High Performance GPS Disciplined Oscillator and Distribution Amplifier with Network Time Protocol Support

Alex Ferro
Dept. of Electrical and Computer Engineering, University of Utah
alex.ferro@utah.edu

***Abstract*— The world runs on time. Modern communications systems require clocks to match among all actors in a system, in many cases to better than 16 parts per billion [1].**

The core of this project is a global positioning system disciplined oscillator expected to be capable of sub 1 part per billion accuracy. Also included is a distribution amplifier system to provide multiple outputs to other equipment and an NTP server to provide local time information. To achieve this level of accuracy, the project was carefully designed to minimize interference from both internal and external sources corrupting the outputs. Due to the time constraints of the project, while the hardware design was carefully done, the software sections did not get completed and the control loop is non-functional. However, preliminary testing has shown that the frequency error goals are not unreasonable with the current system design.

I. Introduction

"Does anybody really know what time it is?
Does anybody really care?" — Chicago

It used to be that all the vast majority of humanity needed as a notion of time was that which could be provided by a glance at the sky. In our modern world this vague notion of time is not all that appropriate. Instead, that much of our modern world depends on time comes as a surprise to many. This need is hidden from view, deep in the inner working of the things we take for granted, the Internet, the cellular phone system, or the GPS system that keeps us from getting lost on unfamiliar trips. In practice, the stringent requirements of many of these systems require the use of atomic clocks in communications installations as nothing else meets the tolerances. However, atomic clocks are not the only way of achieving these tight tolerances. A global positioning system (GPS) disciplined oscillator (GPSDO) can also meet these specifications, while providing automatic self calibration and aging compensation in the natural course of operation. That said, to fly somewhere in a plane does not require one to be able to build a 747. Nor does the use of a cell phone require the ability to understand how one works. However, to those who create these systems that do depend on time, an accurate notion of time can be critical for testing or implementing system functionality.

To provide time and frequency references, there are several different types of reference. Some are cheap, and good enough for low specification use only, while others cost several thousand dollars and up and come close to meeting the very definition of the second itself. This project revolves around a middle ground, that of the disciplined oscillator. In short, a disciplined oscillator is one whose undesirable characteristics are adjusted and suppressed with the use of a better but directly unusable source clock. Currently GPS is commonly used to provide the high accuracy reference clock. This same approach has been taken in this project. The targeted accuracy is 10 parts per billion (ppb) accuracy with 1 ppb expected and sub 500 ppt possible. Preliminary testing has shown that the raw performance of the oscillator stage is likely to exceed 400 parts per trillion (ppt), over the time period the the control loop will be unable to correct.

Besides the high accuracy oscillator, other additions like the built in distribution amplifier and the network time protocol (NTP) server are provided as a number of instruments can use clock references and it can be important that all machines share an accurate time reference.

II. Background

There are several different methods of generating an accurate and tight variance clock source. Each has advantages and disadvantages with respect to price and performance.

A. Quartz Oscillator

A cheap quartz frequency reference, like those used in a watch or many computing systems may be 100 parts per million (ppm) or worse from true, but these references cost only pennies. Even at the upper end of the scale, accuracies of 10 ppm at room temperature are about as tight as is practical for basic quartz oscillators as thermal effects begin to affect the frequency significantly. Above this point, there are a number of compensation schemes that correct to a degree this drift, but none of them are as effective as placing the oscillator inside a temperature controlled chamber.

B. Oven Controlled Crystal Oscillator (OCXO)

A significant upgrade from this is to use a high tolerance frequency reference of this type, and place it in a heated oven to well above ambient to remove the effects of frequency drift. These usually cost two or three orders of magnitude more than a basic crystal. However, for that

price, accuracies of 0.1 to 10 ppm are common, with a change with ambient temperature from 2 to 500 ppb and a variance better than 10 ppt [2].

C. Current Atomic Clocks

The definition of the second is 9,192,631,770 transitions of a cesium atom at 0 K. A clock that measures this transition is by definition accurate, although measurement uncertainties become the bigger problem at that point. In the United States, the holder of the clocks that define the second is the National Institute of Standards and Technology in Boulder, Colorado. Their flagship clock, NIST-F2, has a current measurement uncertainty of 1×10^{-16} [3]. Their previous clock, NIST-F1, has an uncertainty of about 3×10^{-16} [4] and is planned to be run in parallel with NIST-F2 concurrently for the near future [3].

In the commercial sector, atomic clocks are also available. Most commercial standards use rubidium instead of cesium, both for cost reasons, and because the use of rubidium allows simpler microwave cavities and is less susceptible to magnetic fields [5]. These commercial grade standards are accurate to within 50 ppt, and have short term variance near 2–10 ppt. However, long term accuracy is limited by aging effects and commercial standards may drift on the order of 100 ppt per year [6].

D. Disciplined Oscillators

Clock sources in this category have multiple components, consisting of a low accuracy, but tight variance oscillator like an OCXO, and a high accuracy, but large variance or less usable reference signal. In the past, many frequency standards of this type used the 60kHz carrier of WWVB (a NIST broadcast radio station) as a reference signal as NIST controls the frequency to an accuracy better than one ppt. A product like the Spectracom 8165 can lock to WWVB with an accuracy of one ppb and a stability around 100 ppt [7]. Unfortunately, since these systems were manufactured, WWVB changed their modulation to incorporate phase modulation, which causes this class of standard to fail to lock on properly.

Modern general purpose[1] disciplined oscillator systems are usually based around the GPS signal. Since GPS works by solving a system of equations with the current time as a variable, a GPS receiver knows the current time. Additionally, GPS satellites carry atomic clocks onboard and that accuracy is available to the receiver. This accurate, but jittery clock is then used to tune an inaccurate but low variance clock to match. A clock accurate to <100 ppb can be easily created, and well designed units can achieve a short term accuracy of about 100 ppt, with a longer term 24 hour average of <10 ppt. This is accompanied by the tight stability of the internal OCXO, which provides short term stability of <30 ppt including the effects of the control loop [8].

[1]Excluding data stream clock recovery and similar systems.

III. Proposed Work

The primary purpose of this project is to be a high accuracy and high stability local oscillator. As a secondary purpose, a distribution amplifier is integrated capable of driving several external loads, allowing more that one device to receive an accurate clock. As a related goal, not all devices can make use of the same frequency clock, so a phase locked loop (PLL) circuit is also integrated into the circuitry to allow variable output frequencies with high quality. Finally, a local NTP server is planned as a function of the device to make use of the accurate time information from the GPS network as a reference for networked computers and devices with regular clocks.

See Figure 1 on page 7 for a high level block diagram of the system.

A. Rationale

While the correct time is important to many different disciplines, those who build, design, and test communications equipment are acutely aware of the role that accurate frequency measurement plays in the modern world. Systems as critical to our way of life as the cellular network require clocks accurate to between 16 ppb and 50 ppb for large area base stations [1]. Additionally, the GPS constellation that many of us depend on for getting to unfamiliar places would not work without accurate and low variance timing. For this reason, both GPS satellites and many cellular base stations carry atomic clocks onboard to prove high quality sources of time. This need for high accuracy references in modern systems prompts the focus on sub 10 ppb frequency error.

Even in laboratory settings where the absolute frequency is not important, the variance of the frequency may be important. It is also important that multiple instruments which are involved in a process all agree as to frequencies. It may be fine that a 1 MHz signal is 1% off, but if one device is reading 1% high, and the other device is reading 1% low, then the potential exists to get stuck looking for a problem that only exists in the measurement tool. In that case, having both devices agree on the frequency is much more important. The synchronization problem provides a reason to have several outputs to drive a whole lab full of equipment that needs to know what the local “Standard Time” is.

Finally, an NTP server is included to handle the other devices and probable general purpose computing devices that need to know time and a format that corresponds to the way humans view time in days, hours, minutes, and seconds.

IV. Completed Work

Due to some time constraints and a busy schedule, I was unable to complete everything I had planned on as a part of this project. In particular, I was unable to create a working control loop to adjust the frequency of the oscillator. I also dropped the NTP server from the project as I was having trouble with the network interface

in the latest FPGA build that had more important portions fully functional. If the network interface had been working, I had everything all ready to flip the switch on NTP.

Despite these problems, much of the project was functional. My testing shows that the low noise considerations I designed in were effective, and that I am able to get the designed performance out of the critical portions of the design. In particular, even without a control loop, I was able to achieve 400 ppt accuracy by setting the clock once against a borrowed atomic clock, and then returning 4 hours later to measure the offset. Based on that result, I am confident in a well designed control loop being capable of the same results or better over that time period. As for the ancillary features of the project, all of those I needed are functional. My board worked properly and I can measure the power consumption of any individual piece to determine ensure the expected power consumption specifications are met. All of the independent power supplies and point of load regulators are working, and the digital clock tree is providing the system clocks that everything needs.

A. Oscillator

The design goals for the oscillator subsystem were twofold. First, it needed have very stable frequency characteristics over a time period of up to a few hours. Second, I needed to have the ability to adjust the frequency output over time to counteract initial and long term error. These two criteria dictated the use of a VCOCXO (voltage controlled OCXO) as the primary frequency source. The VCOXCO is then adjusted by an 18-bit DAC under processor control with the information from the frequency measurement system.

To minimize undesirable effects on the frequency like temperature coefficients and derived effects, I had originally planned to place the whole subsystem inside of a heated oven at a temperature above ambient. However, my later research found that it would cause more trouble than it would gain in stability, and that instead, I should just enclose the area to minimize thermal transfer. The enclosed area will entirely prevent drafts, and slow down the effects of ambient temperature on the frequency to help the control loop remain locked.

B. Frequency Measurement

To correct the frequency of the oscillator, I am using the traditional frequency counter technique of measuring the number of cycles in a gate interval. For this method, the frequency being measured is the local oscillator, and the gate time is generated by a GPS receiver via its pulse per second output. The results of the last several seconds can then be accumulated and used to adjust my clock source.

C. Distribution Amplifier

As an instrument, this will not be all that useful if it is only capable of driving one external piece of equipment. This is at its most useful when many to all instruments that measure time in a lab are all getting their frequency reference from a single box. Thus, a distribution amplifier stage is present, such that several different instruments may be driven with a single clock without interference from each other. This makes a device like this project useful for a whole lab, not just a single instrument.

D. Phase Locked Loop

To provide additional output agility, some of the external outputs have been connected not to the raw oscillator output, but are instead connected to a PLL frequency multiplier/divider chip. This allows a programmed frequency change such that an instrument that instead of a 10 MHz clock, needs a 27 MHz clock, will still be able to be used with this equipment. The control interface for these PLLs is planned to be over the on board Ethernet link, with status and outputs available on the front panel.

I had to ditch the externally facing PLLs for the demo due to time constraints. The hardware functionality is present, but the software was not ready.

E. NTP Server

Another useful function for a lab is synchronized time keeping at a higher level than just frequency. Systems that log data with timestamps also have a need for the current time. Again, if one instrument records an event happening at onne time in its timestamped log and another instruments records a related event in its log, without synchronized timestamps, it can be difficult to determine if the two systems saw their events at the same time, or not and what the delay was between them[2]. As the GPS module I am using as my frequency reference also outputs the current time, it is fairly easy to use that time information as the source for an authoritative NTP server. I planned on using one of the well known NTP servers written to run on Linux for this as they can directly use the information from the GPS. However, due to some temporary problems[3] near the end of the semester, I had to remove the NTP server for the system functionality as of this paper.

F. Power Distribution and Supply

Since this system has very tight tolerances, the power supply was carefully designed to minimize the influence of noise of the frequency output. To that end, all of the critical power supplies are heavily noise filtered and derived from low noise regulators separate from those used by the non-critical systems. Additionally, the power domains are carefully laid out on the board to minimize coupling between critical and non-critical areas.

As a second consideration, since long term information will be monitored about the relative frequencies of the

[2]Technically this problem is better solved with IEEE1588 (Precision Time Protocol), but that was out of scope for this project and NTP is fine for millisecond alignment.

[3]I broke the Ethernet port configuration in Linux when I rebuilt the system. It was working in earlier builds of the system.

two systems, it is important that momentary power glitches or short outages do not negatively impact the control loop function. To achieve short term holdover, a battery backup system has been implemented to keep the system running during an interruption of around an hour. This will allow the system to be robust around momentary issues with power, including a tripped breaker or a momentary power glitch.

V. Hardware Implementation

To meet these goals, I realized very early in the process, that I would need to make a custom circuit board with all of the different parts on it and encapsulating in its design the low noise principles needed to meet the specifications I was targeting. The overall design and hardware creation took up most of the summer and the first month of the fall. This time was not wasted as I had plenty of time to go through my design and make sure everything was right, and the time to implement each feature in a non-hurried fashion on the PCB. I also made time to have design reviews completed with the engineers at my work. The time spent on design and implementation and the importance I placed on design reviews meant that I had great confidence in my hardware to function once it was built.

In addition to the primary features of the project, during the design stage I thought about my other goals of the project. One of the most important goals that have driven this project forward was that of desiring a useful piece of equipment in my lab once completed. As I was thinking about that goal, I realized that if I used the extra functionality given by the FPGA, I could make a device that could fill more than just the purpose of being a clock in my lab, for much less than it would cost to make a standalone tool for the job. With that in mind, I added some high speed digital to analog converters to the board that I can drive at full rate out of the FPGA. When the software is completed for this, I will have the ability to push up to four channels of high bandwidth analog data from this project. I then took all of the remaining pins left over and connected an expansion header with a number of extra signals, a few power rails, and the shared communications bus on the board. Both of these features will allow this project to provide a greater utility as I continue to make use of my lab.

As a parallel task to the PCB, I enlisted the help of one of the mechanical engineers at my work to design me a sheet metal case to put the project in. I placed this as also a high importance task as it provides a way to protecting the board while I work on it, but even more importantly, since I want this to be a functioning piece of equipment in my lab after I finish it, I need a case to protect it and provide mounting locations. Luckily the timeframe for the housing aligned with the school schedule and I was able to have a housing in time for the demonstration.

Figure 2 on page 8 is an annotated view of my project, with the major sections highlighted.

A. Control

To run the system, I chose to use an Avnet PicoZed module. This module provided most of the useful pins of a Xilinx Zynq on its board to board connectors and freed me from having the problems of working with fine pitch BGA components and the complexities of memory system routing. The choice of a Zynq was motivated by the combination of a powerful FPGA, with the Cortex-A9 processors and the ability to run Linux. Running Linux meant that some portions of the system stack I had to work with got far more complicated, but the other portions were far simpler, and I felt (and still feel) the advantages outweighed the disadvantages. Also in the control section, are the external ports for the Linux system, USB, and Ethernet, as well as the other important piece, the GPS module. This GPS module was fairly cheap, easy to get, and even better, has a reputation for better performance than other devices in a similar price bracket.

B. Oscillator

The oscillator section (marked OSC) is the most critical section of the whole system. I have shown it in this picture without the thermal shielding that would normally cover this area. In this section is the low noise power supply, and digital to analog converter used to adjust the output frequency of the VCOCXO (silver with the black label). I raised this up from the rest of the board to both isolate it from the noisy environment of the main board, as well as to allow the thermal shielding to isolate it from the rest of the system more effectively.

C. Output

The output section consists of several buffer amplifiers that take the 10MHz from the oscillator and provide a local reference for the rest of the board, but also serve as the distribution amplifier to drive the cables that would ultimately go to other units in the lab. Here, the focus was also on low noise, but it was less critical than the oscillator that everything be thermally stable.

D. Power

The power supply requirements of this project were complex. I needed to take an input supply of 12V and convert it into many high current outputs ranging from negative 5 volts to positive 5 volts, all with low added noise. To meet this goal, I turned to heavily filtered point of load post-regulation with switching converter frontends. Each supply has its own current monitor and I produce a supply for every subsystem that is unique and not shared with any other section to avoid cross contamination with internally generated noise. The power subsystem also is designed to keep the batteries charged and switch over to running off of them in the event of a power loss.

E. PLL

This portion of the board contains the phase-locked loop chips that create most of the onboard and external clocks needed. These PLLs were chosen based on the output frequencies needed, and their performance specification to create low jitter clocks at any frequency needed. The external variable frequency clock outputs are from this section of the board.

F. Future Expansion

The future expansion section of the board is composed of the functionality I added during the hardware design stage that is not part of the actual project. This section exists to make this project even more useful in my lab.

VI. Results

I was not able to finish the proposed functionality of the project by the end of the semester. All of the hardware has been completed, but the software is lacking several important features.

A. Problems

On the software front, the biggest missing feature is the control loop to adjust the frequency of my reference oscillator to match the GPS. The reasons for this are twofold. First, I failed to realize that the approach I was using for measuring the frequency was not suitable to use for measuring the small offset that the system actually exhibits. The jitter from this method was far higher than the jitter inherent in what I was measuring. This compounded the software problem I was also having with the control loop. Since the jitter was excessive, when I was looking at the results, I kept thinking that everything was fine as I saw the results changing, never mind that they would always change back, and I would go back to looking at my code to find why it wasn't working, rather than realizing the root cause of poor data.

I also needed to scrap the NTP server as a working portion of the project for the demonstration. This was simply because of a problem with the lastest Linux image I created. I had broken the network port's configuration in the course of trying to get the prerequisites for the control loop to function in place. Once I fix the network port in Linux, it will not be any trouble to get the NTP server running, as all I need to do is flip a switch in the software to get it to run on my board.

The final missing feature to the external PLL outputs. I have every reason to believe that the external outputs will be functional, as I can control the chip that would generate them, but I did not have time to configure the PLL chip to produce outputs for testing.

B. Functional

The large majority of the hardware has been tested as functional. I have not extensively tested the portions of the board in the future expansion category, or the external PLL, but everything else seems to be working at a basic level.

I did not have any unexpected trouble with the process of getting the PicoZed to boot Linux. I have also confirmed that Linux is able to talk to everything on the board that it should be able to.

The most interesting result is the system performance I am getting with the oscillator even without a control loop. I decided to use an oscilloscope in X-Y mode to measure the frequency. The advance of this method is that by measuring the rate of rotation on the screen, the frequency error between the two signals can be determined. From the frequency error, one can divide that by 10MHz to get the proportional error. The other advantage with this method is that it provides a visual aspect of the function of my project.

To determine the frequency error my project is capable of without a working control loop, I manually tuned the DAC by eye comparing it on the X-Y plot with a frequency reference. I then observed the plot while I worked on other projects in the lab for four hours. In that time, the error in the system had not drifted more than 400 ppt relative to the reference. As the reference clock I was using was an ExactTime 6000 built around a rubidium atomic clock that Dr. Neal Patwari loaned me, I am confident that the majority of that drift is due to both my initial setting error and the inherent drift of the project. Based on this error, I am expect that I can achieve this or better over the same time period when I have a functional control loop.

VII. Future Work

Unlike some projects, most of this project actually would remain the same if I were to do it again. Much of that is because I spent far longer than I should have in the hardware design stage and therefore created a system that is built correctly and with the functionality that I really wanted. Additionally, a unit with the same purpose as what I proposed has been on my list of projects for a few years and the senior project gave me a reason to actually get it done. Finally, during the hardware design stage, I was left with the conundrum of what to do with the extra pins on the PicoZed that I did not need for core project. To chew through pins, I added two high speed DAC chips that I could use to synthesize both other pure frequencies, but also for wide bandwidth analog output from the FPGA for either waveform generation or radio transmission baseband. Then, the pins that were left over after even that, were routed to an expansion header, that should I, down the road, need some fast I/O and a clock, I could just build a small module that would do what I need, without a full high speed FPGA, just to handle one one task.

I plan on actually finishing this project over the next few years, starting with the control loop that I was not able to get done. Once the rest of the functionality proposed has been completed, I intend on finishing the high speed DAC cores and creating a waveform playback mode as well as potentially some possible modulations for data

on top of a carrier. Finally, to properly interface with everything, I will need to write a Linux kernel driver, as userspace code is not the right place to handle logic that is coupled this close to the hardware. This will be good skills as well as I will then have experience with Linux kernel driver creation, as well as proper communication between userspace and kernelspace as the system does not fit the requirements of either one without portions in the other.

I see this project as not only being a project to teach me a number of new skills and keep me busy, but also as a useful addition to my bench of equipment.

VIII. Summary

The creation of an extremely high accuracy and low variance frequency standard will require the primary oscillator to be dynamically adjusted to match the atomic clocks that are a part of the GPS system. To be a useful addition to an engineering lab, components have been added around the core of the project to drive more than one external instrument and to provide more than one frequency, but the core is the oscillator itself.

While GPS disciplined oscillators may be bought readymade from a number of manufacturers, this is an educational project that not only will I learn about the challenges of system design down at the ppb level, but I will also have a useful addition to my lab bench that will give me the ability to measure frequency to a tighter degree than any current standard I own.

Even though I had to strip out some of the software functionality to meet the deadline, this was a project that pushed many of my skills to the edges of what I knew, pushing me into some newer territory, and teaching me more about low noise design, high speed routing, and will be a good testbed for further learning about embedded Linux, FPGA, and analog design.

IX. Acknowledgements

- Dr. Cynthia M. Furse — For convincing me that my enthusiasm for this project should not go unpunished.
- Dr. Peter A. Jensen — For convincing me that my enthusiasm for this project should not go unpunished.
- Dr. Neal Patwari — For loaning me the ExactTime 6000 that I used to confirm the performance specification against.
- Dr. Ken Stevens — For providing the opportunity to learn how much time interested me and for overseeing this project.
- The engineers at Aerel Design — Donated many hours of their time to help with design reviews, ideas, and to design and handle the ordering of the sheet metal housing that I designed my project to fit.
- The engineers at rfRanging — Answered many of my questions that my coworkers were unable to answer, tested coax cable samples for me before I ordered larger quantities, and nost importantly, allowed me to use their large reflow oven to solder all of the parts. Even with all of the support from my workplace, without the use of a large reflow oven, this project would not have been possible to assemble without several times more work fixing bad solder joints.

References

[1] M. Weiss. (2012) Telecom requirements for time and frequency synchronization. [Online]. Available: http://www.gps.gov/cgsic/meetings/2012/weiss1.pdf

[2] *10 MHz high-stability ovenized quartz oscillator*, Stanford Research Systems, 4 1995.

[3] American National Standards Institute. (2014, 4) NIST launches a new U.S. time standard: NIST-F2 atomic clock. [Online]. Available: http://www.nist.gov/pml/div688/nist-f2-atomic-clock-040314.cfm

[4] ——. (2013, 2) Primary frequency standards. [Online]. Available: http://www.nist.gov/pml/div688/grp50/primary-frequency-standards.cfm

[5] United States Naval Observatory (USNO). (2006, 3) The USNO rubidium fountain project. [Online]. Available: http://tycho.usno.navy.mil/clockdev/RubidiumFountain.html

[6] *Rubidium frequency standard with low phase noise*, Stanford Research Systems, 3 2013.

[7] *WWVB Disciplined Oscillator*, Spectracom Corporation, 1 1995, rev. 2.2.

[8] *Epsilon Clock Model GPS Disciplined Oscillator*, Spectracom Corporation, 6 2008, rev. D0.

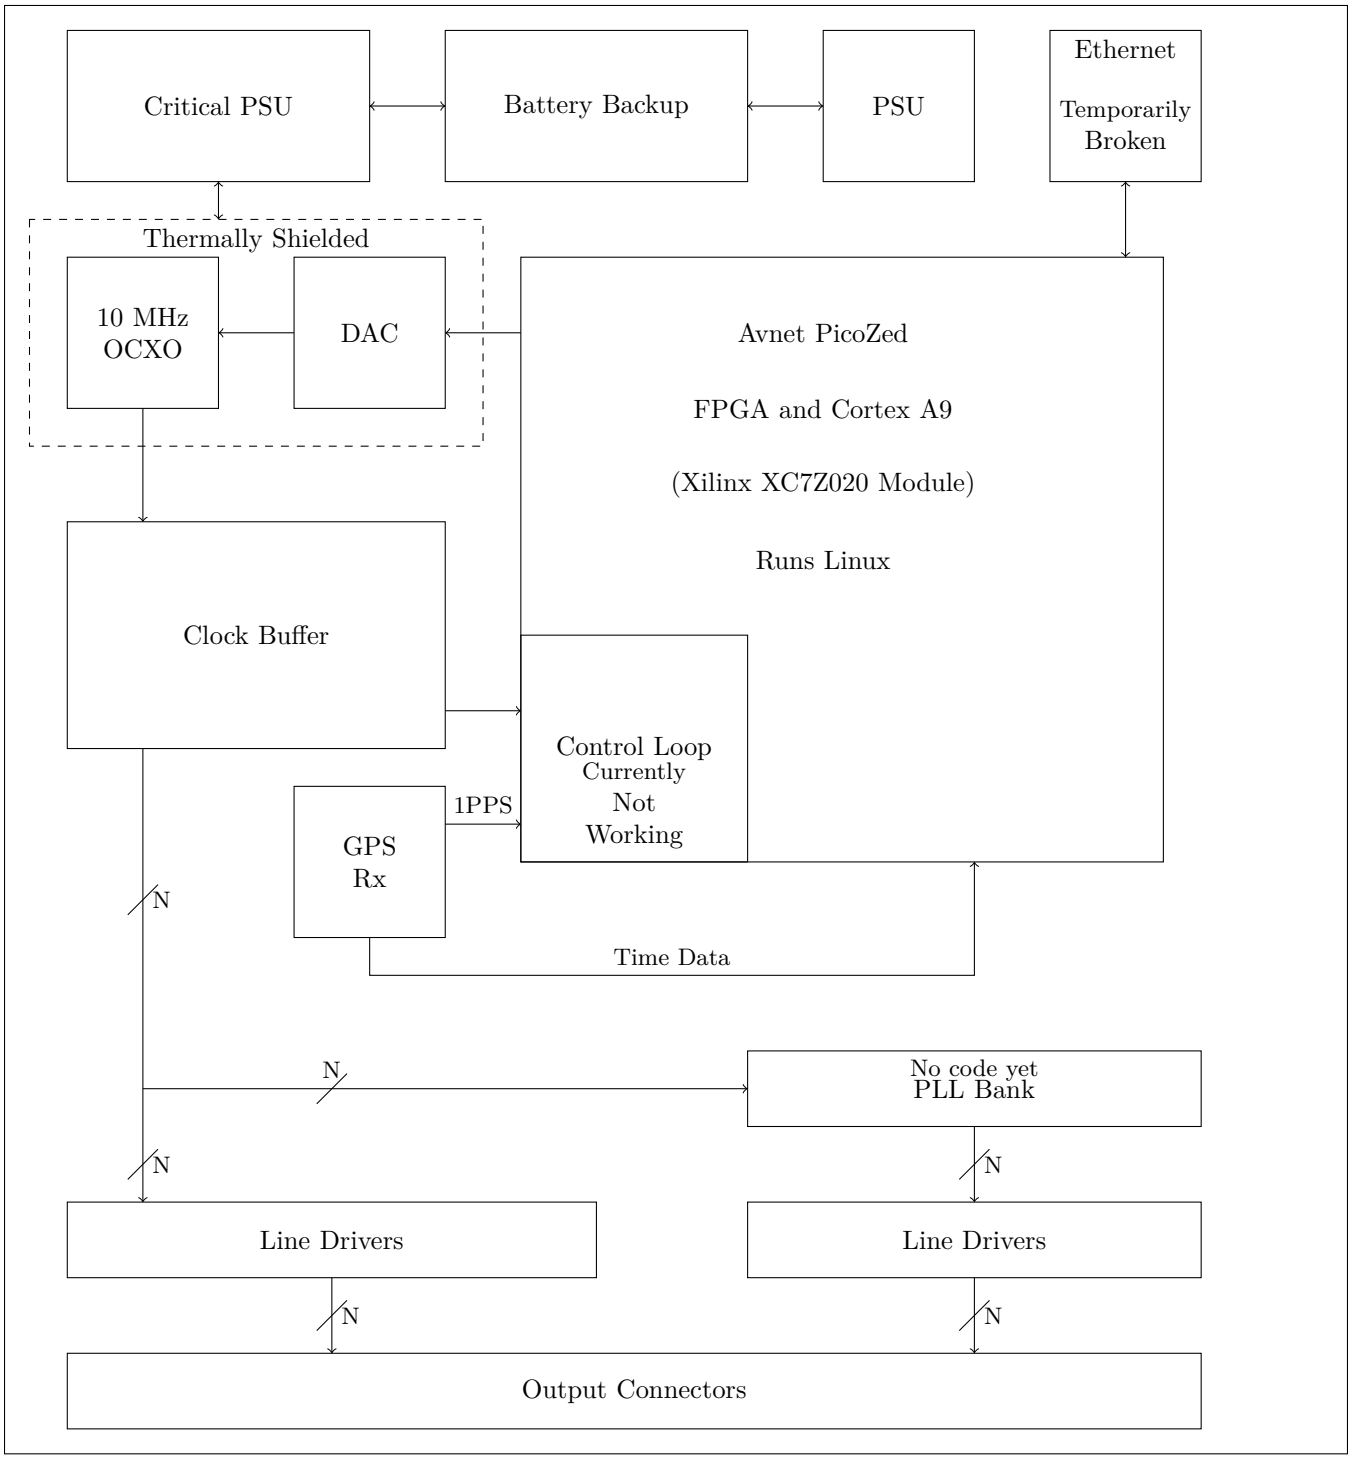

Fig. 1. High level system block diagram

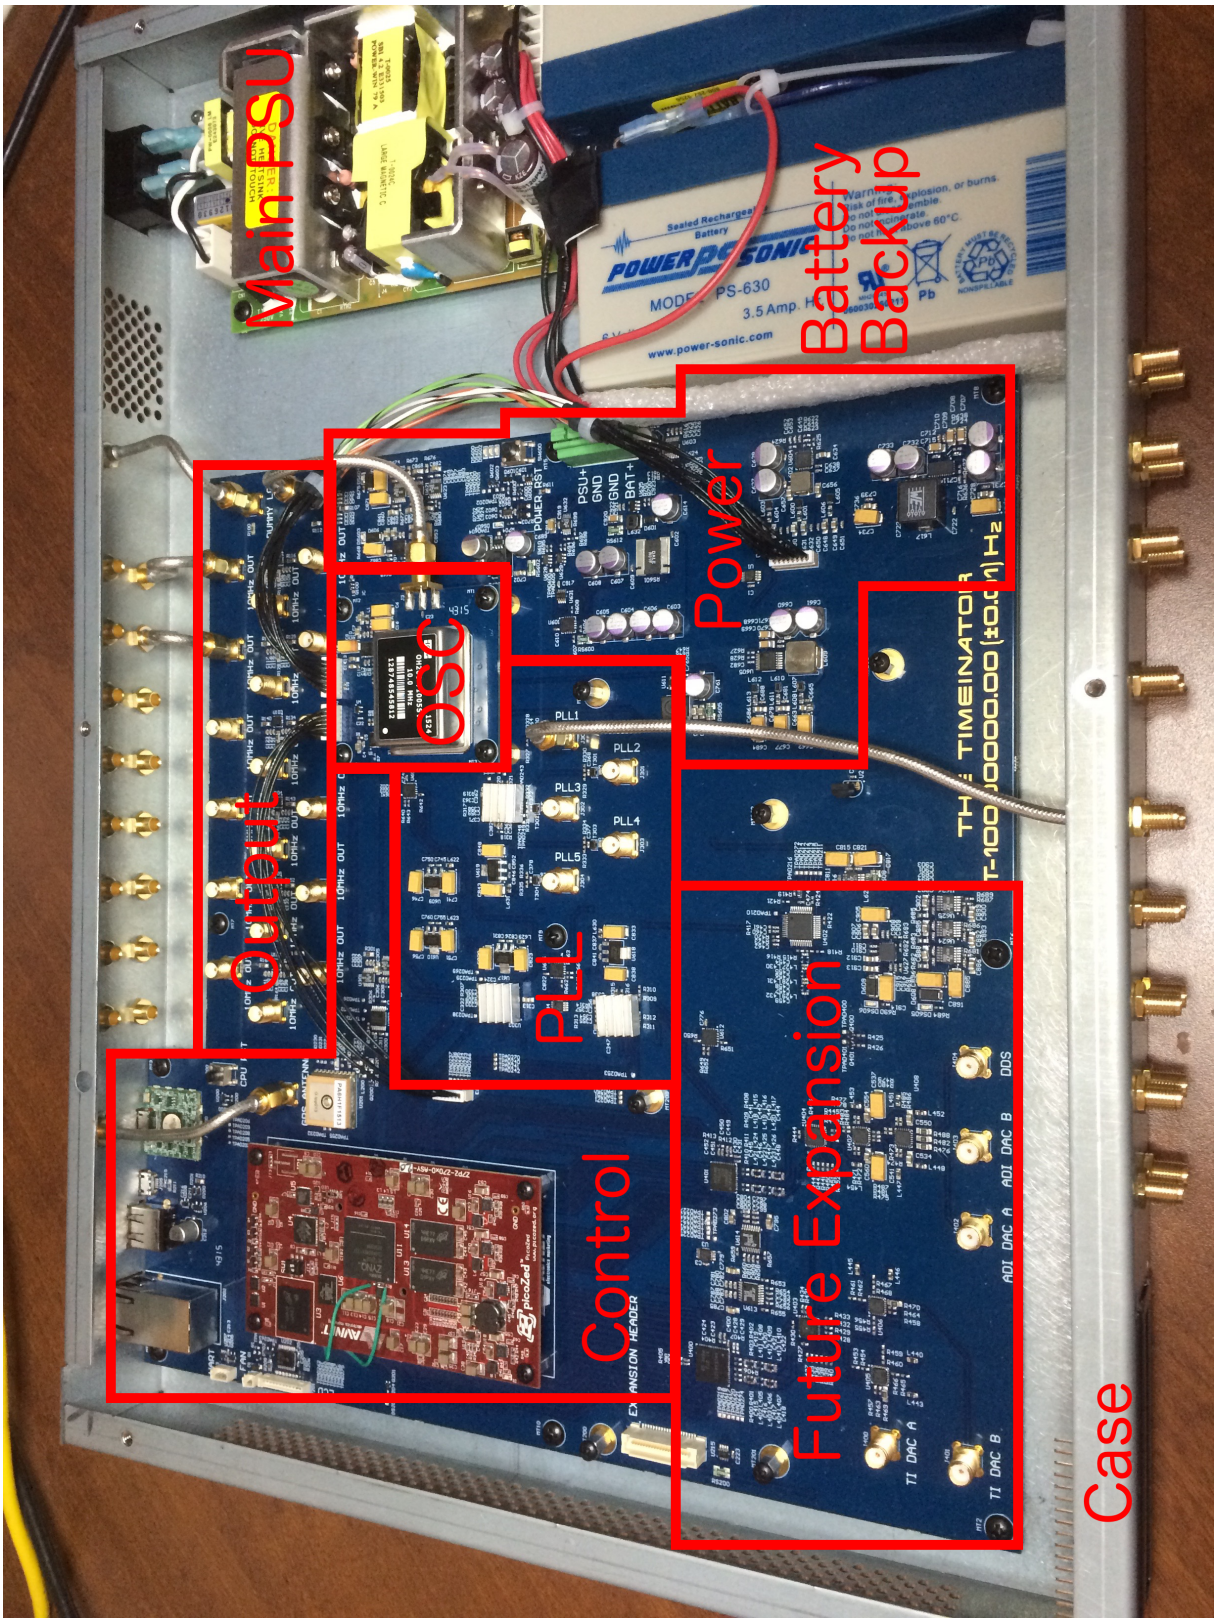

Fig. 2. Annotated view of the project