

A SOFTWARE SUITE FOR VALIDATION OF X2O BOARD

PERFORMANCE

An Undergraduate Research Scholars Thesis

by

JESSICA WILLIAMS

Submitted to the LAUNCH: Undergraduate Research office at
Texas A&M University
in partial fulfillment of the requirements for the designation as an

UNDERGRADUATE RESEARCH SCHOLAR

Approved by
Faculty Research Advisors:

Dr. Alexei Safonov
Dr. Jyh-Charn Liu

May 2023

Majors:

Computer Science
Mathematics

Copyright © 2023. Jessica Williams.

RESEARCH COMPLIANCE CERTIFICATION

Research activities involving the use of human subjects, vertebrate animals, and/or biohazards must be reviewed and approved by the appropriate Texas A&M University regulatory research committee (i.e., IRB, IACUC, IBC) before the activity can commence. This requirement applies to activities conducted at Texas A&M and to activities conducted at non-Texas A&M facilities or institutions. In both cases, students are responsible for working with the relevant Texas A&M research compliance program to ensure and document that all Texas A&M compliance obligations are met before the study begins.

I, Jessica Williams, certify that all research compliance requirements related to this Undergraduate Research Scholars thesis have been addressed with my Faculty Research Advisors prior to the collection of any data used in this final thesis submission.

This project did not require approval from the Texas A&M University Research Compliance & Biosafety office.

TABLE OF CONTENTS

	Page
ABSTRACT	1
ACKNOWLEDGMENTS	2
NOMENCLATURE	3
1. INTRODUCTION.....	4
1.1 Need For the X2O.....	4
1.2 Describing the X2O	6
2. METHODS	8
2.1 Development Methodology	8
2.2 Test Setup and Dependencies.....	8
2.3 Safety Considerations	10
2.4 Test Flow	10
2.5 Data Output and Structuring.....	12
2.6 Failure Detection and Recovery	12
3. RESULTS.....	14
3.1 Physical Data	14
3.2 FPGA Programming	17
3.3 Register Access Speed	18
3.4 Clock Frequency	19
4. CONCLUSION.....	21
4.1 Future Research	21
REFERENCES	23

ABSTRACT

A Software Suite for Validation of X2O Board Performance

Jessica Williams
Department of Computer Science
Department of Mathematics
Texas A&M University

Faculty Research Advisor: Dr. Alexei Safonov
Department of Physics
Texas A&M University

Faculty Research Advisor: Dr. Jyh-Charn Liu
Department of Computer Science
Texas A&M University

As the Large Hadron Collider (LHC) undergoes upgrades to increase luminosity, its detectors will need faster data processing and improved electronics. The X2O is a custom electronics board to support operations in the Endcap Muon System of the Compact Muon Solenoid experiment (CMS). It has a power module, a FPGA, and 120 high speed optical links to receive, process, and communicate data.

When the design of the X2O is fully validated, the boards will go to be manufactured and each board will need to be thoroughly tested to meet the design specifications. A software suite was created to test the X2O. It analyzes physical characteristics like voltage, temperature, power, and current. It also checks programming the FPGA, reading and writing register values, and the reference clock frequencies. Operating parameters were almost all within ranges, and the X2O is ready to move to production. Future work should center around tightening performance bounds, and testing the X2O at other places in the data path it will be implemented.

ACKNOWLEDGMENTS

Contributors

I would like to thank my faculty advisors, Dr. Safonov and Dr. Liu, for their guidance and support throughout the course of this research. I would also like to thank Dr. Gilmore for all of his knowledge, kindness, and being a first line of defense.

Thanks especially to Dr. Juska with CERN for providing direction, being a technical liaison, and always being willing to help debug. Thanks also to the CMS colleagues.

Thanks also go to my friends and colleagues and the department faculty and staff for making my time at Texas A&M University a great experience.

Finally, thanks to my family for their encouragement and love.

The data analysis in this report was conducted in part by Ali Ahmad; special thanks to him for helping test run the software as well.

All other work conducted for the thesis was completed by the student independently.

Funding Sources

Undergraduate research was supported by Texas A&M University. This work was also made possible by those at CERN.

NOMENCLATURE

CERN	European Organization for Nuclear Research
CMS	Compact Muon Solenoid
CSC	Cathode Strip Chamber
FPGA	Field Programmable Gate Array
GEM	Gas Electron Multiplier
LHC	Large Hadron Collider
QSFP	Optical Transceiver plug-in
X2O	Electronics Board for Data Processing

1. INTRODUCTION

1.1 Need For the X2O

1.1.1 High-Luminosity Upgrades

The Large Hadron Collider (LHC) was built in Geneva, Switzerland to collide proton beams at high energies and discover new physics by tracking and detecting products from collisions [1]. The large number of collisions requires a large amount of computing power to process the data. Protons travel through the ring in dense bunches, and pairs of these bunches collide every 25 nanoseconds. Each bunch crossing, as they are called, has many proton-proton collisions that create a cascade with a variety of decay particles depending on the energy of the collision. In order to look for physics beyond the standard model, detector and electronics upgrades are underway to increase the instantaneous luminosity by up to five times prior levels in a series of upgrades known as the High-Luminosity LHC [1]. This will increase the data rate by a similar factor, so the electronics handling the data need to be upgraded as well.

CERN has a long history of creating new frontiers in computing. They first gave rise to the Internet to connect members across countries. They invented distributed computing and many ideas in parallel processing came from implementations of CERN's work. The high-luminosity upgrades are the latest in this long history of computing at the edge of what is possible, and working on ways to capture, process, and store as much data as possible will lead to better physics resolution and more clear answers about the laws governing the universe.

The various stations and experiments on the LHC operate largely independently of each other, as they each have independent collisions and bunch crossings at the core of each of them. Because of the huge amount of data generated, there must be layers of data processing and triggering to keep only the data likely to be interesting. The experiments can compare measured values and high-level probabilities and event distributions, but each experiment does its own reconstruction of the particles and their energies and tracks independently. This allows confirmation, like

the discovery of the Higgs Boson was joint between ATLAS and CMS. They each had different particles and collisions, but constructed results that were indicative of the Higgs Boson and agreed with each other. Reconstructing the tracks and types of particles requires several layers and steps in its own right.

1.1.2 The Compact Muon Solenoid

The Compact Muon Solenoid experiment (CMS) is one of the main experiments on the LHC. It has several detectors for reconstructing the direction and energy of the particles that pass through, such as the muon detectors, silicon tracker, and electromagnetic and hadronic calorimeter [2]. Each layer has its own types of detectors and various qualities it can measure, mostly through electromagnetic and hadronic interactions. There are electronics systems and software tools to interpret the data from each layer of detectors into localized tracks, and then systems to blend the data from the layers into a global picture.

The system of interest for this project is the Endcap Muon System, which detects and reconstructs the path of muons, which are heavier cousins of electrons produced in decays of real or virtual W and Z bosons. This contains layers of Cathode Strip Chambers (CSCs) and Gas Electron Multipliers (GEMs) that produce raw data about interactions on the detectors due to a muon passing through, usually through ionization of the gas[3]. For instance, the CSCs work by having six layers. A passing muon will ionize the gas in the detector as it passes through, and because of the electric field this will create a detectable electromagnetic buildup. If there is four out of six layers hit at once, that track is considered to be a candidate for a muon.

The data is first preprocessed on the on-chamber electronics, such as Cathode Front End Boards and Anode Front End Boards for the CSC detectors. Then the data needs to be combined into muon candidates, which is like a segment of a track that may or may not be a muon. Because electronics noise or spurious signals from non-muon backgrounds like neutron interactions can also result in detectable signals, it is important to determine what is actually a muon and what is not. The CSC and GEM data needs to be combined to have a coherent picture of any possible muons that passed through the endcap, and because muons must be assigned to a specific event,

the system needs time resolution down to the bunch crossing.

Once the data from the endcap is combined and processed, it goes to the Muon Trigger. This combines data from the endcap and barrel regions to determine muon paths for that bunch crossing across the entire experiment. It is important to account for all muons because if they are not accounted for, there could appear to be missing mass or momentum which is a sign of physics not yet understood. The data from the Muon Trigger about how many muons were created and where they went is sent to the L1 Trigger, which is the first experiment-wide determination of what data to keep and what to not. Since quantum physics is probabilistic, there are many common collision results that have been studied extensively and are not interesting. The data from those, or from partially recorded collisions, will be discarded. Interesting and uncommon collision results will be kept and passed on to the next trigger and go further for analysis.

The X2O board will be used in two different locations along the data path. One is towards in the trigger data path to implement the Muon Track Finder and the Global Muon Trigger [4]. These algorithms collate the data from detectors to reconstruct the paths of muons and determine if any were produced in a particular collision. This information is used by higher level triggers (notably the L1 trigger) to decide whether or not to keep the data from the bunch crossing.

The other application that is the focus of this research is its use as a backend system board for GEM and CSC chambers. Here, it provides clock and control signals to the detectors, and reads out their precision data. It performs partial data reconstruction using signals sent from the detectors, performs zero suppression and data packing, and transmits data using high speed optical links. It also provides the clock frequencies to the CSC and GEM frontend computations. It also can read out and communicate with the CSC and GEM boards to configure and monitor their operation.

1.2 Describing the X2O

The electronics and computing systems handling all this data must be carefully designed and their capabilities proven to ensure efficient synchronization when processing data from multiple detectors. The custom backend system designed for accumulating and processing data from

the CSC and GEM muon detectors is the X2O board [4]. It is one roughly 12in x 12in board with three dedicated modules: one for power, another to support a powerful Field-Programmable Gate Array (FPGA), and the last to support the high-speed optical links. The power module contains a Xilinx Zynq system on a chip; this allows control of the X2O from the control computer without the main FPGA being programmed. It is running centos and can be accessed via ssh through the control computer. The key utility of FPGAs is that the firmware can be changed by the user and adapted over time at low time cost, and they can do real time processing because the pipelining allows for fixed latency. The FPGA used is a Xilinx VU13P. The optical module has 30 slots for four-channel optical transceivers for a total of 120 links and dozens of custom high-speed cables connecting them internally.

The ATCA crate that houses the X2O has a number of slots. There are two for control cards, 12 where X2O boards can go, and one shelf manager that handles fan speed. A network switch in one of the control slots provides 10 Gb/s Ethernet to all the boards.

Validating its performance is critical to both ensuring production quality and diagnosing any failures before it is installed. Key metrics include performance functionality, communication reliability, and verifying power and temperature profiles are all within specifications. Low level access has already been implemented through the operating system on the power module of the FPGA; what needs to be done is creating and automating higher level routines for quality assurance and diagnostics [4].

This work builds on and is similar to other testing of electronics boards for the endcap muon system. The optohybrid boards that the X2O controls and communicates had a similar automated testing system developed and implemented. The GEM detectors in the GE2/1 sector will be installed with those optohybrids soon, and they will need X2O boards to support their operation.

2. METHODS

Validating the X2O prior to installation in CMS requires a series of tests. It must meet stringent standards in order to ensure its fitness for use at the LHC, both in terms of data collection and sustainability. The method of research is creating software to analyze different functioning parameters of the X2O, then running that software to collect data. This data will be used as a baseline to diagnose performance issues while running at the LHC as well as for quality control and production testing before installation in the endcap.

2.1 Development Methodology

The testing software suite was developed with an agile methodology. Since most of the factors to be tested can be run independently, features were added incrementally with stable versions. This also allowed for partial data collection throughout the project lifetime. The test is also modular. Since the main test script executes and validates each of the smaller tests, each aspect can be tested and evaluated individually. Many total runs as well as many stress test longer runs of each of the subtests were made. The run parameters for each test are easily configurable using one config file or by manual changing.

2.2 Test Setup and Dependencies

A secondary goal of this research was to create a X2O testing manual to standardize the testing conditions and also make it easy to use. This manual details the procedures and test setup as well as common failures and ways to respond to them. The manual allows for testing to be duplicated in Geneva before board installation in the LHC and to increase the ability to replicate test output. The manual also shows the dependencies for running each of the sub tests individually, and contains acceptable ranges as a guide to interpret the data in greater granularity and detail.

Other materials needed to test the X2O board are 30 100 Gbps QSFP optical transceivers, and 15 optical fan-out cables to connect them. The QSFP receivers must be able to function reliably when received optical power is -11 dbm or less. They were thoroughly tested using a

splitter and optical attenuator, and all are well past that operating parameter. It is also helpful to have an oscilloscope with an optical probe to create eye diagrams, although that is only necessary for in-depth testing or to diagnose unique failures. The board is installed in an ATCA backend crate. This crate is air cooled, and the fan speed is adjustable via software. In operation, there will be multiple X2O boards in one crate in addition to a control board.

The number of X2O and control cards in the ATCA crate varies based on which substation of the endcap it is serving. For the CSC detectors, it will be ten boards in two crates. For ME0, the innermost GEM detectors, it will be 18 boards in 2 crates. For GE 2/1 and GE 1/1, the GEM detectors in the middle sectors, it will be eight boards in one crate each.

All tests in this thesis were conducted on the same X2O board. Production of a few dozen is expected to be over the next year. A protocol for running a cold test on a fresh X2O board was developed and will be used once new boards are built. Upon receiving a new X2O board, a physical inspection would be conducted to make sure there are no shorts. Then it would be installed in the backend crate and turned on. The test scripts would be copied over and the temperature and voltages checked to make sure there are no shorts and no damage occurs to the board during testing. Then QSFP optical transceivers would be inserted in the board, a full load of 30. Optical fibers would be connected so QSFP 0 connects to QSFP 1, QSFP 2 connects to QSFP 3, and so on.

Many firmware versions were used throughout the project, and there will be several in operating conditions. The Endcap Muon System has many sectors and areas where detectors will be that are similar to pie slices. Different firmware will be needed for different detectors, as the number of detectors, their operating parameters, and the geometric data about how to combine the data from those detectors will be different for each combination of detector and sector. There are at least four - CSCs, ME0 which is the innermost GEMs, GE 1/1, and GE 2/1. Lower numbers are closer to the collision point in the center and the higher numbers are further out.

As part of the test stand, most of those firmwares we had partial implementations to mimic a full detector setup. For instance, we connected the X2O to a ME0 optohybrid board. The optohybrid board is the output and preprocessing board for the GEM detectors and was previously

tested at A&M. We also had a GE 2/1 optohybrid board to connect the X2O to. We did not have real detectors or real data, but simulation data is well established and randomly generated data works just as well for analyzing computational capacity.

2.3 Safety Considerations

The X2O board contains highly sensitive electronics. During the physical inspection, it is important to ground yourself fully before touching the board, and to wear a static grounding bracelet. Always using static-proof packaging is critical, as is ensuring proper protection during transportation.

The optical components are also very delicate. It is necessary to plug the QSFPs in the appropriate way and at the proper angle as to not bend any internal metal components. Additionally, the ends of the optical fibers should always be covered when not connected to anything else. Also, optical fiber should not be coiled too tightly or gripped at the edges where it connects, as that can wear away at the casing and expose the raw fiber.

2.4 Test Flow

Once the board is installed and basic operational safety has been ensured, the testing can begin. Operating the test is simply calling one bash script, making it user friendly and ensuring that operating parameters stay the same between tests. All of the software dependencies are in the backend repository, which can easily be copied to the X2O from the controlling computer. The software dependencies are primarily low-level libraries for reading and writing registers to the firmware and getting physical data from the X2O. The software suite itself is contained in a utilities folder with many python and bash scripts to read and analyze data in a structured way. The main script requires no setup to run other than having all the files it needs in the right places. If the user wishes to run a sub-test independently, some of them require the FPGA to be programmed, the optics calibrated, or other dependencies. Those are documented in the user manual also created.

The X2O validation test is run from one central script and is fully automated. It is integrated into the repository for the endcap muon electronics and is portable. The test script is run on the X2O itself. The main script executes various smaller scripts, primarily in python, to test and record

different performance characteristics, and uses exit statuses to determine failures. This allows for a modular flow where one set of parameters is checked and evaluated before moving on to the next set. It also allows the data to be kept together as different sub-tests of a single run. The test will also stop and power off the board if there is a critical safety error, such as temperature or voltage out of range, which is also helped by using exit statuses.

The first test that is run is the physical data such as voltages and temperatures. Running it first allows any critical failures to be detected immediately to minimize risk to the board. This tests temperature, voltage, current, and power for various chips and parts of the X2O. It does the same for any optical transceiver QSFPs that may be attached. The test waits until the physical data stabilizes, then takes 30 reads to get good statistics. A key feature is that if it detects a temperature that is too high, it will exit the test, shut down the X2O and increase the fan speed. A short version that just checks if the temperatures are within operating range was also developed and is called periodically in the main test script. The full data is written to a physical test data file, and the average stabilized values are written to the summary data file.

The next test checks programming the FPGA. It programs it with standard operating firmware, then checks to see if the connection is up. Throughout the testing process, both firmware to support CSC detectors and to support GEM detectors were used. In the future there will be other types of firmware to implement trigger algorithms further along the data path, although this research group is more concerned with the backend detector support implementation. If the chip to chip connection is up, it passes. The chip to chip connection is a link between the operating system for the X2O on the power module, and the FPGA chip on the FPGA module that allows for control and monitoring. It does this test ten times, and each cycle takes about two minutes. If the chip connection is down, it waits three seconds, sends a reset, and checks again. If the chip to chip connection is still down, the test fails.

Next, it checks the register access speeds. These allow for control and monitoring of the board. This is important to ensure it can be quickly programmed. It tests twice - once with static reads and once with random reads and writes. The random read/write should take about twice as

long since it is doing twice the operations, and acceptable values for the average of each operation is around the 50 microsecond range, so a total average of 50 microseconds for the static read and 100 microseconds for each random read and write.

It then checks the stability of clock frequencies. There is one clock that is at 40 MHz, and the rest are at 156.25 MHz. Passing is defined as all clocks having less than 25 parts per million deviation from the stated value. This is an especially critical parameter, as the clocks are used to control the operation of the optical links and the timing constraints are very tight.

The final test is a stress test. The purpose is to establish that the X2O can handle firmware updates in the future that may increase the power and to further characterize the acceptable bounds of operation. The X2O is programmed with a firmware that draws more power and performs more intensive calculations that would be beyond the expected operating needs. Unlike the regular FPGA programming test, it does not check the chip to chip connection since this firmware version does not support that. The fan speed is increased and the physical test is rerun. Again, it is checking to make sure that the temperatures, voltages, powers, and currents all fall in acceptable ranges. If a temperature is outside acceptable, it shuts down the X2O and exits as usual. The data is written just as the regular physical test with all data to one file and the averages to the summary file. The filenames reflect the type of firmware that was used, generally "hot", "csc", or "gem", in addition to the usual date and run number.

2.5 Data Output and Structuring

Data files were designed to comply with CERN database. Each run creates one main summary file and some auxiliary files. The main summary file contains the various tests and if it passed or failed, and the auxiliary contains all the data taken and the various statistics created from it. Files are timestamped to prevent overwrite, and organized by the date they were taken. We manually backed up the data from the x2o onto the control computer.

2.6 Failure Detection and Recovery

Part of the purpose is to identify errors. If there is a critical failure like temperature being too high, the automated test will stop executing and let the user know what happened. If a test fails

in a non-critical fashion, the test will note that in the results but keep going.

Throughout development, all the reasons tests failed were documented and indexed. This will be useful in diagnosing future issues with the boards throughout their life cycle. Most common failures were testing one aspect without the X2O being in the right state. Ensuring that the terminal state matches the firmware also matches the test trying to be run can be hard to keep track of.

3. RESULTS

A completely automated software suite was created for testing the X2O board. Reliable results were generated to be used in comparison with future X2O testing. There is partial data from each test generated throughout the project, as well as full runs and a variety of shorter specific tests. Results have been grouped into each sub test. Most sub tests were run at higher values than are used in the full automated test to get better confidence on the values of the results to ensure as accurate a baseline as possible across many days and firmware variations. Analyzing was complicated by the fact that development changed the code for running the tests throughout the project timeline.

3.1 Physical Data

The software tools to test physical characteristics have been developed, tested, and shown to record values accurately. The purpose of the software is to test different parameters, and it fulfills that purpose. The conditions in the lab are similar to those after installation, so it is a good approximation of expected operating ranges.

Because the board used is still a prototype and not the final production board, some small parameters will change in the final board version. Small measured variations provide for a way to define acceptable variations during production testing of the actual boards.

3.1.1 Voltage

The voltages showed extreme consistency. They were all within five percent of the nominal value. Further, there was less than one percent difference between most reads, even among different days, firmwares, and fan levels. It will be interesting to see if the small deviations are consistent between boards or are products of the commercial components of the board. Notably, the oscillator voltages varied almost not at all.

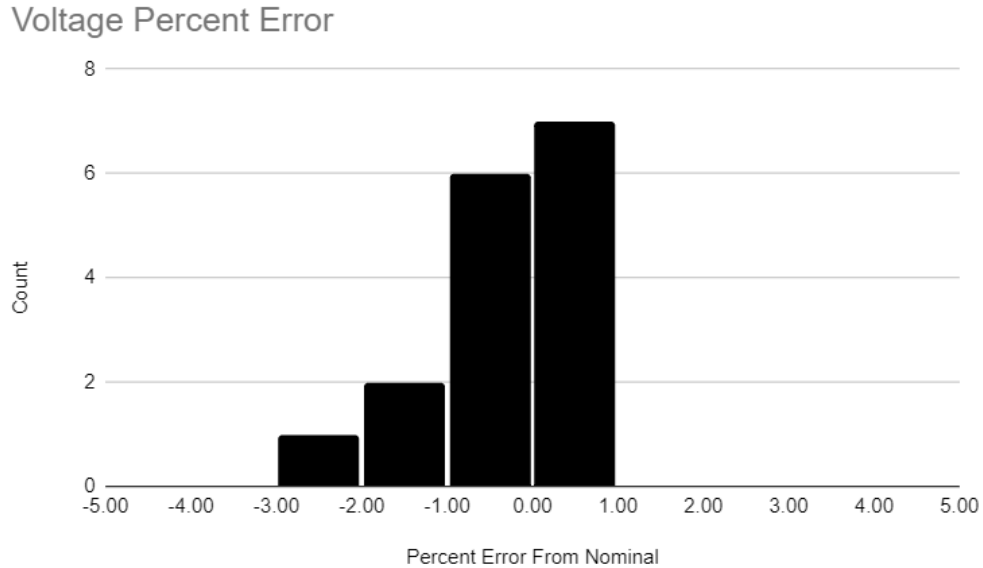


Figure 1: Histogram of Voltage Errors from Nominal

As shown in figure 1, the voltages were all within the operating constraint of 5 percent error. There were dozens of voltages measured and each one is a member of the histogram. Reading between runs and between days varied less than a tenth of a percent, and the values used are an average of reads from several days.

3.1.2 Temperature

Temperature of the FPGA is expected to vary based on firmware and the fan speed. The maximal temperature for safe operation is 90° C, although ideal for long term operation is 70 degrees. If any measurement went above that, the test would automatically shut down and the fan speed increase, although this was never observed. For full operation, the fan speed should likely be increased to level eight or nine since noise will not be an issue and temperature will also likely increase with a full ATCA crate of X2O boards.

Regular Temperature Stabilization

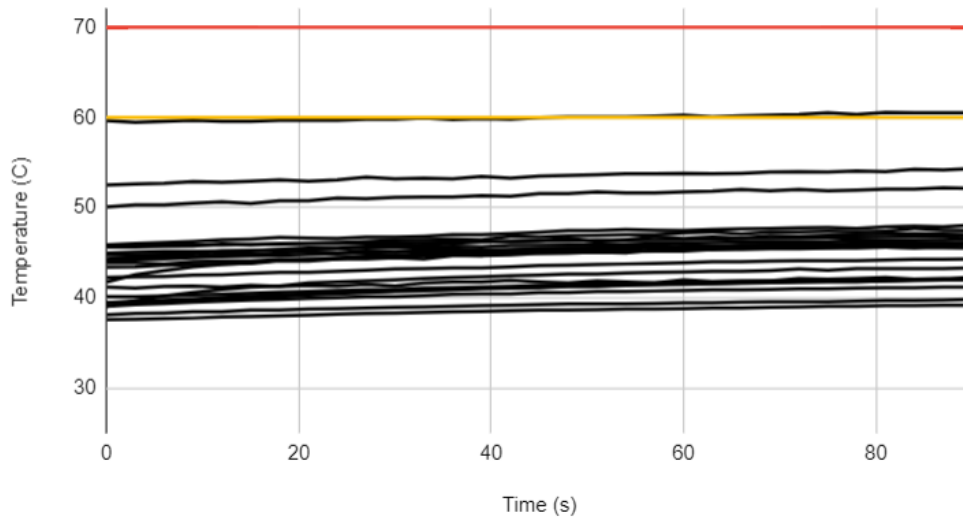


Figure 2: Temperature Stabilization with Regular Firmware

Temperature stabilization under the regular firmware was good. Each line in figure 1 represents a different temperature measurement on the X2O. There was one in the warning zone although it was stable. The warning zone is also substantially below the actual maximum operating temperature of around 80 to 90 degrees for long term operation, but it is an abnormality to take note of.

Stress Temperature Stabilization




Figure 3: Temperature Stabilization with Stress Firmware

Under a higher intensity firmware that drew more power and performed more computations, temperatures were a few degrees higher across the board. They still stabilized within acceptable ranges, and the points of measurement that were higher under the regular were higher on the stress firmware.

There is a broken set of temperature measurements on this X2O board that warrants more exploration. The 0.85 V VCCINT chip consistently read temperature values of 98304 Celsius across all of its six temperature points. This is being looked into by the board designers.

QSFP temperatures are another key metric to monitor throughout board operation. When there was not a full load of 30 optical transceivers, temperatures were in the high 20 to low 30 degree range. With a full load, the QSFPs in the middle were nearing the upper 30 degrees. Although the QSFPs can be safely operated for a long time in the 60 degree range, if they go above 50 for a long time that can be a warning indication and lower temperatures are better for longevity.

3.2 FPGA Programming

The FPGA Programming test was to program the FPGA with a firmware and then check the chip to chip connection. This test was not applied to the high power stress firmware, only the regular firmware intended for physics operation. The test passed if the chip to chip connection was working after programming. This test also does not have a separate log file for the test - its result the number of passed and number of total cycles is simply stored in the summary file. It is also possible to program the FPGA once as part of the software suite.

Approximately every 45 cycles, the chip to chip link would not be up. After waiting for three seconds then resetting, the chip to chip connection would be up. Simply waiting before resetting the link the first time did not work. This is an unusual characteristic, although it does not impede performance. It will be necessary to see if this behavior persists as new boards are created. For now, simply resetting twice has been acceptable.

The tools for developing and assessing the FPGA programming capabilities have been shown to work as required. Several slightly different firmware versions were used throughout the results period, all with good results. The FPGA was programmed several hundreds of times

and functioned after all of them. The main test script tests the FPGA programming ten times. Additionally, several overnight tests of programming the FPGA 500 times were run, all with no failures. Programming the FPGA takes about two minutes independent of the firmware. That is not of concern, as reprogramming will not occur very often during normal operation, only when something goes wrong, it is power cycled, or the board functionality needs to be changed.

3.3 Register Access Speed

There were two subtests implemented within the register access speed, one that simply reads a static value and one that writes and reads a random value. The pass criterion for both of these is that the average operation speed is less than 100 μs .

The average time for a static read was 38 μs and for a random read and write was 66 μs . The operating constraint is that a read operation be less than 50 μs and the read and write operation be less than 100 μs , so this is well within range. The specific register used for reading was simply the board id number. The random numbers were generated before the timing began to ensure the highest accuracy. It has been noted that reading a value takes longer than writing a value, perhaps because it has to send a request and receive data rather than just sending data over.

The test script does 100,000 operations for each of the subtests. In addition, several runs of one million and ten million operations were made, and their results were consistent with the 100,000 operations. Even 10,000 and 1,000 operations were within range, just a couple microseconds higher than average likely due to variation and the startup overhead. Minimal other processes were being run on the X2O at the same time as the register access test in order to prevent variation in timing from process switching. The results were also independent of the firmware version. The table summarizes the results of more than 30 tests of the register access speeds. The standard deviation is low, indicating that the access time is very consistent.

Table 1: Register Access Times μs

	Static	Random
Average	38	66
Deviation	1.1	1.5

3.4 Clock Frequency

On the X2O board, there are 60 reference clocks at 156.25 MHz and one at 40 MHz. The 60 reference clocks correspond to two clocks for each optical transceiver. On this prototype board, only 30 of the reference clocks were installed. In future production boards all 60 will be installed, and the software will easily adapt to that. The factory specification is the specification for the clocks, which is a tolerance of ± 25 ppm error in frequency. The tight frequency bound is important to have high time resolution from the detector to connect data to the appropriate bunch crossing.

The clock frequencies were almost all with one per million of the nominal value. The test software makes 100 reads of the clock frequency to produce good statistics. Most standard deviations for a clock frequency were less than three Hz, which is extremely consistent. In addition, a few runs were made of 1800 reads, and the results were consistent with the 100 read tests.

MGT 108 reference clock 0 notably had the most deviation from its nominal value. On one run, it had a value that was 3000 Hz off from nominal 23 times in 1800 reads. This is still an error of 19 ppm so it is within the specification, although it is about 20 times as much error as any other clock. On another run, it was off by 3 parts per million on one read, and on another 1800 read run it was all within the normal one part per million deviation. No other differences greater than 2 parts per million were observed across any clock on any run.

The below figure is the result of a 1800-read run. It shows a histogram based on the maximum deviance from the nominal value in a single read. It excludes MGT 108 clock 0, as that is an outlier as discussed. Clearly, most clocks never deviate more than one part per million from nominal, which is very well within the operating specification of 25 parts per million error.

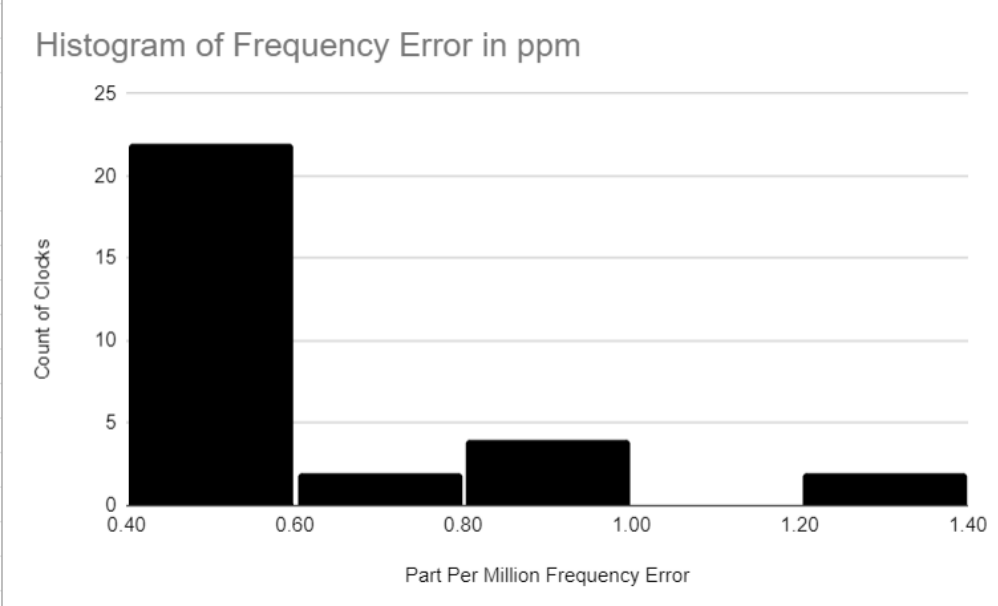


Figure 4: Maximum PPM error of reference clocks

4. CONCLUSION

The X2O board is a custom electronics board that will go in the Large Hadron Collider at CERN. As the beam will be upgraded to have higher luminosity, faster data processing is need. Additionally, new detectors are being installed in the Endcap Muon System of the Compact Muon Solenoid. Both the new detectors and the increased data flow necessitate new electronics to be designed and added into the data path so that more of the data can be captured and processed, leading to more physics information.

A software suite for comprehensively testing the X2O board was created. It is modular and portable, and easily adaptable to new firmware. A prototype X2O board was tested in a variety of operating conditions, and most of the parameters were well within the operating specifications.

The testing implemented checks on many features, such as QSFP temperatures, the FPGA programming, and reference clock values. Preliminary results show parameters were stable and within acceptable operating values. When a full board is created with the rest of the reference clocks and voltage readouts those will have to be evaluated.

The output of the test goes to log files as well as the console. This allows for a quick determination of pass or fail on specific tests as well as globally. The log files form a database to be integrated with CERN hardware information. The data logs are timestamped and each run number has multiple files associated with it, one as a summary and overview and a few for the full results of some subtests.

4.1 Future Research

Next steps would be a final design review of the x2o and moving into production. Once more boards are produced, the software suite will be used to test them. This will allow for better characterization and failure detection. Also having a full board with the missing clocks and voltage measurements will require some minor software changes to integrate and validate those measurements.

This work can also be applied to testing X2Os that will be used later in the data path as track finder and muon trigger algorithms. Although there is different purpose, firmware, and optical configurations, the physical data baseline should be similar. The codebase for testing will also be useful in establishing capabilities and determining suitability for physics applications under that use case.

REFERENCES

- [1] O. Aberle *et al.*, *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*. CERN Yellow Reports: Monographs, Geneva: CERN, 2020.

- [2] G. L. Bayatian *et al.*, *CMS Physics: Technical Design Report Volume 1: Detector Performance and Software*. Technical design report. CMS, Geneva: CERN, 2006. There is an error on cover due to a technical problem for some items.

- [3] J. G. Layter, *The CMS muon project: Technical Design Report*. Technical design report. CMS, Geneva: CERN, 1997.

- [4] M. Glazewska and M. A. Konecki, “Level 1 Muon Triggers for the CMS Experiment at the HL-LHC,” tech. rep., CERN, Geneva, 2022.