STEREO IMPACT

PR Numbers: 1xxx=UCB, 2xxx=Caltech/JPL, 3xxx=UMd, 4xxx=GSFC/SEP, 5xxx=GSFC/Mag, 6xxx=CESR, 7xxx=Keil, 8xxx=ESTEC, 9xxx=MPAe

Assembly: SEP	**SubAssembly: LET**
Component/Part Number:	**Serial Number: FM2**
Originator: Branislav Kecman	**Organization: Caltech**
Phone: (626) 395-4264	**Email: kecman@srl.caltech.edu**

Failure Occurred During (Check one √)

□ Functional test √ Qualification test □ S/C Integration □ Launch operations

Environment when failure occurred:

□ Ambient □ Vibration □ Shock □ Acoustic
□ Thermal □ Vacuum √ Thermal-Vacuum □ EMI/EMC

Problem Description

During thermal vac testing (June/July 2005) near -1 ºC, always as the temperature was rising from a cold soak, LET microprocessor would crash. The problem was eventually traced to LET memory corruption due to a particular bit flip on LET SRAM's data buss.

Analyses Performed to Determine Cause

Since the failure was reproducible only inside the vacuum chamber, analysis of schematics and firmware code had to be done to try to isolate its cause between the hardware and the firmware. Although there are hypotheses, the ultimate cause of the problem still remains unknown. The leading theory is that one of four 0.1 uF by-pass capacitors mounted on the underbelly of SRAM device gets somehow disconnected at low temperature, loses its charge through leakage, and then it reconnects as the temperature goes up, thus creating a momentary glitch on the +3.3V supply line as it sucks the current in order to charge itself up.

During T/V test we had additional diagnostic firmware loaded into the LET microprocessor to investigate potential causes in both hardware and firmware. It took a while to determine the exact thermal and other operating conditions under which the problem occurred. To reliably capture the failure the instrument temperature was typically set to -5 ºC and dwell time to 30 minutes. Then the temperature is slowly raised to +5 ºC, and if the bit flipped during the transition (typically at -1 ºC), a core dump would be obtained for LET memory, and then compared with the "normal" memory core dump. The process was repeated number of times in order to catch the bit-flip in the act (later on, it was also used to verify that the problem has been solved once the suspect SRAM part was replaced).

7/24/2005 e-mail message (bellow) from Rick Cook summarizes the analyses performed:

On Friday we agreed on a conference call to prepare for replacing both the LET MISC ACTEL and the one memory IC in which the corrupted bit occur. A new UMC ACTEL was programmed on Saturday thanks to Igor and Shane. (8/31/2005 NOTE: As it turned out, there was no need to replace the ACTEL part.)

Since it is extremely important that we not replace parts and later discover we still have the same problem and it is s/w after all, we held a review at Caltech on Friday afternoon. I presented my case for why the problem can't be s/w to Alan Cummings, Ed Stone, Dick Mewaldt and Rick Leske. The weakest part of the argument seemed to be my exoneration of the event interrupt and processing which occurs asynchronously. Rick Leske suggested that I check in the core dumps to see if there was evidence of a pending event (i.e., one that had its data transferred into the MISC by interrupt and was awaiting processing by a foreground routine). I did that and found that in every case there was no such pending event.

However, to be more certain that a s/w bug is not at the root of the problem, we performed an experiment in which we turned off all the routine LET processing except a single interrupt (that occurs only once per second), the clock gating s/w, and the short diagnostic routine that detects bit-flips and halts the processor.

We then, with this much-reduced set of code actually running (but all the LET code loaded in memory), we took the system up as usual through the -1 ºC region. (If a s/w bug is causing the bit flips it was known that the synchronous once per second interrupt did not contain it -- due to the apparently random time of bit flip first detections relative to the once per second pulse.)

We did catch a bit flip, the processor halted, and we obtained a core dump. The typical pattern of bit flips at the typical addresses was seen in the core dump. In this test the processor clock was entirely stopped for 99.99+ % of the time since it would "wake up" only once per second in response to the once per second interrupt and execute for only a few 10's of microseconds before turning the clock back off. So one possibility is that the memory corruption actually occurs when none of the memory access signals are active and the clock is stopped.

Over night we have been trying to go even one step further. We have been doing thermal cycles wherein we pass through the critical -1 ºC region with absolutely no processor activity occurring. To do this we turn off all the foreground and interrupt processing, except that which services the command interface. In this case the only time the processor clock runs is following the receipt or transmission of characters on the command interface. By simply not commanding, we ensure that the clock remains stopped. (I verified that indeed this mode works and produces an entirely stopped clock using our EM setup at Caltech.)

It appears that we did catch two bit flip episodes this way, with the clock entirely stopped during the transition through the critical temp range. I was able to verify bit flips in the image obtained from the first episode, but haven't yet seen the most recent episode's core dump. However, I was surprised that the number of bit flips present in the first image was much fewer than typical -- in fact only three.
(I need to spend some more time today studying this image and the later one.)

I had a conversation with someone at Honeywell on Friday regarding the HLX6228 memory devices. I gave him the corrupted bit pattern and address pattern info and he is going to study it with their engineers and report back to me in a phone call scheduled for 11 AM pacific on Monday. He suggested that I review the timing and margins on our memory access signals and I have done this using our EM unit. (There is plenty of margin -- all very conservative.) I also measured the DC state of all the memory access signals in the "clock stopped" state (in which the memory chips are not selected).

I was able to examine a rough block diagram of the memory devices and it appears that the pattern of addresses at which bit flips occurs bears some relationship to the set of address lines which enter the "column" as opposed to "row" address decoders.

The bottom line is that the cause of the bit flips remains a mystery. However, if the bit flipping actually can occur with the MISC clock stopped and the memory access signals static (and memory actually disabled) then it is hard to see how a bad ACTEL can be the root cause.

The open-circuiting bypass cap theory still seems plausible and is the only one so far that attempts to explain the odd temp behavior.

I suggest we hold off on replacing parts until after my discussion with the Honeywell engineers (11AM Monday, Pacific time).

Corrective Action/ Resolution

√ Rework □ Repair □ Use As Is □ Scrap

1. 7/28/2005 Replaced the offending SRAM device U7 (Honeywell P/N HLX6228TSR, S/N 00013, LDC 0314) on FM2 LET Board with S/N 00028 from the same lot, but without the four by-pass capacitors.
2. FM2 LET passed re-qualification T/V test between 7/29/2005 – 7/30/2005.

3. The failed device was sent to Honeywell for failure analysis, along with the one from FM1 SEP Central (PFR 2014). The failure analysis is still pending. So far, they have not been able to reproduce the problem and determine the cause of malfunction.
4. In case of a negative result from the failure analysis, the investigation will continue at Caltech. We will install the failed device on the EM LET Board and instrument it for the test in thermal chamber, possibly with support from Honeywell engineers.

Update 6/2/2006 LR:
Although the failure could easily be replicated at Caltech, Honeywell was unable to reproduce the same thermal conditions to recreate the original failure. There was a schematic review between Caltech and Honeywell in an attempt to determine the root cause prior to any destructive analysis. The suspect device (from PFR 2013) was then delidded and Honeywell found a possible correlation between the failed locations and the spare fuse circuits. Although, at this time, the root cause of the failure is still unknown. This issue will be forwarded to the GSFC parts group for possible further analysis with Honeywell.

Since the SRAM devices were replaced, both FM1 and FM2 SEP Central/LET/HET have successfully passed instrument level and observatory level vibration and thermal vacuum testing. In addition, due to other rework, both instrument assemblies successfully passed an additional single axis vibration and four additional thermal vacuum cycles in May 2006. FM2 also saw one more additional single axis vibration. All flight units have been working nominally since the original repair in July 2005.

Date Action Taken: 7/28/2005 **Retest Results**: Passed
Corrective Action Required/Performed on other Units ☐ Serial Number(s): N/A

Closure Approvals

Subsystem Lead:	Branislav Kecman	Date: 8/31/2005
IMPACT Project Manager:	_____	Date:_____
IMPACT QA:	_____	Date:_____
NASA IMPACT Instrument Manager:	_____	Date:_____