functional designs and sometimes include radiation-hardened-by-design (RHBD) methods. Appropriate methods for evaluating them for radiation effects are not obvious, and similarly it is not obvious how to assist missions in device selection from a radiation-qualification approach.

SOC evaluation is especially problematic because there is limited knowledge for how to evaluate an SOC's functional blocks and, therefore, even more difficult to determine the effects of the interaction of these complex functional blocks in concert. In fact, there is almost no evidence of current radiation-effects-test knowledge applicable to modern microprocessors, bridge chips, and interface buses [1-6]. Since these are not only present in modern SOCs, but also may occur in multiple instantiations, effort to understand test methods for these individual blocks is needed as well as a reasonable understanding of system-level effects which are not commonly addressed by radiation specialists. Furthermore, SOCs are sufficiently complex that full characterization would require an unreasonable amount of time [7]. Instead, this task is working toward a model-based approach for establishing application impact employing limited radiation testing.

In FY10, this task examined the Aeroflex UT699 LEON 3FT and Maestro SOC devices. In the case of the UT699 initial methods to test an RHBD SOC were developed. The findings indicated that an RHBD device may have error modes that are different from those found when testing commercial devices, because the weak elements that usually limit testing of a commercial device are likely to be less sensitive in an RHBD device. This necessitates the development of new test methods (outside of those common to commercial SOCs) in order to accurately identify the leading-order radiation sensitivities of SOC devices, which often are due to single event transient (SET) sensitivity [8]. In the case of Maestro, the efforts were targeted at how to migrate a manufacturer-type test kit to a radiation test. Examination of the test kit codes was carried out and the ten most appropriate test codes were listed. Both Maestro and UT699 efforts indicated that manufacturer involvement is very helpful, especially when unexpected phenomena (which often occur during radiation testing of complex devices) can be examined by collaborative efforts.

The current year's work focuses on identification of the most appropriate SOCs for current efforts and development of appropriate test methods for these devices. We surveyed technologists throughout NASA to determine a good candidate for focused development. The selected device was partially examined to establish an approach for more rigorous subsequent test method development. Finally, during this fiscal year we were able to establish some beneficial collaborative connections that will ensure continued success of our work.

The future of this task will be continued development of test methods appropriate to Freescale processors. This, along with possible follow-on work related to the UT699 will be used to develop the first version of a radiation test and qualification guideline for usage of SOCs. Additional effort will also be directed at collaborative efforts and ensuring appropriate SOCs are being investigated for NASA programs.

1.2 Review of Last Year's Efforts

This report continues the work of last year's SOC devices task. For context, we briefly discuss last year's efforts. The topics discussed here are the UT699 and Maestro work.

1.2.1 UT699 Testing

The UT699 was selected for examination under the SOC devices task because it has considerable functionality appropriate to NASA missions and because it is constructed with RHBD methods, which are also very important for specific space requirements.

When RHBD methods are applied to an SOC many of the most common error modes are removed and testing will uncover previously unexpected error modes. It is common in complex devices to observe error modes that have not been seen before; however, they are usually of incidental interest because other error modes are much more prominent. In the case of the UT699 it was expected that a new error signature, which would have to be captured during testing, would lead the event rate.

The actual radiation testing did uncover such an error mode. The observed error mode was called the "register partial reset" [9]. After collecting data on this error mechanism it was determined that the space rate for this event is only slightly higher than the rate for upsets in the RHBD latches that protect the execution pipeline (where upsets result in unpredictable crashes).

1.2.2 Maestro Efforts

During FY10 the Maestro 49-core multicore processor, under development by Boeing, was studied for appropriate test approaches for evaluating its single event effects (SEE) sensitivity. This work was conducted as part of a larger collaboration with other collaborators carrying out test software development and actual test data collection. Because of the broad interest of this NEPP task, we also examined test approaches under development and monitored test planning efforts.

Developments during FY10 for testing included a series of collaborative reviews of a special test board targeted at measuring parameters during functional operation of Maestro, and development of another test board targeting SEE testing of Maestro by GSFC. Because of the very complex nature of the Maestro device, the test board development is of interest here because it is indicative of future test issues likely to occur with complex SOCs.

1.3 Overview of This Year's Efforts

For FY11, the NEPP SOC devices task focused on moving to a new target family of SOCs. A survey of potential devices of interest was conducted, including general review of the market of SOCs and a focused review of NASA technologist input regarding key devices of interest. The Freescale architecture was selected for the next step, and limited testing was conducted. Some limited additional monitoring of Maestro SEE evaluation efforts was conducted. For this fiscal year, additional collaboration building activities were conducted.

1.3.1 Why Move Off Maestro and UT699?

During our review of last year's efforts, we determined that the task needed to focus on a new architecture for FY11, because (1) the Maestro efforts were largely explored in terms of benefit to this task; (2) the

UT699 efforts covered a large portion of the UT699 functional space; and (3) additional collaboration and device functional structures were also of interest, necessitating expansion to another SOC.

The NEPP SOC devices task is not focused on a particular architecture. Instead, the ultimate goal is to establish radiation testing and qualification approaches to the general category of SOC devices. As such, it is important to survey several different types of architectures and technology approaches. Although we do not intend to close the book on LEON-based processors, it was nevertheless the right time to move on. It should be noted that further LEON-based efforts may come forward in FY12. Also, although it is unknown how Maestro will develop, this task is open to additional work on Maestro or Maestro-like devices in future work.

1.3.2 Review of SOCs of Interest to NASA

Although reported in more detail in Section 3.0, here we briefly review the SOC review effort conducted this year. While reviewing the general needs, we determined that Freescale PowerPC™ and LEON/Sparc architectures are of the most immediate interest. Interest also exists in Maestro and similar high-performance options.

1.3.3 Test and Development Work on the Freescale P2020

Details of P2020 testing are also available in the report on P2020 testing produced by this NEPP task in FY11 [10]. Further information can be found in Section 5.2. This section briefly reviews the efforts conducted this year. This section quickly covers device under test (DUT) information, software development, and proton and heavy ion testing highlights.

1.3.3.1 DUT Information

In order to test the P2020, we needed a test vehicle. For this work we obtained P2020RDB development boards from Freescale. These boards cost much less than most similar test boards we have previously obtained for Freescale device testing. The P2020RDB provides multiple interfaces for high- and low-speed communication.

1.3.3.2 Software Development

The P2020 has the e500 core, which is very similar to the G3 processors, or the PowerPC 750, and JPL has experience testing this processor core. The existing software for testing the earlier Freescale processors translated well to the P2020, with the exception of hardware configuration and operation. For hardware details, some minor software was developed this year. The final test software was able to evaluate upset rates in one of the two cores, the L1 and L2 caches and, to a limited extent, the sensitivity of the memory management unit and coherency bus.

1.3.3.3 Proton and Heavy Ion Testing

Testing was conducted using the software approach described in the previous section. For this reason, the test results are limited to the microprocessor sensitivity.

DUT preparation was important, particularly for heavy ion testing. The P2020 has an embedded heat spreader which was removed over the silicon to enable heavy ions to penetrate to the die. In addition to the heat spreader, there is filler between it and the silicon. For this year, we established that a maximum LET of about 14 MeV-cm^2/mg would be possible without removing the filler material. Previous Freescale devices are known to have saturated the cross-section by LET = 14 [11].

Proton testing was conducted using a P2020RDB with no modifications. Testing was performed at Lawrence Berkeley National Laboratory (LBNL) using the 88-inch cyclotron, delivering 55 MeV protons. Care was taken to limit proton exposure to the SOC and to protect the surrounding electronics. During

testing, upsets of the L1 cache resulted in cache misses. Upsets of the L2 cache were easily measured because it was placed in RAM mode and the EDAC was disabled. One general purpose register (GPR) upset was observed.

Heavy ion testing was conducted with ions from LETs 1 to 14 MeV-cm^2/mg. Testing was performed on one of the P2020RDB boards, with the P2020 device milled to remove the heat spreader. Testing showed that the L1 and L2 caches were very similar (though not identical) in sensitivity. Due to a software bug, no GPR results were captured.

1.3.4 P5020 Development and Test Effort

The P5020 was selected as a more modern version of the P2020 (it is a 64-bit processor). Section 5.3 briefly covers the results of P5020 efforts. Additional information may be obtained from the test report for P2020 and P5020 from September 2011 [12].

For this year, the primary effort on the P5020 was limited to device preparation and limited radiation test efforts. Further work will be conducted in FY12.

1.3.5 Continued Monitoring of Maestro

NEPP has been interested in ongoing efforts to establish radiation sensitivity of Boeing's Maestro processor. This year there was a collaboration established to perform radiation testing of Maestro which was supported outside of NEPP but monitored for information gained relevant to this NEPP task. Results of the work on Maestro this year influence the recommendations for RHBD and multicore approaches and recommendations for future work, but are not discussed in detail.

1.3.6 Collaboration

Further development of collaboration opportunities were pursued this year. This was conducted in two directions. First, we looked for manufacturer involvement where new insight on structure and test methods would be possible. Second we pursued information internal to NASA to identify SOCs of interest. This project also benefits from input based on peer review, but for FY11 we did not directly seek this type of collaboration.

Manufacturer involvement this year provided a unique opportunity to understand how commercial architectures are likely to inject into aerospace applications. This follows primarily from Freescale-based processor and system architectures. Some discussion of this is presented in this report. Further, additional manufacturer input regarding LEON devices such as the UT699 has been invaluable for determining how NEPP can collaborate to evaluate a device and understand the impact of SEE on that device. Because of the nature of such collaborations, finding unexpected error modes may be very important for device development but may also put the manufacturer's program at unfair risk. All elements of the complex interaction between NASA and device manufacturers must be carefully examined in order to achieve a robust ongoing relationship with these important collaborators.

Internal NASA discussions on appropriate SOCs for use in a wide array of NASA missions were instrumental this year in determining, the relative importance of different types of SOCs for study. As a result of survey efforts we were able to isolate the three most important classes of SOCs at this time. Although interesting options from Intel and small PIC or 8051-based SOCs may be of interest to some NASA programs, these do not translate to the general needs and thus have been tabled in favor of LEON-based aerospace devices, commercial and special Freescale devices, and specially developed high-performance computing devices.

2.0 SEVEN-STEP APPROACH

During FY11, it became clear that this NEPP task has several thrusts that need to be clearly delineated. Some of them are clearly targeted at radiation test methods, some are targeted at collaborative methods, and so forth. Currently, we have identified seven primary thrusts:

1. Collaboration with manufacturers and users
2. On-chip peripheral approach and prioritization
3. Fault tolerant device test approaches
4. RHBD device challenges to test development
5. Multicore device unique techniques
6. General test methods
7. Collecting results from sample testing

Each of these is important for establishing the appropriate methods to be used for radiation qualification of SOCs, and each of these should be developed.

2.1 Collaboration with Manufacturers and Users

Collaboration with manufacturers and users is an important thrust of this task. The target devices are so complex that proper operation of the device can be difficult without support. In addition, RHBD devices typically have key structures that the manufacturer wants to demonstrate are able to handle SEE; thus, manufacturers may be interested in collaboration. However, these complex devices do not always lend themselves to easy identification of important use cases or key structures. Therefore, it is important to obtain assistance from both device manufacturers and application users.

During FY10 we tested the Aeroflex UT699. In both FY10 and FY11 we collaborated on test methods for the Maestro 49-core multicore microprocessor. These devices highlighted the difficulty of testing RHBD structures. The Maestro device work highlighted the trouble associated with very complex devices. The lessons learned are included at the end of this section.

Collaboration with users is important for identifying devices of interest and how to characterize them. For FY11 we established partners for selection of devices. In the future we will also discuss findings of test efforts in order to establish the relative value of the approaches developed for this NEPP task, which include:

1. Fault tolerant structures will get overwhelmed during high-rate testing
2. Test algorithms must test many different types of processor and peripheral operations
3. In cases related to errors that had not been previous observed, it may be necessary to modify algorithms to isolate the error signature
4. Extracting the signature of SEE sensitivity through overwhelmed fault tolerant (FT) devices requires carefully establishing a level of sensitivity achievable during testing
5. Many types of peripherals are of interest to users, including slow-speed interfaces like UART.

2.2 On-Chip Peripheral Approach and Prioritization

Because of the plethora of on-chip resources and limited testing resources, prioritization is very important for selecting appropriate peripheral testing. In addition, methods must be developed to enable testing of the highest priority peripherals. We have established methods for testing several of the peripherals. A summary of the findings thus far can be found in Table 2-1.

Table 2-1. Findings on methods to test key structures

Structure	Hardware Approach	Software Approach	Priority
SpaceWire Ports	Loopback connector	Full duplex transfer using designated data buffers	High
UART	IO with external console	Not fully developed	Medium
MMU	Use on-board memory connection	External memory test	High
Processor to Processor Bus	n/a	Establish message transfer during test	High
Ethernet	Ethernet cable to external computer	Message transfer (IO) during transfer – under development	Medium
Others	TBD	TBD	TBD

2.3 Fault Tolerant Device Test Approaches

Although FT device approaches were mentioned in Section 2.1, they also make up a key category for special consideration. FT methods make testing difficult and analysis tricky because the test team often gets locked into characterizing the sensitivity of the FT-protected elements. This type of examination is of limited benefit. Additionally, proper configuration of FT is not always obvious for user applications and care should be taken to point out the appropriate configuration to users.

2.4 RHBD Device Challenges to Test Development

During FY11, testing was targeted at non-RHBD devices in order to accelerate technology discussions. However, the impact of RHBD methods applied to SOCs is of considerable interest to this task. In FY11, we were able to present results of the UT699 testing which highlighted a vulnerability that may be similar to the types of challenges likely to appear in RHBD devices. These difficulties require effort to develop techniques to isolate measurements necessary for RHBD devices [5].

We are particularly interested in RHBD devices, because their error modes are not expected to be the standard register or SRAM storage cell vulnerabilities. It is expected that RHBD devices will be designed against errors in such structures either through hardening or error mitigation methods. Once these are removed from the error candidate list, we are left with storage bits with very small numbers of bits, or SET sensitivity on clock edges and significant frequency dependence [8].

2.5 Multicore Device Unique Techniques

Modern microprocessors have multiple cores, and this hardware trend is likely to continue for two reasons. First, the sizes of the cores are getting smaller and on-chip resources are making communications easier to handle if distinct boundaries are made on the chip. Second, the processing capability of individual cores is now high enough to be sufficient for most operations, so the desire to perform multiple complex operations simultaneously is a bigger driver for technology development.

It is expected that the current trend of multicore processors will continue and further integration of interconnect peripherals on chip will become more important. Thus, a major thrust of this effort is to identify methods for determining SEE sensitivity of individual processor cores, and to distinguish this sensitivity from that of the communications networks that enable the cores to communicate with each other. This topic is expected to be further developed in FY12 as more detailed information about multicore communications test methods and sensitivity is collected.

2.6 General Test Methods

General test methods are very important for SOC efforts, because trying to bring together a test system for a target device and ensuring adequate testing is performed to meet program needs can entail considerable resource expense. Reducing unnecessary testing and improving the quality of used test methods can result in significant increased test value.

Some general test methods have been established based on earlier work, such as the microprocessor test guideline developed for NEPP [11]. However, many of the general test methods for SOCs require rethinking earlier approaches, because modern devices are more similar to computers than individual microprocessors. We have established several areas for test development, and each is addressed in the following subsections.

2.6.1 Test Hardware Selection

Test hardware should be limited to available demonstration boards, preferably with the option to provide off-board power to the DUT. This ensures that all the desired hardware interfaces are available, and that desired application-like software can be used on the target board.

Because spacecraft rarely need audio/visual capabilities, or high-density in-system data storage (such as a hard drive), it is not necessary to obtain full demonstration units, and in many cases a much simpler limited-capability board may be sufficient. One example of a limited capability board that seems sufficient for this work is the P2020RDB discussed later in this report. Boards built recently as COM Express modules may be another beneficial option.

DUT mounting and tool access (i.e., mill or acid etcher) to the DUT position are also very important. If it is possible to obtain a package layout of the DUT before procuring test boards, the information may be used in selecting the test board or choosing the plan for beam exposures.

As a final hardware selection criterion, access to the DUT for beam exposure should also be taken into account. Specifically, if angles must be used, an unobstructed path to the DUT at the desired angle should be verified. Also, if the beam cannot be safely collimated to only strike the DUT, selection can also take into account the type of structures that are near the DUT. Also, depending on the relative radiation hardness of the DUT, peripheral circuitry may provide undesired background events – especially in proton testing [13].

2.6.2 Test Hardware Operation

During testing there will be several challenges to test hardware operation. The first is making sure that the test board is powered in a controlled way. In fact, depending on the technology node, it may be possible to completely ignore this area. If knowledge of the manufacturer and technology indicates that there is no significant high-current or device-damage risk on the DUT, it may be reasonable to use an ATX supply to power a test board (note that this should only be used if there is no other option, or if sufficient opportunity exists to redesign the power approach on-site in the event of unexpected behavior or board damage).

Another major operation problem is very long communications cables. In the case of the Indiana University Cyclotron Facility (IUCF) (recently renamed the Integrated Science and Accelerator Technology (ISAT) Hall), for example, you will need 60-foot cables between the beam room and the counting room. At the NASA Space Radiation Effects Laboratory (NSRL), this number can go up to 200 feet. Because of this, testers should plan for communication on interfaces that support such distances.

2.6.3 Test Device Preparation

Test devices must be prepared for most beam sources. Even in proton facilities, heat sinks and cases require modification to support energy determination. For heavy ion facilities such as TAMU, devices must be mechanically modified. For devices in ceramic packages with metal covers, the covers must be removed. For flip-chip devices with the substrate exposed for contact to a heat sink, they may need to be milled (depending on the required LET range), and they may require verification of reliable operation without a heat sink, or modification of the heat sink approach to support only partial coverage of the DUT.

Another type of package that is more common in modern devices is for a flip-chip to also have a heat spreader installed over it for contact with the heat sink. In these types of packages, recent work has proven that significant risk may exist when modifying devices. The recommended approach is to determine how hot the device runs during nominal testing. If the temperature stays below 50°C (chosen arbitrarily), it is likely ok to remove the heat spreader and use some sort of cooling line. If the device quickly reaches temperatures in excess of 100°C, it may be likely that permanent damage will occur if thermal regulation is lost even for a brief period – on these devices we recommend only removing the heat spreader over small regions of the die, and then attaching a modified heat sink to the remaining portion. If the temperature is between 50ºC and 100ºC, some in-between variation may be attempted, but careful monitoring is recommended during operation.

2.6.4 Software for Basic DUT Sensitizing

Based on the recommendations for processor testing in [13], it is recommended that basic DUT sensitizing be performed with assembly language based test code that has complete operational control of the DUT (and may actually be built into boot code, but in all cases must not allow calls to library functions outside of the code written for the test; i.e., it cannot use printf).

Basic structures on the DUT are best tested by directly operating them as desired for determining sensitivity. In the case of registers or caches, operation can easily be targeted at static data storage or data transfer operations. This can be contrasted with running an operating system with a special test program developed in C (perhaps with inline assembly), which relies on a programming environment and threading support built into the operating system to ensure the test program does not have complete control over the processor. In addition test programs developed in higher level languages (here C is such a language) use library calls and operating system I/O operations which can expose test software to significantly more SEE sensitivity than desired during testing.

The main structure of the SOCs of interest here is that of one or more processors with support buses and peripherals. Thus, the main thrust of testing can target the processor(s) and the approach mentioned here can easily be extended to key bus or peripheral features.

2.6.5 Software for Application-Like DUT Sensitizing

The method suggested in Section 2.6.4 has some distinct risks when it comes to predicting complex system behavior. Because the basic sensitivities are known, some approach must be taken to establish system-level sensitivity. The standard approach is to take the basic sensitivity for each primary static element in the system (for example, there may be one type of structure for registers and another for memory bits) and multiply it by the number of elements in the part.

There are two problems with this approach. The first, which results in conservative estimates, is that a processor is rarely sensitive to all of the bits in the design. Instead, application usage greatly affects the importance of most of the bits. The second, which may result in underestimated rates, is that not all of the underlying structures relevant for normal operation may be tested. A very good example of this is a buffer that is not clearly documented but may hold synchronization data for the MMU in the case of multiple

processor cores working in similar memory ranges. This buffer may not be documented in a way where its relevance to the system can be correctly included in estimates of system sensitivity.

Thus we suggest at least a basic effort to test SOCs with operating systems similar to potential flight applications. Because such systems are likely not available, a reasonable alternative is to run a Linux operating system with a test program written in C or a similar language. A multicore test program would be recommended for a multicore SOC.

2.6.6 Test Matrix

Care must be taken to choose the test matrix appropriate to the SOC of interest. In the case of commercial devices, cell-level sensitivity of very sensitive elements often will dominate any event rate calculations, and these are likely to reach saturation at very low LET (or at least have such a high event rate due to low LET ions that higher LETs can be ignored).

A good example of this is the Freescale processors discussed in [1]. The SEE behavior of these devices is limited by upsets to their registers and caches (and some other structures built out of the same SRAM cells that make up the registers and caches such as the branch history table). These SRAM structures have been observed to reach saturation at an LET of about 8 MeV-cm^2/mg. Because of this low saturation LET, these devices do not generally need package modifications (unless there is a heat spreader or heat sink in the way) in order to be exposed to LETs over 10 MeV-cm^2/mg using 40 MeV/amu Argon at TAMU.

The primary areas of interest in the test matrix are:

1. LET
2. Multiple DUTs
3. Operating voltage
4. Tilt angles
5. Rotation angles
6. Operating frequency (especially RHBD devices)
7. Test data patterns (and location tracking – i.e. does address 5 only show “1” to “0” errors while address 15 only shows “0” to “1” errors?)
8. Software packages
9. Beam flux (especially important in devices with error-correction capability)
10. Beam fluence per run (especially important because there may be unknown error-correction or redundancy structures impacting results).

2.7 Collecting Results from Sample Testing

The last major point of the SOC radiation evaluation approach is the collection of results from sample testing. This is a very important part of the approach because it is instrumental in verifying that the recommended testing yields operational results consistent with expectations. It is likely that when establishing SOC test approaches, many ideas will not work when applied to real devices.

This task has targeted collection of test data on the Aeroflex UT699 and Freescale P2020 and P5020 devices. These devices represent the RHBD Application Specific Integrated Circuit (ASIC) family (UT699) and commercial multicore processors (Freescale devices). Results on the former have shown that RHBD device testing suffers from difficult trade-offs between time-per-beam-run and fault tolerance

limitations. Results from the latter have shown that modern commercial devices are difficult to mechanically prepare for testing and are likely to become even more difficult.

3.0 KEY TYPES OF SOCS

One major effort this year was a review of prospective SOCs that might be of interest to NASA programs in the near term. In this section we discuss the review effort and findings.

3.1 Background of SOC options

This task seeks to find a good study area that applies to many existing or potential NASA programs while also staying cutting edge and collecting a stable of SOCs of interest to provide depth and breadth to the work. When choosing candidate devices we hope to address all of the following elements, although several different SOCs may be required in order to achieve all of these parameters.

1. Devices used to run computer systems (i.e., containing processors similar to flight options)
2. Devices of modern processing, such as 45 nm, 32 nm, or smaller feature size
3. Devices built with significant RHBD and/or fault tolerance features
4. Devices specifically designed for the aerospace market
5. Devices that can enable very low power systems
6. Devices that integrate many types of peripherals
7. Devices with multiple or many processor cores
8. Devices with cutting edge on-chip organization or communications systems.

3.1.1 Key Input on SOC Survey

For this survey we relied on publically available information, existing study areas, and direct communication with collaborators. Specifically, collaborators who provided direct input include Raphael Some (JPL), Paul Marshall (GSFC), Wesley Powell (GSFC), Yutao He (JPL), and Greg Cardell (JPL).

3.1.2 Key Types of Devices

SOCs can be generally divided into two types. The first is the genuine system-on-a-chip that seeks to provide all resources necessary for operation on the die. The second (and more common in the commercial marketplace) is the SOC that is the result of merging some of a computer's resources into a single die. The latter is perhaps a more pragmatic version of the former, as applied to personal computer systems [14].

We found that both types of SOCs have potential for use in NASA missions because the purer SOC (that has most or all functions on a single die) is well-suited for very low power missions. However, most of the structures on these devices are actually rather simple from a technology standpoint, and there may be a few very specialized (and potentially complex) peripherals that can make studying these devices problematic for the goals of this work.

The more common modern definition is that which is most likely to suit the needs of this task, because the goal is to establish general methods to handle qualification of SOCs. In this regard older technology with chip-specific peripherals cannot lead to a good general guideline. For this reason, the findings of the survey of potential candidates highlighted commercial SOCs (as opposed to RHBD SOCs which have been the principal subjects earlier in the work of this task). It should be noted, however, that the more traditional SOC should be handled in a general way in the guideline and at least one such device should be examined under this effort in the future, if possible.

3.2 Devices Examined for Potential Study

Based on a survey of available SOCs, and utilizing the inclusive definition of SOC, we constructed a list of potential devices for study in FY11. This list was based on examination of commercially available complex integrated circuits (ICs), devices available from aerospace sources, key devices under development for space programs, and devices with highly integrated single-chip capability. The list of devices we used to drive discussion and assist in selection of an appropriate device for additional work on this NEPP task is presented below.

3.2.1 List of Potential SOCs for Study

1. Freescale MPC8641D
2. Freescale P4080
3. Freescale P2020
4. Freescale P5020
5. Atmel AT7913
6. SiLabs C8051F310
7. Freescale PowerQUICC II (82xx)
8. National Semiconductor LM98640QML
9. Intel 80579
10. Texas Instruments AM389x
11. Marvell-based SOCs
12. STMicroelectronics STM32F103x
13. Analog Devices ADSP-BF542
14. STI Cell Processor
15. Aeroflex UT699
16. Opera Maestro (Boeing)

3.3 Key Devices for this Task

We narrowed the list provided in Section 3.2.1 after focused consideration of technologies appropriate for current and future NASA missions. A subset of devices has been identified for the current work on this task. This subset consists of devices of current and ongoing study and consists of the following: Freescale P2020, Freescale P5020, Aeroflex UT699, and the Opera program's Maestro. In addition we expect future expansion of the list to include the Atmel AT7913 and SiLabs C8051F310x (see more details in Section 6.0).

3.3.1 Freescale P2020

The Freescale P2020 has become an important vehicle for work on the NEPP SOC devices task, because the P2020 is readily available in a low-cost package. For this year's work, P2020 efforts were conducted using the Freescale P2020RDB development board. This board is a low-cost vehicle for testing of the P2020 with protons or heavy ions. Heavy ion LET is limited to about 15 MeV-cm^2/mg due to the structure of the part and difficulties with delidding. However, this system provides access to the peripherals and hardware interfaces of the greatest interest to this task, while keeping individual part costs to a reasonable level.

3.3.1.1 P2020 Architecture

The block diagram of the P2020 is presented in Figure 3.3-1 below. As shown in the block diagram, the P2020 has two e500 (32-bit) Freescale processor cores tied together with a system bus and mated to a memory controller and a host of high- and low-speed communication interfaces. This architecture provides a relatively easy system for porting existing Freescale processor test algorithms over to a multiprocessor system with built-in peripherals.

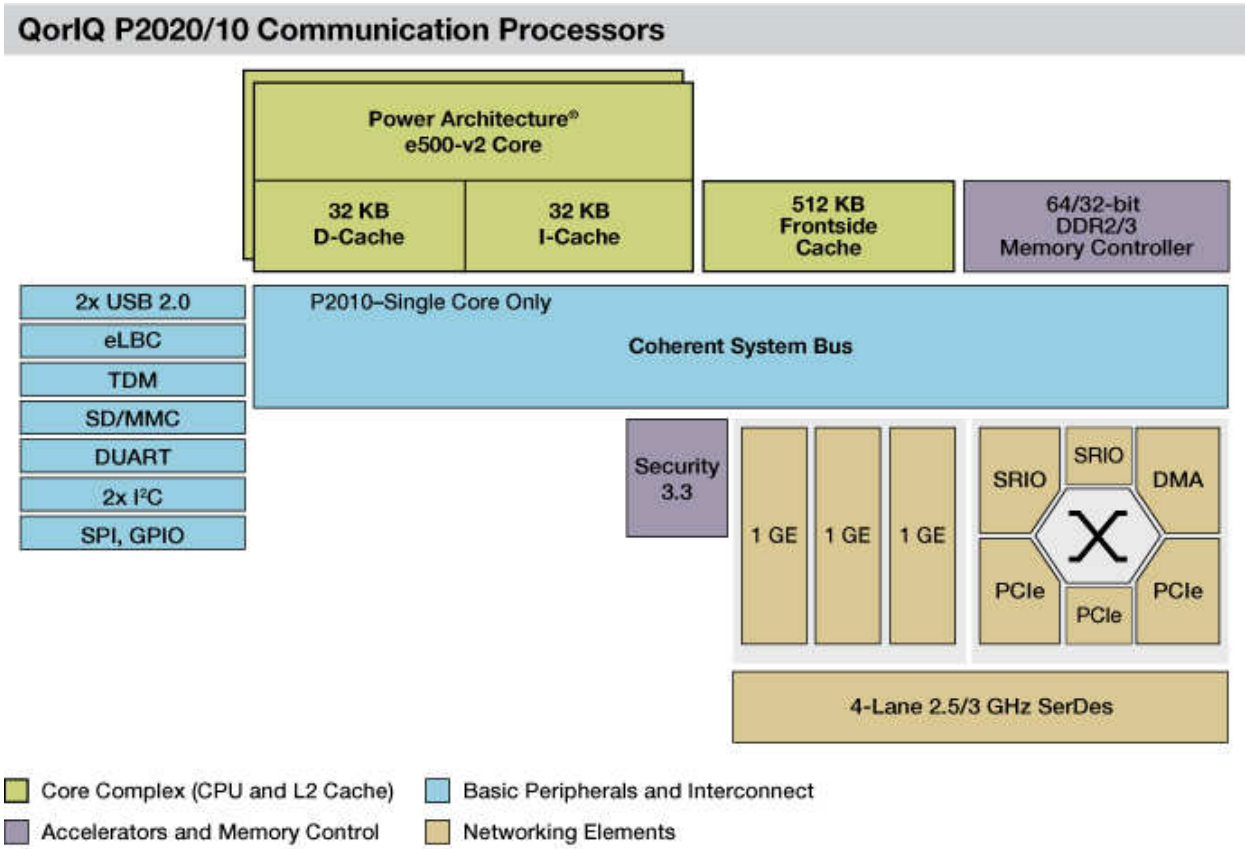

Figure 3.3-1. Block diagram of the P2020 device.

3.3.1.2 The P2020RDB Development System

The P2020 was also selected because obtaining test devices was relatively easy and the test systems were easy to work with. In Figure 3.3-2 the P2020RDB development system is seen from the back of the box showing the plethora of communications connections (including the SD card port which connects directly to the P2020 SOC). See Section 5.0 for more details on this year's testing.

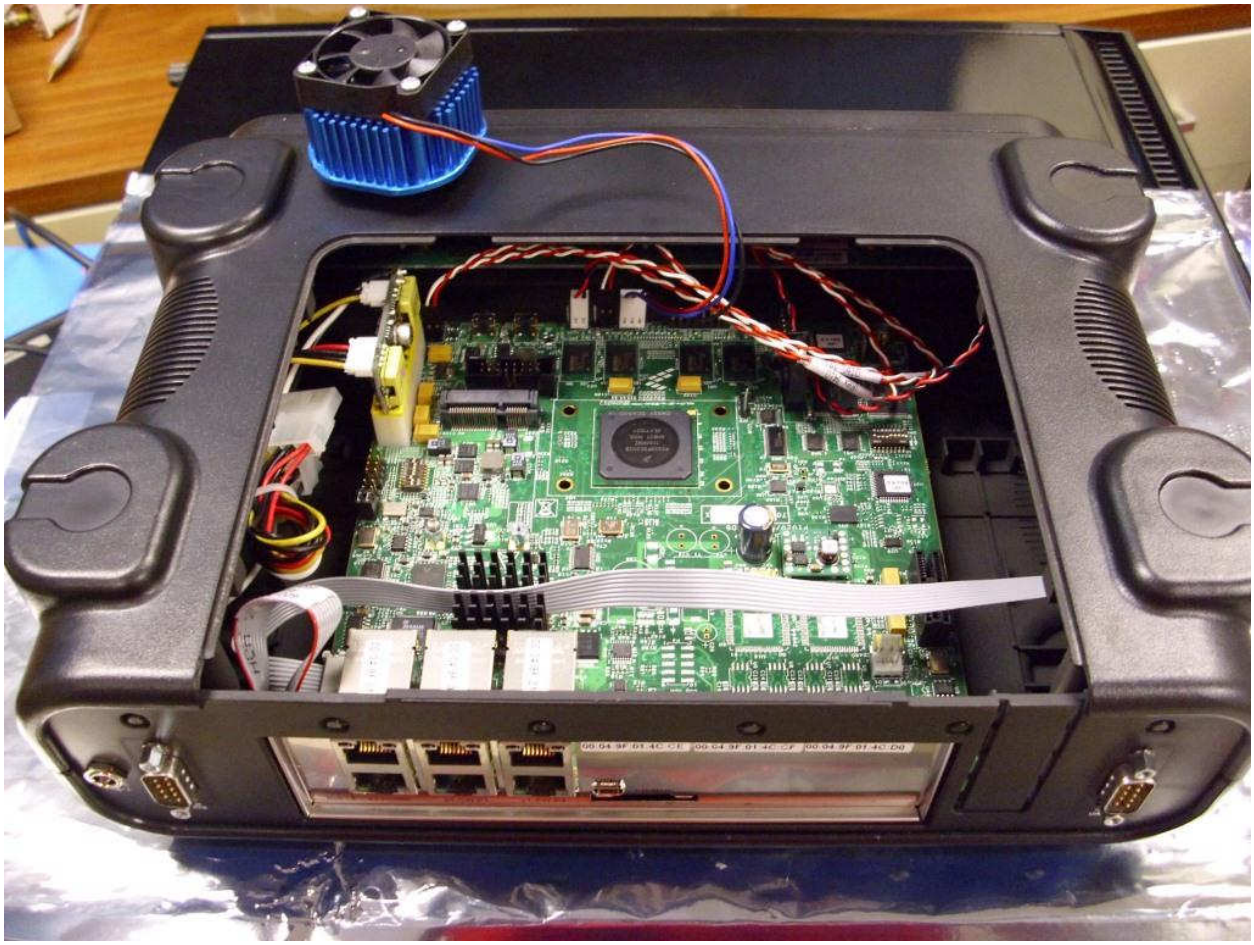

Figure 3.3-2. The P2020RDB development system is an inexpensive but very functional platform for the P2020.

3.3.2 Freescale P5020

The Freescale P5020 is a higher functionality version of the P2020, with 64-bit processor cores and more peripheral ports. The P5020 is intended to support a full computer system (while the P2020 is more suited to purpose-programmed communication boxes or application systems). The block diagram of the P5020 is shown in Figure 3.3-3. Compared to the P2020, there are clearly more peripherals, and they are designed for interfacing through the CoreNet system rather than the P2020's "Coherent System Bus". The P5020 also uses the 64-bit e5500 processor core, while the P2020 uses the 32-bit e500 processor core.

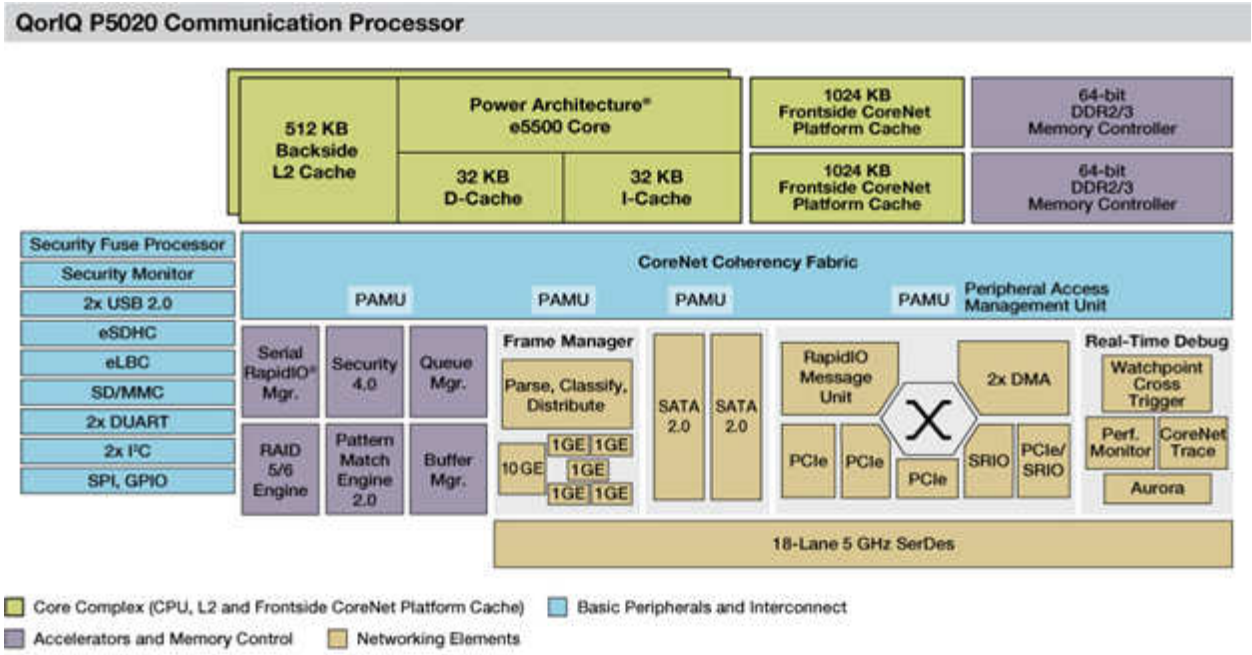

Figure 3.3-3. Block diagram of the P5020 processor.

A limited amount of testing of the P5020 was performed this year. This testing required careful consideration of device cooling. Difficulties with device preparation and complexities with the structure of the P5020 inherently limited test capability for FY11. Testing was performed using the P5020DS (Development System), and removing a portion of the DUT's heat spreader, but only over the portion of the chip that covers the structure of interest for testing (in this case the processor cores). The P5020DS with DUT ready for irradiation is shown in Figure 3.3-4. See Section 5.0 for more details on this year's testing.

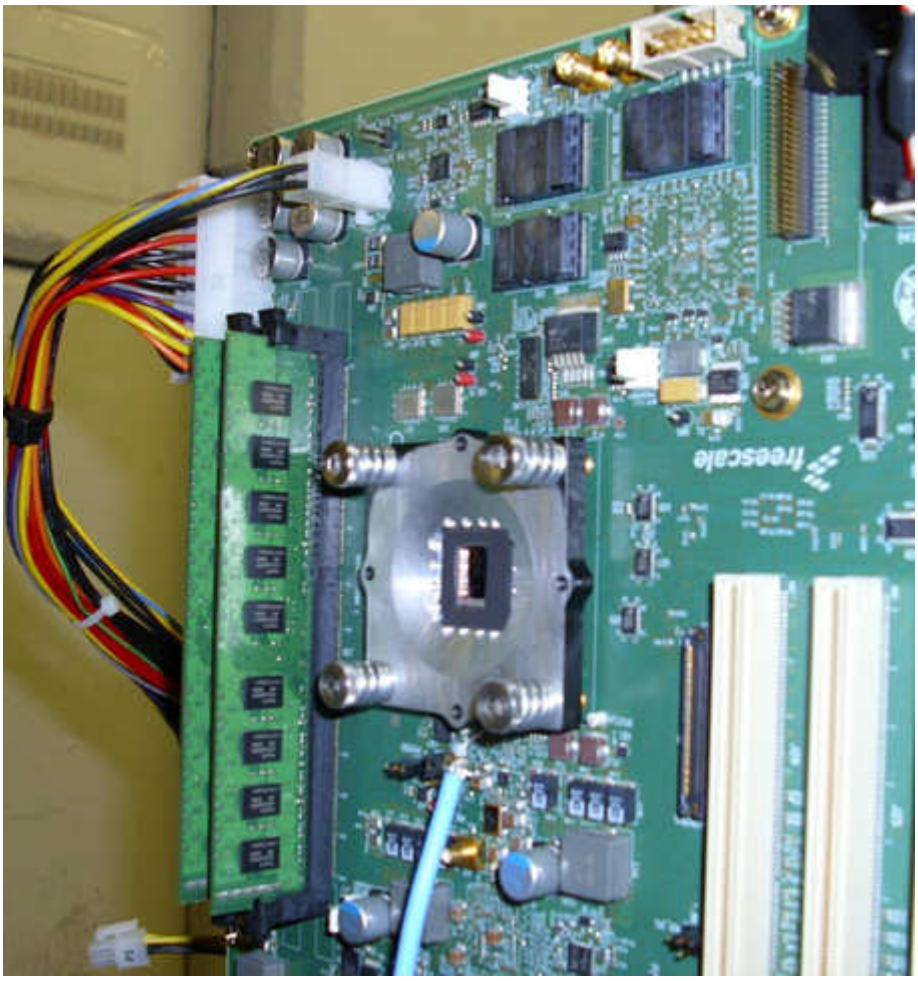

Figure 3.3-4. An example of a P5020DS system with processor ready for exposure to heavy ions.

3.3.3 Aeroflex UT699

Tested in 2010 under this task, the Aeroflex UT699 is a good vehicle for examining the radiation performance of an RHBD SOC, and for examining the relationship between NASA testing and manufacturer involvement. The UT699 was tested between 2008 and 2010 by Aeroflex with NASA involvement, and by NASA with Aeroflex involvement. The collaboration increased the chances of success during testing and increased the ability to examine an unexpected error mode observed during testing called the "register partial reset."

Also of interest to this task were the specific nature of the testing methods for this part and the test plan structure. It was observed that standard microprocessor test methods required modification to deal with an error mode that was not expected. This error mode highlighted the need to directly observe the test data and compare to on-chip error monitors, and to understand the fault structure of the FT device as it impacts error handling. A further finding was that in the case of RHBD and FT devices there is a possibility that an error may occur that sits at the edge of detectability with a reasonable test run time – i.e., we observed events that took ~1 hour to observe and event-elimination based on FT (e.g., that overwhelming EDAC may have caused the event) can only be done by increasing the run time, which is unreasonable. Note that although increasing flux can be used to show flux dependence, it cannot reliably be used to eliminate a possible underlying event sensitivity without flux dependence. However, increasing the flux and not observing flux dependence would indicate a lack of flux dependence.

3.3.4 Opera Maestro

The Opera program's Maestro (made by Boeing) has been examined as a case study under this effort. In FY11 a JPL program under NASA Enabling Technology Development and Demonstration (ETDD) program performed SEE testing of Maestro, utilizing development efforts spread across several collaborators spanning over several years, including this NEPP task. Maestro is a 49-core many-core microprocessor with five on-chip networks connecting the cores to each other and the chip's many peripherals.

During FY11 test efforts indicated that the significant collaborative and manufacturer involvement (both from Maestro and Tilera) were invaluable to successful testing of the device. Maestro performed very much as expected; however, due to constraints on the test system and goals of the program, many interesting error modes were below the level of sensitivity of the testing. Understanding the manufacturer's debugging tools proved invaluable to examine the methods under which test tiles stopped functioning.

Significant understanding of how to force the test tiles to remain functional was also extremely important for test system development. This tricky detail may have significant impact on other SOC testing, so a few more words are in order. One school of thought on testing is to reset the DUT into a known state very often (~100 ms) and allow the processor to run for a brief period in order to observe evolution of the system state. Unfortunately, because of the very weak structure of the caches, it is virtually guaranteed that some of the tile cache bits will be affected, even in a short period of time. Individual bits will not yield errors during execution (however, Maestro does have a problem where L1 data cache bits will incorrectly cause interrupts), but double-bit errors will yield errors, and any observed errors in the data stream must take into account the possibility that they were caused by double-bit errors. However, the chip itself can actually reach a much higher cache scrub rate than once per 100 ms by performing its own periodic scrubbing. This method was used during testing, and a data downlink rate of only ~20 kbps was sufficient for downloading tile-level test results.

Without utilizing the ability of the test device to identify and clear out errors, the data log would have included all errors and it would have been necessary to post-process the results to isolate potential signatures, in order to divide events into categories—this is a daunting job. But the counter to arguing for on-site test design is that it necessarily leads to ill-defined test plans. However, most complex device tests necessarily result in multiple efforts to perform testing, and a well-defined test plan may result in an ill-defined approach to get the most important data.

4.0 TEST METHODS UNDER DEVELOPMENT

Determining radiation performance of an SOC is not a straightforward problem. Because SOCs include a microprocessor it is reasonable to assume the problem is more difficult than the microprocessor evaluation case. Microprocessor testing is outlined in [1–6] and a guideline can be found in [11]. By looking at all these efforts, it is clear that there are many different ways to evaluate a microprocessor, and the application sensitivity can range at least a factor of 100 from the error rate predicted by using another application's observed sensitivity during radiation testing. As a result, we expect that the radiation sensitivity of an SOC will have many dependencies.

Several different types of test methods can be used for microprocessors, and they generally can be applied to SOCs in a straightforward manner. Specific methods will be described later in this section. But first it is important to point out that the methods do not directly translate to application sensitivity (unless the target application was the application tested). In order to apply test data it is necessary to establish models to take each test method's test results and find application sensitivity. The combination of a test method and the model approach used to apply its data to establish application sensitivity can together be referred to as a test approach or test method. If two or more approaches are used, then the difference in the predicted application rate among the approaches is referred to as "model dependence". When only one approach is used, the commonly observed model dependency due to the given model can also be used to estimate the model dependency of using the current approach.

Some of the information here also is presented in Section 2.6 since general test methods are developed as one of the seven major thrusts of this effort. This section is intended to provide a more conceptual discussion of how to test from which some of the earlier information can be derived.

4.1 Basic Methods

Several common methods for testing can be used. They include methods to test individual low-level structures, testing of specific circuit elements, functional testing, and application-like testing. These methods largely cover all of the likely methods a test group may want use. In this section we discuss each method and the pros and cons associated.

4.1.1 Low-Level Structure Testing

Fundamentally, ICs are made of very few unique low-level structures. If the radiation sensitivity of each of the unique low-level structures is known, models can be used to extrapolate the behavior of the device. Testing only needs to determine the SEE response of the different types of structures used by the manufacturer.

Examples of low level structures are the following:

1. SRAM cells
2. Latches
3. Combinatorial paths
4. Linear devices/sensors
5. PLLs
6. State machines.

4.1.2 Circuit Elements/Blocks Testing

In some cases modeling the low-level structure contributions is not as meaningful as directly examining target elements or blocks for SEE examination. The two main cases where this happens are: (a) sensitivity

is dependent on ill-defined portions of a circuit such as connection length between nodes, or (b) normalization of the measurement is difficult.

Circuit elements are expected to be more important for SEE response when individual low level structures are hardened against SEE. For example, if the individual cells of a L1 cache are hardened to have a high threshold LET for upset, then SETs in the control lines might be more important for SEE response.

The second case where we may want to concentrate on the block-level, discussed above, is when normalization is difficult. An example of this is data transfer. We may want to test an Ethernet buffer by transferring data through it at a particular throughput and normalize errors to the number of transferred bytes. Fundamentally, the buffer is built of control registers and SRAM cells, and the error rate is possibly be due to the number of SRAM cells used or it may be due to only a subset of them; in a real application it may be the case that the SRAM cells only hold data 1% of the time or the data held only shows a small portion of the errors. In this case it may be very difficult to normalize the error rate to the size of the Ethernet buffer. So, instead, we may want to take a block-level approach and treat the device structure as a functional unit.

The Spacewire results presented for the UT699 in [5] are an example of a circuit element/block test.

Key elements of interest to this task (i.e., for which test methods are being developed) are the following:

1. Microprocessors
2. Bus data paths
3. Bus controllers
4. Memory managers
5. Communication structures.

4.1.3 Application-Like Testing

The principal problem with the test methods discussed above is that they may not be inclusive of all the structures in a device. That is, they cannot be used to directly provide an upper limit on the SEE response (rate and significance of errors). In order to provide this limit, an application-like test can be performed. In this case an application that is very similar to the flight application is used to sensitize the DUT during irradiation. Since it is flight-like, the observed error rates will establish the nominal error response. If desired, lower operating voltages can be used to provide a worst-case limit to the device response.

4.2 Test Applications

In order to perform testing of the methods discussed above, an application must be selected which sensitizes the DUT to radiation. In this section we briefly discuss applications appropriate for SEE testing of SOC devices.

4.2.1 DUT Tests Itself

The DUT-tests-itself method utilizes software executing on the DUT to provide sensitivity of resources to SEE. This type of test application is very useful for establishing the actual sensitivity of the SOC to SEE under conditions where code is executing. The specific type of code used to provide the DUT sensitivity can be as test-specific or flight-like as desired. In fact, test code could consist of flight software performing calculations or communications identical to flight.

The key limitation of this approach is that the test device is also the controlling device. As such, significant risk of mixing desired observations with upsets of the test system exists. A few options exist to limit the risk that the device corrupts test data. One is to write the test code at the machine level and shut

down the operating system; this increases the control wielded by the test engineers and reduces misleading data. Another method is to provide many different ways to test the same target and, therefore, provide independent means of verifying any observed rates.

4.2.2 Lock-Step Testing/Test Vectors

One method of testing a complex device is to have two duplicate devices and have a cycle-by-cycle comparison of the operation of the devices [15]. Any deviation in the operation of the two devices would indicate occurrence of an SEE. Typically, this method requires special hardware for operating the devices together and for comparing their operation.

This method clearly identifies the events that impact a device, and the approach is well-defined. The method can be unambiguously used to test any desired type of application or targeted element in the device. This method may be particularly beneficial in the extension to SOCs (because it is very difficult to define what it means to test an SOC), while this approach would clearly identify when operations do or do not provide identical responses.

Unfortunately, this approach has significant problems with modern devices. In particular, two key problems stand at the front. The first is that modern devices are often designed to respond to errors with recovery mechanisms, and this test approach would incorrectly count a recovery attempt as an error. The second is that modern devices, and especially SOCs, are very likely to internally operate for significant amounts of time without having an observable external signature. The latter is likely to limit the ability of the test approach to reliably identify when a SEE occurred. Although these two problems are big enough to discount this method altogether, there are other problems with this approach.

Other problems with this approach include the following. Testing SOCs under this method would require development of custom test boards with two mounted devices. This is a very difficult engineering task and probably far too expensive for a reasonable qualification approach. Another problem is that the IO devices in an SOC cannot really be tested in this way. Similar to the previous paragraph, many IO devices are actually designed to correct transmission errors. Yet another issue is that internal configuration of an SOC may be changed without producing an external signature, but some applications will have sensitivity to the changed configuration.

Based on the discussion here, the lock-step approach to testing cannot reasonably be recommended as a test method on an SOC.

4.2.3 Debug Hardware and Software

Another method for testing an SOC is to use debug interfaces provided by the manufacturer. These interfaces may utilize special debug ports (such as in the UT699 debug support unit) or be specialized software and hardware designed to take advantage of JTAG ports.

Utilizing debug hardware and software can be invaluable for determining how the microprocessor or peripherals are responding to the beam, but they cannot provide an execution environment similar to actual flight conditions. In particular, this equipment and software rely on external control of the processor, and they do not provide either the same type of clocking or the same type of duty cycle as flight conditions. (It should be noted that debug hardware and software can be used for diagnostic purposes after an unexpected behavior is observed. However, if hardware or software is used in this way care must be taken to ensure the debug operations did not change the device SEE response.)

Because of the indicated limitations, it is recommended that test personnel be familiar with debug hardware and software for debugging purposes. For test purposes, these tools should only be used for static testing such as the determination of the cross-section for cache or register bits.

4.2.4 Hardware Interaction and Development

SOC testing methods discussed earlier in this section primarily cover the operation of the core of the SOC, containing the microprocessor(s), main communications bus, and memory controller. However, a major portion of SEE evaluation of an SOC rests in understanding the sensitivity of the peripheral devices. This should be separated from the core elements because the peripherals can largely be tested separately from the core elements unless the core elements are so upset-prone that their errors will interfere with observing the error signatures of the peripheral elements.

All peripherals needed for most programs that would use SOCs should be examined. Methods are needed to establish radiation sensitivity of the following peripheral devices.

1. UART interfaces
2. Ethernet
3. General SERDES ports.
4. PCI bus
5. Spacewire bus

Testing of these SOC components only makes sense in the present context by operating the interface in a controlled system. The SEE sensitivity of the system while it is not running the given interface must be understood before adding the operation of the peripheral. In order to test peripherals we then connect the desired port to an external monitoring system. This external system will be port-dependent and may require software or hardware development to observe nominal and SEE-impacted transmissions. It is possible that loopback interfaces could be used instead of external hardware.

By keeping the peripheral testing separate from the core we orthogonalize a lot of the problems involved in testing the SOC. It is then easy to add the peripheral operation to a running test setup.

Common problems with the approach for testing peripherals presented here include the following. First, since testing the peripheral is not part of the core approach, it is possible that no testing of peripherals would be carried out because the initial test that establishes the operation of the SOC in a radiation environment appears to provide results that appear sufficient for the interested program. The second problem is that the method is most likely to be implemented in a simple system as opposed to a fully functional system. In this case, it is not clear if for a full system the peripheral use would be significantly different and more or less susceptible to radiation effects. Finally, a model approach (with associated model dependence) will be required to extrapolate the test results to the impact in a fully functional system.

4.3 Test Data Approach

Although not formally addressed for this task yet, it is important to discuss the approaches that can be used to convert test data into flight-rate predictions. Once a test method and application are selected, data can be gathered on the target type of structure on the DUT. The gathered data refers to the SEE sensitivity of some portion of the device under some specific operating conditions. In order to apply these results to predict space rates, models must be used to translate the observed sensitivity into a flight application rate.

For example, if we observe a space rate of 1×10^{-2} register SEU/device day, how does this translate into an actual error rate in an executing software program. An example of how this sensitivity could be much different than the estimated rate is if the operating system only allows five local registers to be used (or only compiles to use more than those five in a relatively small percentage of the time). In this case, only five out of 32 registers are expected to cause problems if upset, and then only when they store data (as opposed to providing scratch-pad space). Thus, only about 16% of the register bit upsets in the target

device will lead to flight-application sensitivity and perhaps only half to a third of those bits will be sensitive at any given time, leading to an overall reduction to about 5% of the maximum rate.

5.0 TEST DATA

Once we selected the target SOCs for this year, testing was performed at LBNL for proton SEE and TAMU for heavy ion SEE. These are reported in more detail in [10] and [12]. A brief review of the testing is presented here.

5.1 General Results this Year

During testing this year we were able to establish the viability of both the P2020 and P5020 as test candidates to highlight the goals of the NEPP SOC task, and also as viable devices for both heavy ion and proton testing. The latter is highlighted through the ability to identify viable test hardware and prepare DUTs for exposure. In addition, significant debugging resources exist to enable debugging of test hardware.

5.2 Testing the P2020

The P2020 device was tested during FY11 for both proton and heavy ion SEE sensitivity. Results this year are limited to preliminary test operation verification and limited microprocessor-based findings. Since the P2020 must also utilize the on-chip microprocessor bus and the MMU, these preliminary results provide general evidence of the SEE sensitivity of these other elements of the SOC. This section discusses the test preparations, and proton and heavy ion test efforts conducted this year.

5.2.1 Test Approach

The general test approach for testing of the P2020 is to use the microprocessor to test itself. This approach works best if the test runs are very short, or the information is reported regularly. This approach is outlined in Section 4.2.1.

The actual operation of the DUT was low-level structural testing. The DUT structures targeted for this testing were the SRAM cells used in the caches and the latches used in the microprocessor registers. This data is expected to apply to the majority of the structures on the device, because only the phase-locked loop (PLL), clock tree, and high-speed IO devices are likely to be constructed of other low-level elements. It is possible that the pipeline of the microprocessor and the FPU may be significantly different in radiation response. They likely need to be tested separately in the future.

5.2.2 Test Software Plan and Approach

The P2020 L2 cache is separate from the e500 cores and is a separate functional block, which is independently configurable. One configuration mode it supports is to be mapped directly to the memory manager as on-chip RAM. In this mode it can also be programmed to not use EDAC. The only difficulty in achieving this arrangement is that the translation lookaside buffers (TLB) must be used to configure how the microprocessor accesses the SRAM.

5.2.2.1 Processor-Tests-Itself

As stated, we tested the P2020 using the approach where the processor tests itself. In this system, we use the microprocessor to do the following operations: (1) It configures the processor and SOC configuration registers. (2) It sensitizes all the test targets. (3) It determines the delay and performs and maintenance during the delay. And (4) it collects data from the test targets after the delay is achieved.

5.2.2.2 Test Procedure

The procedure for testing the P2020 follows the diagram in Figure 5.2-1. The main concept behind the test procedure is that the device must function minimally for the full *N*-second delay built into the procedure. For this reason, it is important that the device not have a high probability of crashing during

that wait. For most testing we used N = ~30 seconds. However, for proton testing we needed to extend this to ~300 seconds to reduce periodic downtime when we needed higher fluence exposures.

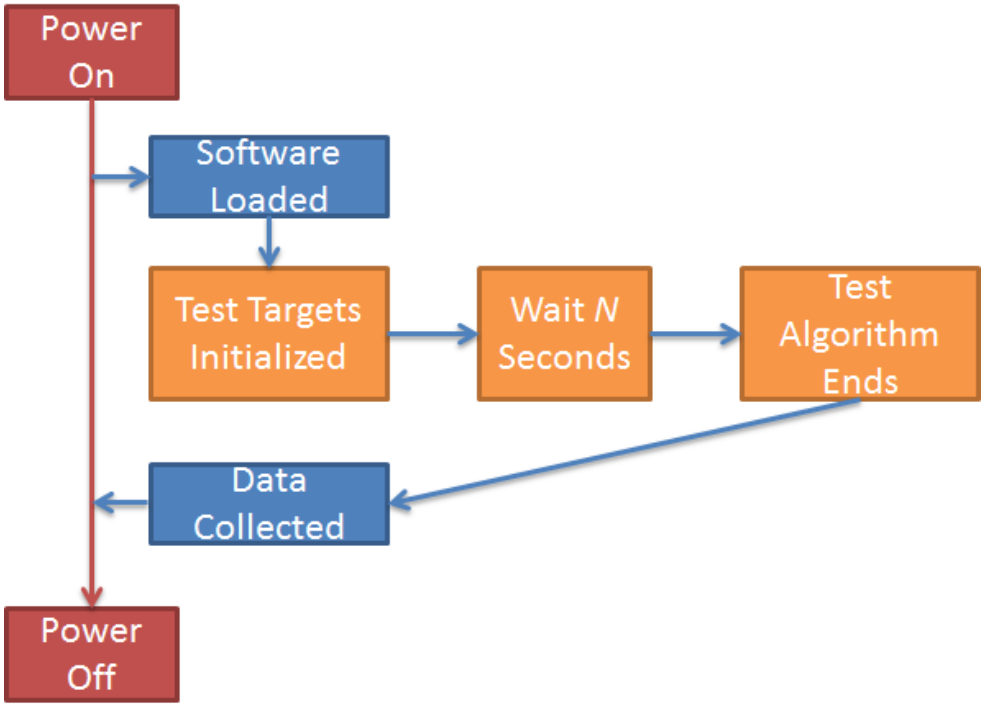

Figure 5.2-1. Test procedure for P2020 testing—delays of 30 to 300 seconds were used.

5.2.2.3 Future Upgrades Needed

During this testing key upgrades for future work were identified. More details on these elements can be found in Section 6.0. The list of identified items is given here.

1. Both e500 cores must be operating.
2. Message passing between cores is needed to test this interface.
3. A more complete MMU evaluation method is needed.
4. Additional work to enable testing of the PCI and/or Ethernet interface is needed.

5.2.3 Proton Testing

5.2.3.1 Exposure Levels

All testing was conducted utilizing a single operational configuration: the DUT was exposed to 55 MeV protons to an exposure level of 2.36e11 protons/cm^2.

5.2.3.2 Key Observations

The primary observations during this testing were the following. The L2 cache readily received SBUs. The L1 cache readily lost cache lines, but did not show single bit upsets. Since the L1 cache is parity protected, some SEU handling was expected; however, the L1 cache parity protection was disabled during testing and upsets caused cache misses instead of exceptions. The microprocessor register bits were not more sensitive than the L1 and L2 caches.

In addition to the observed upsets and associated relative sensitivities, we can say that many structures did not contribute as significant sources of upsets. The memory management unit did not seem to cause upsets, and the coherency unit did not show significant upset chance either.

5.2.3.3 Cache and Register Results

During proton testing the L1 and L2 caches were observed to have SBUs. The results of this testing is presented in Table 5-1. The number of L1 cache upsets was determined by counting the number of invalidated cache lines.

Table 5-1. Results from Proton Testing.

Structure	# of bits	Exposure (#/cm²)	SBUs	Cross-Section (cm²)
GPRs	704	2.36e11	1	0.6e-14
L1 Cache	~278,528*	2.36e11	779	1.2e-14
L2 Cache	4,194,304	2.36e11	10,079	1.0e-14

* - The number of bits in the L1 cache is only approximate because the testing was sensitive to some of the cache tag bits.

5.2.4 Heavy Ion Testing

The P2020 was also tested for heavy ion sensitivity. The level of development in the P2020 test system at the time of this heavy ion testing was not very mature, so the type of data collected is very similar to the limited test type available for the P2020.

5.2.4.1 Device Preparation

The P2020 copper heat spreader is about 300 μm, which is sufficient to block all particle beams provided at TAMU. Because of this, we milled out the heat spreader. The resulting P2020RDB test board was then vulnerable to heat-related test issues. Because the test software did not require a significant portion of the DUT to be activated during testing, this limited the power dissipation and related heating. Testing was performed with a dry nitrogen line providing air circulation, but the gas was not cooled.

5.2.4.2 Ion Exposure Levels

The P2020RDB with milled out heat spreader was exposed to ions at TAMU. The exposure levels, separated by the error type used to sensitize the target structures, is given in Table 5-2.

Table 5-2. Ions and exposures for P2020 at TAMU.

Ion	LET (MeV-cm²/mg)	Error Type	Exposure (#/cm²)
Ne	1.4	‘0’ to ‘1’	1.0e7
Ne	1.4	‘1’ to ‘0’	6.4e6
Ne	2	‘0’ to ‘1’	5.9e5
Ne	2	‘1’ to ‘0’	5.8e5
Ne	4	‘0’ to ‘1’	2.5e6
Ne	4	‘1’ to ‘0’	1.6e6
Ar	5.2	‘0’ to ‘1’	3.0e5
Ar	5.2	‘1’ to ‘0’	1.5e5
Ar	9	‘0’ to ‘1’	1.2e5
Ar	9	‘1’ to ‘0’	1.1e5
Ar	13.3	‘0’ to ‘1’	1.0e5
Ar	13.3	‘1’ to ‘0’	1.2e5

5.2.4.3 Test Results

The P2020 results for L1 and L2 caches are presented here. The data were collected for '0' to '1' and '1' to '0' errors. The results are plotted in Figure 5.2-2.

Both caches show little difference between '0' to '1' and '1' to '0' upset directions. The L1 cache shows slightly higher sensitivity than the L2, but this could be due to an error in the normalization (i.e., it is possible that the number of sensitive bits was not counted correctly).

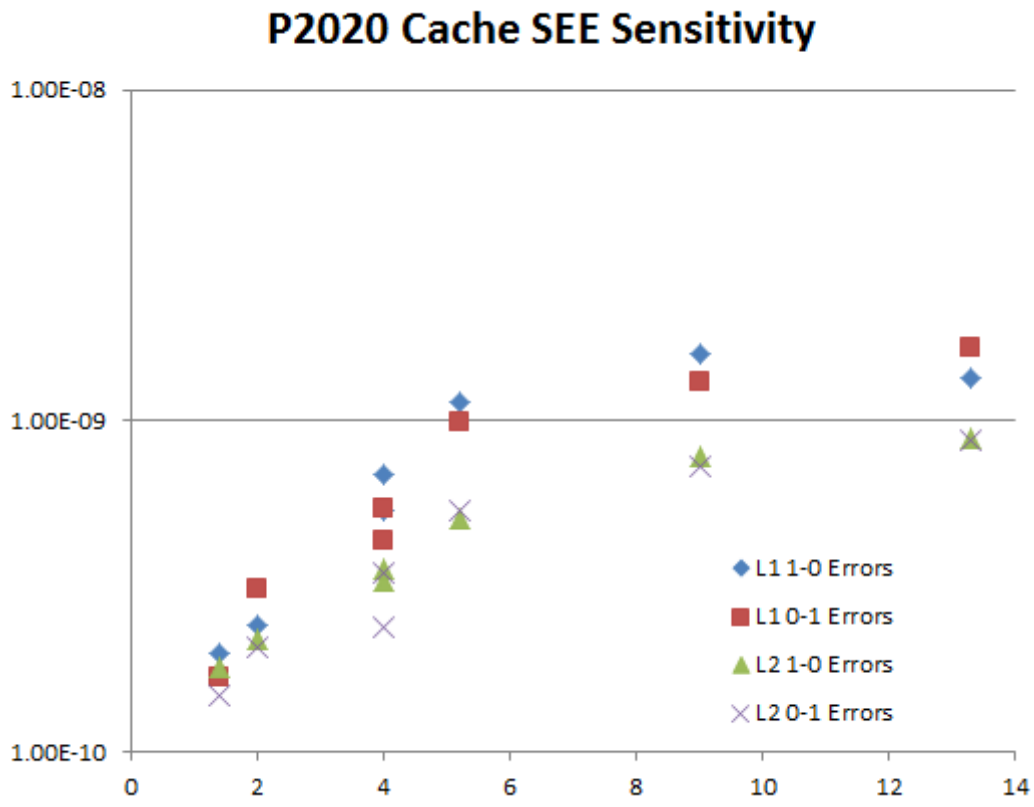

Figure 5.2-2. Results of L1 and L2 cache testing for the P2020.

5.3 Testing the P5020

The P5020 processor was also tested to a limited extent this fiscal year. Some data were collected on the L2 cache. This information is discussed briefly but is also available in a test report submitted to NEPP earlier this year [12].

5.3.1 Device Preparation

The P5020 is required to remain in contact with a heat sink throughout any testing. For this work we exposed the portions of the die required for testing and left the remainder of the heat spreader intact to support heat transfer.

5.3.2 Test Results

Cache results for the P5020 are presented briefly here but may also be found in [12]; see Figure 5-3.

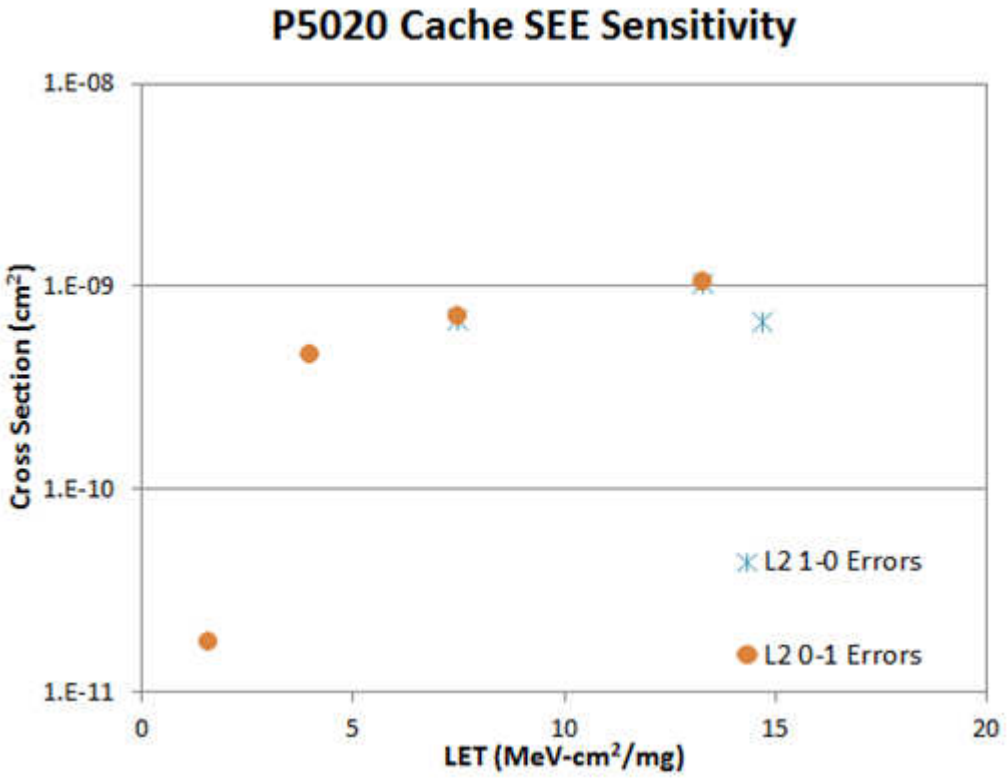

Figure 5-3. P5020 L2 SEE sensitivity, showing most of the data support similar sensitivity for 0 to 1 and 1 to 0 errors.

6.0 FUTURE WORK

FY11 is the second year of this task. The principal goal for this task is to establish a radiation testing and qualification guideline appropriate for SOCs of interest to NASA. As such work will continue to develop understanding of the issues of importance, including the appropriateness of test approaches and the benefit of proposed methods to real NASA programs.

6.1 P5020 Work

The P5020 work started this fiscal year is only an initial effort, and a more complete examination and test regime is needed for this device. For future work we expect to develop or obtain test devices, perform software and hardware development, and perform radiation testing.

6.1.1 Test Devices

JPL is currently pursuing collaborative opportunities for test devices. Efforts this year indicate that P5020 availability and preparation for testing may pose significant challenges in the near future. Work will continue to obtain and prepare devices for radiation test work.

6.1.2 Software and Hardware Development

The P5020 is similar to the P2020. We intend to migrate most of the test software and approach applied to the P2020 for use on the P5020. In this regard, the future work suggested for the P2020 (see below) will apply to the P5020. Unfortunately the translation is not exactly the same, and efforts to directly translate the L2 cache test methods and operation of internal timing systems turned out to not work correctly.

Additional software and hardware development will be necessary to perform the same type of testing on the P5020. Some of this development was performed for the P5020 testing reported here.

6.1.3 P5020 Testing

Additional heavy ion testing of the P5020 is expected under this task in the near future.

6.2 P2020 Work

The P2020 test methods and findings are not yet indicative of use as an SOC. Because of its relative cost, the P2020 is the best candidate for development of Freescale device test methods. Because of this, the P2020 work will be extended in future work in order to perform testing of additional on-chip functional blocks. This work will directly feed into review of the test methods for SOC devices and determination of the applicability of those methods to real user applications.

6.2.1 P2020 Test Upgrades

A major thrust of work in FY12 will be development of the existing P2020 test hardware and software in order to enable testing of more of the resources on the P2020. As indicated in Section 5.2.2, a series of upgrades is needed.

6.2.1.1 Software Work

Key software upgrades planned for the P2020 include the following:

1. Enable multi-core testing
2. Sensitize the “coherent system bus”
3. Write software to support key elements of hardware for testing.

6.2.1.2 Hardware Work

Hardware work must be continued to provide the support necessary for testing key peripherals. Some of the most appropriate peripherals are the following:

1. MMU
2. Ethernet.

6.2.2 P2020 Testing

Testing of the P2020 will continue in the future to explore the upgrades discussed. This testing will verify the applicability of the upgrades and provide insight into additional limitations of the testing and qualification approach when applied to more of the SOC structure than was tested this year.

6.3 Future SOCs

Under the current work there is significant interest in testing of SOCs and sufficient material to provide a sound foundation for the guideline to be produced in a later report from this task. However, significant additional work should be conducted and in future work.

6.3.1 Atmel AT7913

This device would provide excellent counterpoint to the efforts to test the UT699 performed under this task. The Atmel AT7913 is a LEON 2FT. It would provide a platform to validate findings from the Aeroflex UT699.

6.3.2 SiLabs C80501F310x

This device is a full SOC and its use would enable this task to take first steps to help establish what to do in the case of devices that are much more like a traditional SOC than the majority of devices discussed under this work.

6.4 Maestro Efforts

Although NEPP is not directly involved in continued work on Maestro, work on this device clearly falls into the scope of this task. Limited additional radiation evaluation of Maestro is expected. Where possible, available results will be incorporated into the material researched here.

6.4.1 JPL Testing

During FY11, JPL performed testing of Maestro for NASA's Enabling Technology Development and Demonstration (ETDD) program. The results are currently under review. It is expected that much of this information will be available for inclusion into this task in FY12.

6.4.2 Other Testing

Some other groups may be working on radiation testing of Maestro. Where possible, this task will continue to monitor released radiation test results and include any developed knowledge related to test methods and results.

6.5 Collaborative Review

As mentioned in the seven-element approach, collaborative efforts are very important to this effort. Review of efforts is a key part of the collaborative needs. In particular, since application of SOC test data is necessarily model-dependent, it makes the most sense to verify that the model dependencies are limited

in scope by verifying that the findings here are relevant to actual users. The two main areas where a direct review will benefit the task greatly are the development of test methods, and the extension of test data to actual system impact.

7.0 REFERENCES

[1] Guertin, S. M., Irom, F., "Recent Results for PowerPC Processor and Bridge Chip Testing", *IEEE Radiation Effects Data Workshop*. 2010.

[2] Manet, P., Falmagne, S., Garnier, J., Berger, G., Legat, J.-D., "SEE evaluation of a low-power 1µm-SOI 80C51 for extremely harsh environments", *Proceedings of the 2009 Conference on Radiation and Its Effects on Components and Systems (RADECS)*. 2009.

[3] Michalak, S. E., DuBois, A. J., Storlie, C. B., Quinn, H. M., Rust, W. N., DuBois, D. H., Modl, D. G., Manuzzato, A., Blanchard, S. P., "Neutron Beam Testing of High Performance Computing Hardware", *IEEE Radiation Effects Data Workshop*. 2011.

[4] Lintz, J. P., Hoffmann, L. F., Smith, M. J., Van Cleave, R. T., Cizmarik, R. R., "Single event effects hardening and characterization of Honeywell's pass 3 RHPPC processor integrated circuit", *IEEE Radiation Effects Data Workshop*. 2007.

[5] Guertin, S. M., Hafer, C., Griffith, S., "Investigation of Low Cross Section Events in the RHBD/FT UT699 Leon 3FT", *IEEE Radiation Effects Data Workshop*. 2010.

[6] Allen, G. R., Swift, G. M., Miller, G., "Upset Characterization and Test Methodology of the PowerPC405 Hard-Core Processor Embedded in Xilinx Field Programmable Gate Arrays", *IEEE Radiation Effects Data Workshop*. 2007.

[7] LaBel, K. A., Ladbury, R. L., Cohn, L. M., Oldham, T. R., "Radiation Test Challenges for Scaled Commercial Memories", *IEEE Transactions on Nuclear Science*. 2007.

[8] Allen, G. R. "FPGAs—Working with a Commercial Consortium: Tales of Xilinx Virtex-4 and SIRF Radiation Testing", *NEPP Electronics Technology Workshop*. 2010.

[9] Hafer, C., Griffith, S., Guertin, S., Nagy, J., Sievert, F., Gaisler, J., Habinc, S., "LEON 3FT Processor Radiation Effects Data", *Radiation Effects Data Workshop*. 2009.

[10] Test Report on P2020 Proton Testing from March 2011, NEPP Report.

[11] Irom, F. "Guideline for Ground Radiation Testing of Microprocessors in the Space Radiation Environment", *JPL Publication 08-13*. 2008.

[12] Test Report on P2020 and P5020 Heavy Ion Testing from September 2011, NEPP Report.

[13] Guertin, S. and F. Irom, "Processor SEE Test Design", *Single Event Effects Symposium*. 2009.

[14] Wikipedia article on http://en.wikipedia.org/wiki/System_on_a_chip, (03/2012).

[15] Koga, R., Kolasinski, W. A., Marra, M. T., Hanna, W. A., "Techniques of Microprocessor Testing and SEU-Rate Prediction", *IEEE Trans. Nucl. Sci.*, 32(6), 4219–24. 1985.

APPENDIX A. ACRONYMS AND ABBREVIATIONS

AMFRT	Architecture for Microprocessor Functional Radiation Testing
ASIC	Application Specific Integrated Circuit
LBNL	Lawrence Berkeley National Laboratory
BTK	Board Test Kit
CCR	Cache Condition Register
DUT	device under test
EDAC	error detection and correction
FPU	Floating Point Unit
ETDD	Enabling Technology Development and Demonstration (NASA Program)
FT	fault tolerant
FY	fiscal year
GPR	General Purpose Register
IC	Integerated Circuit
I/O (IO)	input/output
IUCF	Indiana University Cyclotron Facility (see also ISAT)
ISAT	The Integrated Science and Accelerator Technology Hall (formerly IUCF)
JPL	Jet Propulsion Laboratory
JTAG	Joint Test Action Group (also refers to hardware used to operate a port defined under the group)
LET	linear energy transfer
MMU	memory management unit
NASA	National Aeronautics and Space Administration
NEPP	NASA Electronic Parts and Packaging Program
NSRL	NASA Space Radiation Laboratory
PLL	Phase-Locked Loop
RHBD	radiation hardened by design
SEE	single-event effects
SOC	system on a chip
TAMU	Texas A&M University Cyclotron
UART	Universal Asynchronous Receiver/Transmitter
TLB	Translation Lookaside Buffer