53601

Dr. Rod Barto
3312 Moonlight
El Paso, Texas 79904
915-755-4744
email: rod.barto@worldnet.att.net

April 30, 2001

Report on Mars Odyssey Independent Assessment Team Activities

Background and Chronology of Events

During the week of March 26, 2001, I was asked by Rich Katz, NASA-GSFC, to participate on the Mars Odyssey Independent Assessment Team (IAT) that would investigate the implications of the failure of an Actel RP1280 FPGA, which occurred on the SIRTF spacecraft, on the Mars Odyssey spacecraft that was set to launch on April 7, 2001. The other members of the IAT were:

Rich Katz, NASA-GSFC,
Don Mayer, Aerospace Corporation,
Jon Osborne, Aerospace Corporation,
Jerry Soden, Sandia National Laboratory.

We were provided with review materials from JPL and Lockheed Martin (LMA) that would be discussed at a meeting on April 2, 2001.

On Monday, April 2, the IAT met with JPL and LMA personnel at the LMA facility in Denver, Colorado. In the discussions, the IAT realized that there had been two FPGA failures, rather than one (this had also been suggested in the review materials), and that the first failure occurred in the Mars Odyssey system some 12 to 18 months prior to the SIRTF failure. Both failed FPGAs came from a wafer lot labeled U1H466 made on the Matsushita fab line for Actel, packaged at SEI in San Diego, California, tested by Actel, and subsequently shipped to LMA for programming. These parts were the first lot of parts delivered to customers using this flow.

The discussions revealed that none of the failure investigation methods pursued yielded any cause for either failure. However, the IAT learned that the built-in Actel FPGA internal probing mechanisms, which might have given some idea as to the cause of the failures, had not been used. Both of the failed FPGAs exhibited high current draw, with only the SIRTF part having failed functionally. On the Problem/Failure Report (P/FR) for the Odyssey part, the failure was considered to be "normal fallout". In investigating the failures, it was apparently thought, by LMA and/or JPL, that the high current draw was the result of a failure at a specific location on the FPGA die, so both dies were subjected to a liquid crystal test that would show hot spots. In neither case were the results conclusive, and the SIRTF part was subjected to further deprocessing which revealed the presence of a particle contaminant that was proposed by LMA as the failure cause. This proposal was withdrawn on the day of the meeting. The testing was totally destructive for the SIRTF part and partially destructive for the Odyssey part, so the internal FPGA probing mechanisms could no longer be used.

Even though none of the investigation methods pursued yielded any failure cause, several failure causes were proposed and discussed in the meeting.

- Tien Nguyen, JPL, presented the results of his experiments into electrostatic discharge (ESD) and electrical overstress (EOS) effects on Actel FPGAs. His results appeared to be inconclusive for the following reasons:

  - His finding that only a programmed part was susceptible to ESD/EOS was contradicted by his not knowing that the tests he performed did not actually show that the parts were functional;
  - He apparently chose ESD testing over EOS testing because of the equipment available, so it is not clear what EOS testing was done.

  Never the less, the conclusion of Mr. Nguyen's report was that ESD/EOS was the most probable cause of failure, even though there was no evidence seen in the SIRTF part of ESD or EOS damage, based on the symptoms seen in the failed parts being the same as the symptoms induced in the lab.

- Since the search for hot spots on the dies (sites where the excess currents were passed) were inconclusive, failures of a more global nature were proposed. Predominant among these was the possibility that one or more of the charge pumps in the FPGA had failed. Such a failure could cause high current draw and functional failure. However, no tests were run to determine whether the charge pumps were running after the failure and before the deprocessing.

- It was speculated that the testing done on this part lot did not give the expected reliability demonstration, i.e., that while it was believed that the testing done to qualify the part placed the lot in the flat part of the reliability curve, in fact the qualification had not actually been completed, and the lot was still in the initial downward slope of the curve in which infant mortality failures could be expected. By this theory, the failures seen were part of the drop-out that would be seen in the qualification process. This appears to be contradicted by 22 parts from this lot having passed a 1000 hour life test at Actel.

- It was noted that the life test parts had been programmed at Actel, while the failed parts had been programmed at LMA. It was speculated that if ESD or EOS damage was the failure cause, that the LMA handling procedures were inadequate. It was further noted that, while the Actel published failure rate for their FPGAs is 9 FITs, LMA reports 600 FITs for the parts, and that furthermore the LMA failure rate holds not only for parts from the lot in question, but also for parts in their Titan rocket program which uses Actel FPGAs from other lots and flows.

Based on concerns about LMAs' handling procedures, and the possibility that EOS could have arisen from higher voltages on the PC board on which the failed part had been installed, the IAT requested to see the LMA programming area and the SIRTF and Odyssey schematics. These requests started in the late morning and continued until mid-afternoon, at which time we were given a tour of the programming area. During the tour, the grounding of the programmer was verified, and it was determined that the ground lug on the programming module was not tied to ground, but that otherwise the grounding seemed to meet requirements. The LMA employee in charge of the programmer could not remember when the programmer calibration test had last been run, but stated that it probably was a year or more prior to the meeting. The delay in seeing the programmer apparently resulted from this employee having to go off site to obtain the software for the programmer that he expected we would want to see. As to the schematics, however, we were told that they were considered LMA proprietary, and it would have to be determined whether we could see them.

After the discussion, the meeting was closed so that the IAT could discuss their findings. By this time, Jerry Soden had left for home, and was not a party to the IAT discussions, nor some of the later portions of the meeting after the tour. The IAT discussions began by Rich Katz asking each member of the team if they would sign a "consent to launch" form, by which was meant, would we put our names on the line as saying that the Odyssey system was fit to launch. No member would agree to this, and the discussion resulted in the following findings being released to LMA and JPL:

1. Two parts from the flight lot have failed at Lockheed Martin.
2. No credible cause for either failure has been identified.
3. JPL and Lockheed Martin do not agree on even speculative causes of failure.
4. Lockheed Martin has not shown electronics reliability consistent with a high-cost, high-profile spacecraft for the mission lifetime.
5. The 600 FIT rate (which is considered optimistic) for the Actel FPGAs at 60 °C predicts approximately 2 failures over the mission lifetime.
6. Based on solid engineering principles, no member of the IAT will sign a "Consent to Launch" form. A launch decision must be made by management.

After releasing the findings, Rich Katz and I informed LMA and JPL that we wished to return on Tuesday, April 3, to review the schematics and run the programmer certification test, to which LMA agreed. Messrs. Osborne and Mayer returned home on Monday evening.

At about 9:00 AM on Tuesday, April 3, Rich Katz and I arrived at LMA, and from the cafeteria in the SSB (the building in which the IAT meeting had been held) called various LMA personnel, restating our request to see the schematics and verify the programmer. We were told that our requests would be granted, and that we should wait in the cafeteria for further instructions. We were told that an Actel Field Application Engineer (FAE) would be on site at about 12:00 to run the calibration test on the programmer and that we were welcome to view the procedure. We waited, having sporadic contact with LMA personnel, and shortly before noon were met by Jeff Wetch,

Actel FAE, who stated that he had arrived earlier that morning to run the calibration test. He told us that the programmer passed the test, and offered his assessment that the LMA ESD procedures were in the top 10% of facilities he had reviewed. We continued waiting to see the schematics until about 2:30, at which point we left for our hotel to begin writing our report.

On Wednesday, April 4, Rich Katz and I stayed in our hotel writing our report, while letting LMA know that we were still available to see the schematics, if it was possible. At the end of the day, we had still not seen the schematics, but did release a report and left for our homes.

At some point after the Monday meeting, the liquid crystal photographs of the Odyssey failure, that had been made perhaps a year before and at the time were considered to be inconclusive, were reexamined. It was now determined that the photographs showed a high current path between power and ground, so that ESD or EOS was a possible failure cause. It is not clear why the same determination was not made at the time the photographs were taken. Also, the damage was in the core of the IC and not in the I/O area, which is inconsistent with Actel experience.

On Friday, April 6, I received an electronic copy of the schematics at about 7:30 AM MST. Mr. Katz and I performed independent reviews and noted numerous instances of poor design techniques. Neither of us considered the design to be space flight quality. At about 11:00 AM MST we began a teleconference with all the members of the IAT to finalize our report. Don Mayer, chairman of the IAT, released a draft report that contained several speculations as to what the failure causes for the two parts might have been and assigned relative probabilities to them. Neither Mr. Katz nor I believed that the technical evidence presented offered a basis for any of the causes, and that it was therefore specious to say that any of them were more probable than any of the others. Messrs. Osborne and Soden left the discussion late in the afternoon, and the telecon continued between Mr. Mayer, Mr. Katz, and myself until 10:30 PM MST. By that time, Mr. Katz and I had been largely successful in stripping speculations from the report, leaving mostly only that which we knew based on technical evidence.

The decision to launch Mars Odyssey was made by NASA, and it was launched on April 7.

As a postscript, on April 20, I was informed by Mr. Katz that a third part from the lot in question had failed on the SORCE program at LASP, University of Colorado. The failure occurred on March 7, 2000. The PFR stated that the reset input signal on the part had been released, but that the clock circuits it implemented were not running, i.e., held in reset.

Discussion

Although many speculative failure conjectures were put forth, in the absence of any technical evidence to support them, they remain conjectures, not facts. When examining the theories and evidence, it seems that each theory is contradicted by some piece of evidence. A table discussing the theories is given below.

Failure Cause	Contradicted by
Programmer error at LMA: If the programmer sequence resulted in voltage spikes during programming, an ESD or EOS event might have propagated into the core electronics.	• The certification test run by Jeff Wetch, Actel FAE • The third failure at LASP
ESD/EOS during handling outside LMA: If the parts were ESD damaged during handling at SEI, Actel, or other location, it may have created a latent defect that would worsen with time.	The life test was run at Actel on parts programmed at Actel, with no failures.
ESD/EOS during handling at LMA: If the parts were ESD damaged during handling, for example in the receiving, programming or test areas, it may have created a latent defect that would worsen with time.	• Mr. Wetch's assessment of LMAs' ESD procedures. • The third failure at LASP
Manufacturing defect (random)	Not likely for 3 unrelated failures
Manufacturing defect (systemic): An error during processing may have affected all devices on a wafer or wafer lot. If most defects were found, but some escaped, and worsened with time, it might generate the observed symptoms.	The life test at Actel
Board design error: It is possible that some error in board design may cause damage to the FPGAs, for example by inappropriate sequencing of supply voltages during power up. This might cause EOS problems that could worsen with time or with each repetitive operation.	Not likely for 3 unrelated failures
FPGA circuit error: It is conceivable that some circuit design problem could cause a latent error, perhaps by overstressing some area of the circuit during power up, for example. This would be design dependent. Related to board design error issue.	Not likely for 3 unrelated failures
Inadequate qualification program: Additional testing may show these 3 failures to be simply infant mortality	The life test at Actel

Conclusion and Recommendations

Regardless of how improbable any of the failure causes are, or how self-contradictory the evidence is, the fact remains that 3 parts from a very small population failed with no apparent cause. While this clearly disparages their use, it is not clear where in the process flow the culprit is hiding. It seems reasonable, therefore, to take more careful measures in dealing with these parts, and parts like them, if their continued use is desired. The following are recommended:

- Parts from this lot should be subjected to extensive life tests, in an attempt to push them to failure. The failures should be carefully examined and resolved.
- The philosophy behind qualification and burn-in tests should be examined. The testing required for an upgraded commercial part might be different from testing required on a part made on a class S line.
- NASA should more closely monitor contractors to assure correct FPGA usage and handling, and make resolving FPGA failures a high priority.
- NASA should more closely monitor FPGA fabrication, packaging, and testing programs.

Dr. Rod Barto 3312 Moonlight El Paso, Texas 79904
915-755-4744 rod.barto@worldnet.att.net