TEST AND VERIFICATION PLAN

Proposal for an ESS (Environmental Stress Screening) Program for the EVLA Correlator

TVP Document: **A25010N0006**

Revision: DRAFT

B. Carlson, June 6, 2006

National Research Council Canada
Herzberg Institute of Astrophysics
Dominion Radio Astrophysical Observatory

P.O. Box 248, 717 White Lake Rd
Penticton, B.C., Canada
V2A 6J9

Table of Contents

1 REVISION HISTORY 5

2 INTRODUCTION 6

3 BACKGROUND 7

4 CORRELATOR CHIP TEST 10

4.1 TEST #1—ROOM TEMPERATURE FUNCTIONAL TEST 11
4.2 TEST #2—TEMPERATURE CYCLING 11
4.3 TEST #3—ROOM TEMPERATURE ELECTRICAL TEST FOR 1 HOUR 14
4.4 TEST #4—HIGH TEMPERATURE ELECTRICAL STRESS TEST 14
4.5 TEST #5—POWERED TEMPERATURE CYCLE 15
4.6 TEST EFFECTIVENESS EVALUATION 15
4.7 LIFETIME DEGRADATION EVALUATION 15
4.8 TEST BATCH SIZE 16
4.9 ALTERNATIVE TEST STRATEGY 16
4.10 CORRELATOR CHIP PROTOTYPE TESTING 16
4.11 PRODUCTION AUDIT TESTING: FPGA MFG TESTS 17

5 BOARD TEST 18

5.1 TEST #1—ROOM TEMPERATURE FUNCTIONAL TEST 18
5.2 TEST #2—TEMPERATURE CYCLING AND MECHANICAL SHOCK 18
5.3 TEST #3—ROOM TEMPERATURE ELECTRICAL TEST FOR 1 HOUR 19
5.4 TEST #4—HIGH TEMPERATURE ELECTRICAL TEST 19
5.5 TEST #5—POWERED TEMPERATURE CYCLE 20
5.6 TEST EFFECTIVENESS EVALUATION 20
5.7 LIFETIME DEGRADATION EVALUATION 21
5.8 TEST BATCH SIZE 21
5.9 BOARD PROTOTYPE TESTING 21

6 RACK TEST 22

7 REFERENCES 24

List of Figures

FIGURE 4-1 INDUSTRAL FREEZER WITH –85 ^{O}C CAPABILITY. 11
FIGURE 4-2 INDUSTRIAL OVEN WITH +300 ^{O}C CAPABILITY. 12
FIGURE 4-3 MIL-STD-883 METHOD 1010 TEMPERATURE CYCLING CURVE. 13
FIGURE 6-1 RACK SHAKER TABLE TO SIMULATE GROUND TRANSPORT MECHANICAL MOTION. 23

List of Abbreviations and Acronyms

ALMA – Atacama Large Millimeter Array. A correlator for ALMA has similar speed and bandwidth requirements as the EVLA Correlator.

ASIC – Application Specific Integrated Circuit.

BGA – Ball Grid Array.

DRAO – Dominion Radio Astrophysical Observatory, Penticton, BC, Canada.

ESS – Environmental Stress Screening. "The goal of ESS is to cause defects resulting from manufacturing mistakes to become detectable failures without doing damage to otherwise good material." [Hnatek, 1995]. Screens are executed on 100% of the product.

EVLA – Expanded Very Large Array.

HALT – Highly Accelerated Life Test.

HASS – Highly Accelerate Stress Screen.

NRAO – National Radio Astronomy Observatory.

UUT – Unit Under Test.

PCB – Printed Circuit Board.

PQFP – Plastic Quad Flat Pack.

TSMC – Taiwan Semiconductor Manufacturing Corp. The manufacturer of the custom correlator chip ASIC.

ZIF – Zero Insertion Force.

1 Revision History

Revision	Date	Changes/Notes	Author
DRAFT	June 6, 2006	Initial DRAFT release for review	B. Carlson

2 Introduction

This document outlines a proposed ESS (Environmental Stress Screening) program for the EVLA Correlator.

The purpose of an ESS program is to detect fabrication and manufacturing defects—failures that are part of the early portion of the bathtub curve—to improve installed system reliability, thereby reducing overall costs. It is important to note that ESS does not make products more reliable, rather it makes the population of products more reliable by weeding out defects in the early going.

[Davis & Davis, 1989] found that the early portion (infant mortality portion) of the bathtub curve (failure rate vs time) follows the equation:

$$e^{-t/MTBF}$$

For example, for the Baseline Board with a calculated MTBF of 56,000 hours, if it is run under normal operating conditions for 100 hours before shipment, we can expect to see 99.8% of the failures in the field (i.e. only 0.2% of the failures that are going to occur are detected before shipment). 100 hours of run time would seem not to be sufficient to weed out failures, if failures are purely due to normal wear-out mechanisms.

[Scheiber, 2001] notes that it is ironic that more reliable units with higher MTBFs have shallower curves and therefore it is harder to detect failures before shipment. Makes sense—reliable products don't fail as often and therefore it is harder to find the failures in the early going.

Conversely, with unreliable products, the curve is steeper, and therefore it is easier to detect failures. This means that manufacturing defects, if taken into account in the MTBF calculation, should be relatively easy to find if the test conditions are designed properly. For example, if the true MTBF of a particular Baseline Board is 5,000 hours (1/2 year) due to one or more hidden PCB-level or chip-level defects, then running the board during an ESS screen for an equivalent 10,000 hours means that 86.5% of the defects should be found before it is shipped. Conversely, running the board for 100 hours under normal operating conditions (i.e. no ESS) finds only 2% of the failures—not very effective.

This document defines ESS tests for the correlator chip ASIC, for boards, and for racks. It is hoped that the tests proposed in this document find a high percentage of defects before shipment to the VLA site, and are cost effective and appropriate for this project.

3 Background

We have no formal experience with *effective* ESS testing for custom ASICs or for circuit boards with BGA devices of the complexity and technology employed in the EVLA Correlator. Burn-in testing has been used for past projects, but a survey of the literature obtained through the CISTI library indicates that traditional burn-in is not effective [Scheiber, 2001] and, if not handled properly, can be damaging and induce many more defects than it finds [Pecht, 1999].

There is some conflicting information, but it seems that the most effective single screen in the least amount of time is temperature cycling with rapid (>40 ^{o}C/min) air-transfer temperature ramps with devices not under power. There is a curve in MIL-STD 883, Method 1010, which shows cycles for integrated circuits going from –65 ^{o}C to +150 ^{o}C, within a few minutes, and with a minimum 10-minute dwell time. For boards, the JEDEC standard JESD22-A105C, shows slower temperature cycles with concurrent power cycles, 5 min ON, 5 min OFF, ranging over a less extreme temperature from –40 ^{o}C to +80 ^{o}C, with 15 minute temperature ramps. Although, this JEDEC standard is for circuits that would normally see temperature and power cycles, and the standard indicates it is a destructive test.

Multi-axis, broad-spectrum vibration is also listed as an effective screen but is relatively expensive. It can be replaced by a less effective [Scheiber, 2001] mechanical shock test (drop test), but the induced acceleration impulse must be quantified to ensure effectiveness and to ensure that undue damage is not done to the UUT. I can imagine that a quantifiable, easy and cheap to build, and easy-to-use drop tester could be built.

The goal of an ESS program is to find silicon chip, PCB, and PCB assembly defects that may not show up for several months or years in the actual system. Correlator boards will live their lives in a relatively “cushy” existence, however, defects may eventually make themselves known and result in a perception of unreliability, be annoying, and eventually be expensive or impossible to fix due to lack of spares. Thus, it is not necessary to employ tests on production units that mimic extremes that the units will never see in their lifetime—it is likely necessary to employ tests to find defects that may eventually show up.

The best test strategy is likely HALT (Highly Accelerated Life Testing) followed by HASS (Highly Accelerated Stress Screening). There seems to be some controversy, and the biggest proponents of HALT/HASS testing seem to be companies that sell such equipment, but this strategy seems to be the best. HALT establishes the survival limits of the product, and HASS is a full production screen that uses the results of HALT to establish test limits.

HALT is also used to find and mitigate module design defects through iterative testing and re-design, and is most effective for those products that will see mechanical and environmental extremes in their lifetimes (portable test equipment, lap-tops, cell phones etc.). Both of these are not likely necessary for the EVLA Correlator since all of the

modules are rack-mount and significant work is already being done to ensure robust mechanical design for the required environment.

Manufacturers quantify the cost-effectiveness of HALT/HASS with a cost-benefit analysis. For widgets that are sold, HALT/HASS can result in decreased warranty service costs and improved customer satisfaction; a quantifiable analysis can be performed, and refined with field data over time. While both of these apply somewhat to the EVLA Correlator, they are hard to quantify and therefore it is hard to justify a large investment in HALT/HASS testing.

HALT/HASS testing is very expensive. The environmental "shake-and-bake" chamber costs ~$120k USD, plus installation costs for power (the unit has high instantaneous power requirements for the heater, ~ 36kW) and for liquid nitrogen containers. Additionally, there are continuous operating costs for liquid nitrogen. Total installation and operating costs could be $200k USD. For each type of board, testing requires 5 samples, driven to destruction, to gather sufficient statistical data for HASS development. For Baseline Boards and Station Boards with replacement costs of ~$25k and $15k respectively, and Fanout Boards at ~$800 each this is a total board cost of about $200k USD. Total costs for HALT testing is therefore about $400k USD. Additionally, each board is HASS tested, taking a minimum of about 0.75 hr per (large board) to test. For all ~170 Baseline Boards, and ~140 Station Boards, this is ~$17k USD in technician time (at $75/hr).

HALT testing, if contracted-out could be reduced to $65k plus board costs, but there are a lot of hidden costs associated with such an approach (test bed setup and shipping, travel, complexity), and it is not a preferred approach, mostly because of the complexity of scheduling and logistics. HASS testing would have to be contracted-out as well, and that cost is not known but could be quite high (based on the HALT estimate), about $1700/hr, or $395k for 0.75 hrs per board assuming 1 board in the chamber at a time (possibly it could be a factor of 5 or 10 better if the test chamber is big enough). This is not to mention the shipping time and costs, and hassles going over the border and back (the HALT quote from BreconRidge is for a company in Boston).

Total cost for a HALT/HASS testing program for the EVLA Correlator is about $420k USD, ~$265k just for HALT if it is contracted-out, plus HASS tests. It is hard to justify this expense.

This analysis does not include the need to effectively schedule testing to mesh with product build stages. It can't be done on the Stage 1 prototypes because there aren't enough of them. It can't be done on Stage 2 prototypes for the same reason (only 2 Baseline Boards). It would have to be done on Stage 3 prototypes, and then it is too late to make any major design modifications, unless they are total design blunders. This may be the case for portable electronics that gets bumped around, but it is unlikely the case for rack-mount equipment living a cushy lifetime.

The only major concern for a destructive defect that could find its way into the system is the custom correlator chip. Although, it is still being produced by TSMC and a packager

that “should” have good quality for the same reason. However, since it is a BGA device, it is prudent to develop an effective pre-board-stuff screen to minimize board re-heat cycles (to replace defective devices), of which there are a limited number allowed on any one circuit board.

An important point precipitated from reading the literature is that an ESS program, to be effective, must be suitable for the particular product, and must entail feedback and adjustment of parameters as more is learned to ensure continued effectiveness. A starting point must be defined, and then refined over time. Since there aren’t many units (compared to volume production in industry) being produced or tested for the EVLA Correlator, and the “starting point” is unknown (except for that which might be quantified by expensive HALT), a reasonable approach to testing must be taken.

What is believed to be a relatively inexpensive and reasonable approach to an effective ESS program, developed from research into the testing literature, is described in the remainder of this test plan document.

4 Correlator Chip Test

The correlator chip is a standard-cell ASIC fabricated in 0.13 um CMOS. It is perhaps the most important area for ESS since there are so many of them in the system, and even a small percentage of failures can have big consequences in board re-work costs and subsequent board reliability degradation and early end of life.

The chips are delivered to us from the scan testing facility (DA-Test), after having undergone scan test. The scan test ensures that all logic in the chip is exercised at least once, except that it does not cover the ripple accumulators. The scan test is run at low speed, and high temperature. There is some concern, expressed by the author, and from the experience of chip failures in the ALMA correlator (the design of which doesn't include scan testing, but rather functional testing), and expressed by NRAO ALMA correlator engineers, that this is not good enough (ignoring that one way or another ripple accumulator testing must be performed before the chips are soldered on boards). We have no data to confirm that scan testing is not good enough, and indeed some of the literature indicates that it *is* good enough provided the company (i.e. TSMC and the packaging house) adheres to industry quality control procedures.

It would seem that a reasonable pre-board-stuff testing procedure of the correlator chips is in order, provided it doesn't cost very much. Not "costing very much" is difficult to accurately quantify since it is not known, apriori, what the defect percentage might be and therefore what the cost impact on production modules will be. However, if we assume that 0.1% of the chips are bad (not including ripple accumulators that have to be tested in any case), then, out of 12k chips, 12 might be bad, and this could result in 12 boards that require a re-work to replace a bad chip. Each re-work costs 1/3rd of the board's number of re-works or ~$25k/3=$8.3k. If there are 12 bad chips, if not found they might end up costing us 12x$8.3k=$99.6k.

The ALMA chip experience (same design contractor [ISine], same fab house [TSMC], different package [PQFP], different technology [0.25 um], different test [functional vs. scan testing]) was that one lot had a 2.23% failure rate, most were in the ~0.2% failure rate regime, one lot was 0.72% failure, and one lot had a 0% failure rate. The weighted average failure rate across lots was 0.277%. Using this failure rate number, and the cost of defects developed above, results in a potential cost of $275k if these (33.24) bad chips are not found before they are soldered on boards. This, in itself, could more than pay for an environmental chamber, not to mention the time and hassle of replacing chips on up to 33 boards!

Nevertheless, the environmental chamber is still expensive in absolute project dollars terms. It is therefore advantageous if a relatively inexpensive and effective test method was available. Fortunately, this is possibly the case, as described below.

The proposed testing procedure, and the rational for each test is as follows. These tests are based on [Hnatek, 1995], [Scheiber, 2001], MIL-STD-883, and JEDEC JESD22-A105C.

4.1 Test #1—Room Temperature Functional Test

This test is run on the chips by plugging them into one or more Baseline Boards with BGA sockets for the chip. This test finds ripple accumulator defects not covered by scan insertion and is the first full-speed test of the chips. It also forms the baseline to determine the effectiveness of ESS testing (i.e. only chips functioning at speed go onto further testing, so any failures found are due to ESS testing). The test need only be run long enough to establish correct chip functioning.

4.2 Test #2—Temperature Cycling

Temperature cycling, with no power applied, is performed between –65 ^{o}C and +150 ^{o}C.

Chips can be cycled in quite large batches, easily a few hundred at a time, with an aluminum carrier, essentially a glorified layered cookie sheet that secures the chips in place and facilitates maximum thermal transfer to the chips. To do this, a hot box and a cold box are required, and the test procedure requires transfer of chips back and forth between the hot and cold box for each cycle, with maximum 1 minute transition time. Commercial freezers capable of –85 ^{o}C and ovens capable of +300 ^{o}C are relatively inexpensive:

- An industrial freezer from Freezer Concepts (freezer-concepts.com), model C85-3 has a 3 ft^3 capacity, is temperature programmable from –18 ^{o}C to –85 ^{o}C, has a closed cycle refrigeration unit, and requires no liquid nitrogen. Its inside dimensions are 24” x 12” x 18” deep—it can therefore hold chips as well as a small quantity of even the largest correlator boards. The price is $4980 USD. A picture of this freezer is shown below:

Figure 4-1 Industral freezer with –85 ^{o}C capability.

This freezer would easily fit in the DRAO lab (outside dimensions 34" W x 27" D x 47" H), and it is powered off single phase 110 VAC. This freezer is normally used for biology and chemical specimens. It may be necessary to load the freezer with a cold mass so that insertion of the hot specimen minimizes the freezer compartment temperature change. A fan is necessary to facilitate fast thermal transfer; this may require a slightly larger unit (the C85-9 has 47" x 16" x 20" interior dimensions).

- An industrial forced-air oven that might work well is the Cascade Tek (cascadetek.com) model TFO-5. It is capable of Tamb+5 ^{o}C to +300 ^{o}C operation, has a programmable temperature with 4.9 ft^3 capacity, and has inside dimensions 23" W x 20" D x 18.5" H, so it can fit chips as well as a few large boards. The list price of this oven is $3341 USD. This oven is shown below:

Figure 4-2 Industrial oven with +300 ^{o}C capability.

This oven could easily fit in the lab and runs off single phase 230 VAC, 9 A with a 2000 W heater.

Using separate hot and cold boxes is cost effective since it achieves effective and fast temperature changes, is inexpensive, and for the volume of the EVLA Correlator is cost-effective.

Note that Tenney Environmental (thermalproductsolutions.com/Tenney/thermal-shock.asp) makes dual and triple chamber boxes that are built for exactly the purpose of this kind of temperature cycling. The TUJR is an upright floor model with a temperature range of –75 ^{o}C to +200 ^{o}C, with a 1.25 ft^3 capacity. It is priced at $5300 USD. This unit automatically transfers samples back and forth between the hot and cold sections. It is large enough for correlator chip testing, but is not large enough for board testing. Further

investigation is required to determine if a unit like this is more cost effective than purchasing separate ovens and freezers.

This test is based on MIL STD-883 Method 5004/1010 for Class Level B[1] devices, "condition C". Several sources indicate that rapid temperature cycles with air thermal transfer with no power applied is the single most effective technique for finding defects in both silicon and boards. It induces differential thermal expansion and contraction that is very effective at finding defects. (Thermal shocking with liquid is recommended only for audit testing and not ESS because it can induce failures that are otherwise not present; i.e. the temperature change is too fast.)

The MIL STD-883 temperature cycle curve is as follows:

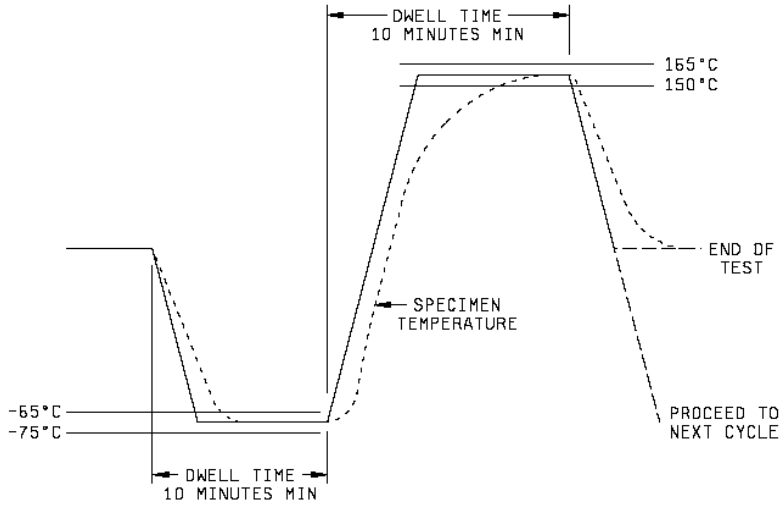

Figure 4-3 MIL-STD-883 Method 1010 temperature cycling curve.

MIL-STD-883 is also referenced in JEDEC standard JESD47D "Stress-Test-Driven Qualification of Integrated Circuits".

The number of temperature cycles that are required for each batch of chips to form an effective screen should be minimized to minimize testing time. Provided that a fast temperature ramp is used, at least ~40 ^{o}C/minute, and that the dwell time is long enough to reach the temperature extreme, only one or two cycles are required. The oven contains fans, but the freezer doesn't (I don't think so, anyway). It may be necessary to set up a fan in the freezer to facilitate fast thermal transfer.

Initial recommendation for the test is 1.5 cycles, which is:

- Start at room temp.

[1] High reliability, military grade but not space-qualified.

- -65 ^{o}C for 10 minutes.
- +150 ^{o}C for 10 minutes.
- -65 ^{o}C for 10 minutes.
- End at room temperature.

This test ensures rapid positive and negative temperature swings between the two extremes. Note that MIL-STD-883 Method 5004/1010 indicates 10 temperature cycles for Class Level B devices. Only 1.5 cycles are used as the initial recommendation, because there is an indication that it is a sufficient test provided the rate of change of temperature is at least 40 ^{o}C/minute [Scheiber, 2001].

4.3 Test #3—Room Temperature Electrical Test for 1 hour

In this test, the chips are loaded into a test board (Baseline Boards with monolithic BGA socket), and operated at full speed and fully functional at room temperature ambient for ~1 hr, with normal core and I/O voltages. Testing should perform the full battery of tests to properly exercise the chips. This test forms the second ESS baseline to determine the effectiveness of Test #2.

4.4 Test #4—High Temperature Electrical Stress Test

In this test, the temperature is set for high temperature (Tj~125 ^{o}C[2]) for some period of time, likely 4 hrs[3]. During this test, the core voltage and I/O voltages are set to their maximum rated values (or 20% overvoltage, whichever is less) to induce electro-migration, hot-carrier injection, and oxide defect failures [Hnatek, 1995].

The test boards have to be jury-rigged to facilitate easily changing the +2.5 V I/O voltage trimming resistors[4], and the Accel chip has to be re-programmed to set a higher core voltage, probably 1.2 V.

Note that MIL-STD-883 Method 5004/1010 indicates 160 hours at Tj~125 ^{o}C, which seems generally accepted to be equivalent to 10,000 hours of operation at Tj~50 ^{o}C [Pecht, 1999][Hnatek, 1995], although no electrical overstress is indicated. 160 hours burn-in is likely too expensive and time consuming for this project, given the perceived effectiveness of the temperature cycling test.

The chips are continuously exercised and functionally tested during this test. The room ambient temperature will likely have to be increased, and the cooling fans decreased in

[2] This temperature may exceed the operating temperature of the z-axis BGA socket pads, and may need to be dropped to 100 ^{o}C.
[3] 8 hours is probably better, but it is highly desirable to completely test a batch of chips from start to finish in 1 day. The defect rate will be used to indicate if more time is necessary.
[4] It may not be possible or necessary to increase I/O voltage for this test. Further research is required.

speed during this test. This test can most effectively take place in the DRAO lab's environmental chamber.

4.5 Test #5—Powered Temperature Cycle

In this test, the core and I/O voltages are set back to normal, and the room temperature is ramped from the high temperature of Test #4 to –10 ^{o}C over 15-25 minutes, held at low temperature for 30 minutes, and then brought back up to room temperature. (Changing board voltage settings, in particular the 2.5 V I/O setting will require a power cycle.) Fans will run at full speed for this test. The chips are continuously exercised and functionally tested during this test.

There is no exact rational for this final test (although it is similar to JEDEC JESD22-A105C, except for the power cycling part); it is one last "comfort test" that the chip works over a wide temperature range at normal operating voltage. If it is easy, during this test, the clock frequency of the chip should be increased somewhat (say 10%, but determined experimentally) to test for speed margin.

4.6 Test Effectiveness Evaluation

The durations of tests, the number of temperature cycles, and even the temperature extremes may need to be modified depending on the failure rate.

If the failure rate is too high, it is likely not due to the tests and could indicate a bad chip lot. In this case, I should think that tests should be repeated and even have more and longer durations to see if more chips fail, or if the first screen did the trick.

If the failure rate is too low, the tests should be modified until the 0.1% failure rate is reached. This should be done incrementally by doubling the number of cycles or test duration as the case may be. If the failure rate does not reach 0.1% at 24 hours of total test time for a batch, it indicates that the screen is not effective in finding failures and that chip and process quality is very high. It is hard to imagine the screen being more severe, and so it may then be appropriate to perform audit testing, rather than testing every single chip (although, every single chip still has to be tested functionally, so Test #1, and Test #2, followed by Test #1 should be done on every chip).

4.7 Lifetime Degradation Evaluation

A statistically significant number of chips, say two boards worth or 128 chips, should be evaluated for lifetime degradation. This is a method employed for HASS lifetime degradation evaluation. In this test, the full set of tests is performed ~10 times to see if the failure rate increases over performing the tests just once. If there is no failure rate increase, then it can safely be said that $<1/10^{th}$ of the chips' lifetime has been used by the tests. If the failure rate increases, then it is an indication that the tests could be too severe and are using up too much chip lifetime, or that there is some low-level process defect that the single ESS is not finding.

Re-evaluation and re-development of the ESS may be needed, depending on the outcome of this test.

4.8 Test Batch Size

I think a reasonable test batch size is 256 chips, or 4 test boards worth, valued at about $41k. This is enough quantity to start to get some failure rate statistics after the first batch, and at the 0.1% level, within the first 4 or 5 batches, to allow adjustment of test parameters in the early going as described.

The monolithic BGA socket will have to be modified slightly so that every time the board is loaded with chips, it is not necessary to install (and later remove) 256 tiny screws. Likely a single plate with a spring over each chip site, can be used to quickly install and remove chips from the test board. This plate must allow for an air gap for cooling, or must thermally couple the chips to the plate.

Four Baseline Boards, required as test vehicles, will also have to be built. Likely these boards are built in Stage 2.

4.9 Alternative Test Strategy

If the monolithic BGA socket is not reliable then the testing strategy will have to be re-thought. The ripple-accumulator test still has to be performed, and requires that we develop a go/no-go single device chip tester with a ZIF socket. Chips are tested at DA-Test with this tester, retested here with a copy of the same chip tester, undergo Test #2 temperature cycle test, and then set on a hot plate at 125 ^{o}C, while chips are tested, one-by-one, in the tester. Each chip takes probably 30 seconds (including insertion and extraction) to test in the tester, and so it will likely take a week or two to test all chips.

4.10 Correlator Chip Prototype Testing

200 untested prototype chips will be delivered for testing and evaluation at DRAO, the first few of them in June 2006. Once the chips are functionally tested, it is prudent to test them according to the testing strategy outlined in the previous sub-sections, to give some early idea of defect percentage, and suitability of the ESS testing strategy.

The first Baseline Board prototype with the monolithic BGA socket is a suitable electrical test vehicle for the chips. The tests should initially be performed on only ~10 chips. This is not a statistically significant number of chips, but it will give us an early indication if we are doing something destructive, or if there is some design flaw or process defect. We need to be careful not to destroy these precious prototypes since they are needed for prototype correlators.

Subsequent batches of chips could be tested 64 at a time, provided there are no failures in the first 10 chips.

4.11 Production Audit Testing: FPGA Mfg Tests

It is interesting to note what FPGA manufacturers do for chip testing. Altera publishes a quarterly report (http://www.altera.com/literature/rr/rr.pdf) indicating the results of reliability audit testing of their chips. A summary of tests is as follows:

- Lifetest (JESD22-A108). Junction temperature 130 ^{o}C, 20 % overvoltage, 25 to 60 units for a total 1000 to 2000 lifetest-hours.
- High temperature storage at 150 ^{o}C, batch size of 25-100, for total stress hours of about 1000. (JESD22-A103).
- Accelerated moisture resistance (JESD-A101 or A110). Autoclave 120 ^{o}C, saturated steam environment at 2 atmospheres; unbiased HAST, 130 ^{o}C, 85% relative humidity.
- Reflow simulation and moisture preconditioning. Moisture soak devices, followed by pre-bake of 150 ^{o}C for 2 minutes, following by 100% convention reflow simulation 3 times, with peak temperature of 220 ^{o}C.
- Temperature humidity bias (THB) test. 85 ^{o}C, 85% R.H., with voltage applied. 1000-2000 stress hours.
- HAST. THB test under pressure, 130 ^{o}C, 85% R.H. 100 stress hours.
- Temperature cycling using dual-chamber (hot/cold boxes). MIL-STD-883 "condition K" (0 ^{o}C to +125 ^{o}C) for very large die, and "condition B" (-55 ^{o}C to +125 ^{o}C) for other devices. Devices are electrically tested after 500 and 1000 cycles.
- Solder joint reliability. Measured by cycling devices on a PCB from 0 ^{o}C to 100 ^{o}C at 2 cycles per hour, under power. 5000 cycles.

Clearly it will be very expensive to do these kinds of audit tests on correlator chips. It may be possible to contract an external company to do these tests, although special PCBs have to be developed for many tests, and each test requires about 60-100 chips. This possibility should be investigated to determine cost and feasibility. John Wu, the Altera FAE in Vancouver, provided the following names of testing labs to do these kinds of tests:

- Minco, Austin TX.
- Austin Sem-Conductor, Austin TX.
- QP Labs, San Jose, CA.
- Stellar, Vallencia, CA.

Altera won't tell us what they do for full production screen testing, except to say that testing guarantees that devices meet data sheet requirements.

5 Board Test

Assembled boards are tested in a similar manner and for a similar rational as the correlator chip. Both the industrial freezer and oven discussed in the previous section have compartments that are large enough to accommodate a few of even the largest boards. A multi-board carrier will have to be developed that allows the boards to be safely and quickly transferred from the freezer to the oven and back.

Board tests are used to find both chip-level and board-level defects, and to allow for gathering of statistics to evaluate the effectiveness and suitability of the ESS.

Tests are defined in the following sub-sections.

5.1 Test #1—Room Temperature Functional Test

Boards arrive from the assembly house (BreconRidge) and are plugged into the functional tester only long enough to verify that they are working. Any defects found at this stage must be repaired, and noted, before moving on. No heat-sinks are attached, since this test aims to find silicon, PCB, and assembly defects and installation of the heat-sink(s) at this point may require removal for repair later.

This forms the baseline to evaluate the effectiveness of the ESS program.

5.2 Test #2—Temperature Cycling and Mechanical Shock

Temperature cycling, with no power applied, is performed between –40 ^{o}C and +100 ^{o}C. A few boards at a time are fastened to a carrier that allows them to be manually transferred, quickly and easily back and forth between the oven and the freezer. A double shift can be run whereby both the oven and the freezer are loaded with boards, and then swapped at the end of the dwell time.

In the short time while the boards are being transferred between the oven and the freezer (and the other way), the board carrier holding the boards undergoes a mechanical shock of ~70 g peak, consecutively on all three axes, in both directions. This shock is introduced with the used of a controlled drop test setup, pre-calibrated with an accelerometer. This test should take only 10 seconds or so to execute, and is performed when the boards are at their temperature extremes. This test is optional, and depends on suitability and effectiveness (i.e. it requires some evaluation to determine if it should be done).

The soak time in each location needs to be empirically determined by running some tests with thermo-couples, but likely 10 minutes is sufficient. The test is the same as for the correlator chips, and will have to be refined somewhat as statistics are gathered.

- Start at room temp. Mechanical shock (x,y).

- -40 ^{o}C for 10 minutes.
- Mechanical shock (-x,-y). Transfer to oven.
- +100 ^{o}C for 10 minutes.
- Mechanical shock (z,-z). Transfer to freezer.
- -40 ^{o}C for 10 minutes.
- End at room temperature.

This test ensures rapid positive and negative temperature swings between the two extremes, with some mechanical shock to help "loosen things up". Alternatively, a commercial multi-axis broad-spectrum vibration table could be purchased for ~$22k, and each board could be vibration tested when transitioning between low/high temperatures.

5.3 Test #3—Room Temperature Electrical Test for 1 hour

Boards are loaded into test crates and racks and run at full speed at room temperature ambient for 1 hour. A full suite of tests is run to exercise the boards fully to find any defects. No heat sinks are attached. If any defects are found, they are noted, repaired and Test #2 is re-run before proceeding.

The test crates and racks are located in the DRAO lab's environmental room for this test, in anticipation of Test #4.

5.4 Test #4—High Temperature Electrical Test

Boards are run at elevated temperature for 4 hours (8 hours—overnight—is better provided that there is protection in place for thermal runaway) in the DRAO environmental room. The room temperature is increased to +50 ^{o}C, and board temperature sensors and fan control feed back should be used to operate the boards at a junction temperature of 100-125 ^{o}C for the duration of this test. The boards must be active at full speed and running a full suite of tests for this test. The heat sink(s) is not attached for this test.

This test could be run for 160 hours at Tj~125 ^{o}C, equivalent to 10,000 hours of normal operation. However, the perceived effectiveness of the temperature cycle test likely eliminates the need for this kind of prolonged burn-in test [Scheiber, 2001].

If a failure occurs during this test, it must be determined if it is a hard or soft failure. If it is a soft failure, then investigation needs to be done to reveal the cause of the failure before corrective action is taken (e.g. is it a timing failure, indicating not enough timing margin?) The temperature should only be adjusted (decreased) if it is not possible to fix the problem with a design change. If it is a hard failure (something broke), then it could

be a genuine defect, or it could indicate that some device is actually operating outside its normal operating range, in which case the temperature should be adjusted.

5.5 Test #5—Powered Temperature Cycle

Boards are visually inspected for any obvious damage that might not be detected with functional tests (e.g. cracked decoupling capacitor solder joints). Heat sinks are attached and secured in their final configuration, along with thermal overload protection, before boards undergo this test. This test is meant to ensure that the boards operate satisfactorily over a range of temperatures they are likely to see in their lifetime, and follows a similar vein as JEDEC JESD22-A105C, except the temperature extremes are not as great.

If possible, the test rack should be setup to shake somewhat while this test is occurring by attaching an unbalanced, rotating load to the rack, simulating exaggerated fan vibration.

The boards are installed in the test racks in the DRAO environmental room. The room is set to transition between –10 ^{o}C and +50 ^{o}C as rapidly as possible (likely about 1 ^{o}C/minute) at least 1.5 cycles over 2-3 hours. Boards are active and running a full suite of tests while this test is running. The test cycle is as follows:

- Start at room temperature.
- Ramp to the high temperature.
- Wait 10 minutes at the high temperature.
- Ramp to the low temperature.
- Wait 10 minutes at the low temperature.
- Ramp to the high temperature.
- Wait 10 minutes at the high temperature.
- Ramp back to room temperature.

The boards will normally be powered on and off by the thermal overload protectors when high temperature is reached. Power cycles at low temperature, every 10 minutes with a 50% duty cycle, should be performed (JESD22-A105C).

5.6 Test Effectiveness Evaluation

The durations of tests, the number of temperature cycles, and even the temperature extremes may need to be modified depending on the failure rate.

If the failure rate is too high, the causes should be evaluated to determine whether the failure is a true defect or whether the test needs to be re-evaluated.

If the failure rate is too low, the tests may need to be modified until the 3%[5] failure rate is reached. This should be done incrementally by doubling the number of cycles or test duration as the case may be. If the failure rate does not reach 3% at 24 hours of total test time for a batch of boards, it indicates that the screen is not effective in finding failures and that board fabrication and assembly process quality is very high. It is hard to imagine the screen being more severe without becoming destructive. In this case, 100% screening could be reduced to an audit test (20%?), or the screen could continue since there are not that many boards to test anyway.

5.7 Lifetime Degradation Evaluation

A statistically significant number of boards, say 4 of each type, should be evaluated for lifetime degradation. This is a method employed for HASS lifetime degradation evaluation. In this test, the full set of tests is performed 10X to see if the failure rate increases over performing the tests just once. If there is no failure rate increase, then it can safely be said that $<1/10^{th}$ of the boards' lifetime has been used by the tests. If the failure rate increases, then it is an indication that the tests could be too severe and are using up too much board lifetime, or that there is some low-level process defect that the single ESS is not finding.

Re-evaluation and re-development of the ESS may be needed, depending on the outcome of this test.

5.8 Test Batch Size

I think a reasonable test batch size is 4 large boards (that fit in the freezer and the oven), and perhaps 10 or 20 small boards (e.g. Fanout Boards). This is enough quantity to start to get some failure rate statistics after the first couple of batches, and to allow adjustment of test parameters in the early going as described.

5.9 Board Prototype Testing

Stage 2 boards should undergo the prescribed testing to get some early data on the effectiveness and viability of the tests. Since there are only 2 Stage 2 Baseline Boards, these should be tested 1 at a time, in case there is some destruct mechanism at play, since the boards are required for prototype correlators.

[5] 3% failure rate is a subjective number, but is chosen to be the same number that BreconRidge considers as an average percentage of boards that can be expected to be unfixable defects from the manufacturing process.

6 Rack Test

Correlator racks will be fully outfitted with sub-racks, backplanes, cables etc. in Penticton, and then shipped by ground transport to the VLA site. The racks will not be shipped with boards in them, however it is important to test the racks to ensure that they will arrive at the site undamaged, and without screws coming loose or parts breaking over the temperature extremes expected during transport.

The first possibility is to do nothing, and just ensure that everything is designed as well as possible. This includes the use of split ring spring washers on all fasteners, adequate support brackets etc. The problem is with this is that it is an open-loop solution; it does not allow for corrective action until it is too late.

The second possibility is to fully assemble a rack, ship it to the site, have it shipped back, evaluate the rack for damage, and take mitigating action on the remaining racks. The problem with this approach is that it does not test every rack. This could be done with accelerometers attached to the rack so that we can determine what has happened to it so as to assist with the design of a shaker (below).

The third possibility is to bolt a couple of racks at a time to the back of the NRC pick-up truck and "drive around" with them on the bumpy roads out to the LAR hangar and back a few times. This is hard to quantify, may be too severe, and costs manpower time.

The fourth possibility is to build a shaker table for this purpose. This shaker table contains a motor connected to a crank arm, driving a rolling platform running over rumble strips. The rack is bolted to the platform, accelerometers are attached to the rack, and the test is run for a desired duration, probably 24 hours or so per rack. This table should be able to be easily built, and adjusted to the intensity that is "about right" by adjusting the motor speed, the crank arm length, and the rumble strip width and depth. There is likely a standard specification that indicates the g forces in all 3 axes and shake frequency that is expected on ground transport.

A diagram of this shaker table is shown in Figure 6-1. Tests are performed in the DRAO lab environmental room over the temperature range of –10 ^{o}C to +50 ^{o}C (approximately 1 cycle every 2 or 3 hours) for a period of 24 hours[6] for each rack.

A final alternative is to buy a shaker table that is meant for this purpose. This possibility should be investigated before a shaker table is built.

The first fully built prototype rack must be tested before other racks are constructed to allow for corrective action, if necessary.

[6] 48 hours would be better since 48 hours is likely the duration of "moving ship time" to the VLA site.

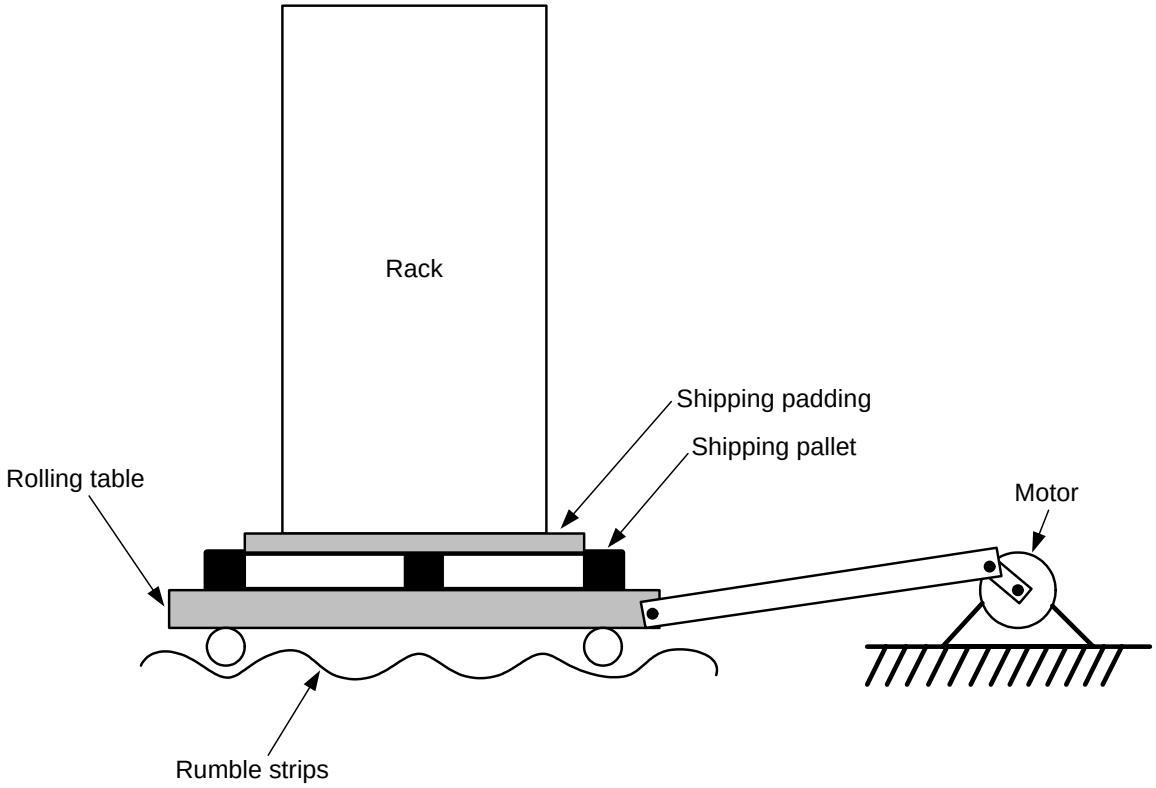

Figure 6-1 Rack shaker table to simulate ground transport mechanical motion.

7 References

Altera Reliability Report 43, Q4 2005, http://www.altera.com/literature/rr/rr.pdf

Davis, Don and Brendan Davis. “The Economics of Stress Screening,” Proceedings of Nepcon West, Cahners Exposition Group, 1989, p. 1813.

Hnatek, Eugene R., “Integrated Circuit Quality and Reliability”, 2nd edition, 1995, Marcel Dekker Inc.

Pecht, Michael G. et. al. “Guidebook for Managing Silicon Chip Reliability”, 1999, CRC Press LLC.

Scheiber, Stephen F., “Building a Successful Board-Test Strategy”, 2001, Butterworth-Heinemann.