

A Second Look at Chip Scale Atomic Clocks for Long Term Precision Timing

Four Years in the Field

Alan T. Gardner and John A. Collins
Woods Hole Oceanographic Institution
Woods Hole, MA, USA
agardner@whoi.edu

Abstract—We compare the performance of 38 Chip Scale Atomic Clocks manufactured by Microsemi Corporation (originally Symmetricom Inc.) used in a fleet of Ocean Bottom Seismographs. CSACs have been available commercially since 2011, promising accuracy close to that of conventional “low power” Rubidium oscillators at one fortieth the power of the most efficient Rubidium. Ocean bottom seismology provides a particularly challenging requirement for timing: millisecond level or better accuracy over a one year interval or longer. While results from an early batch of CSACs have been largely positive, later units have not performed as well. The CSAC specifications have changed, reflecting a decrease in reliability and accuracy of more recent units. We will discuss the impacts of the degraded aging for long-term autonomous experiments. For the majority of applications this change is not very important. However since aging-related phase error increases with the square of time, applications requiring long-term accurate phase such as ocean bottom seismology may be significantly impacted. We will also compare the CSAC to the SISMTB Microprocessor Compensated Crystal Oscillator, designed and manufactured by Seascan, Inc. using lab testing to determine their performance relative to the range of CSAC performance.

Keywords— Underwater timing, precision timing, low-power timing, chip-scale atomic clock, microprocessor-compensated crystal oscillator, ocean bottom seismology, ocean acoustics, ocean observatories, oceanographic instrumentation.

I. INTRODUCTION

When Symmetricom Inc. (now Microsemi Corporation) commercially released the Chip Scale Atomic Clock (CSAC) in 2011 it changed the landscape of low-power precision timing [1]. It promised performance better than the best crystal oscillators for more than an order of magnitude less power than a compact Rubidium or ovenized crystal oscillator in a very small, affordable package. Many applications in Oceanography stand to benefit from such a device, as we often must maintain accurate time over long periods without access to GPS and with limited power and weight budgets. A good overview of the potential benefits is presented in [2].

Ocean Bottom Seismographs (OBS) have one of the most stringent requirements for timing. They are often deployed at deep sites for intervals of around a year. Ideally the clock in an OBS would be accurate to within 1-10 milliseconds at all times throughout the year. This suggests a clock stability of roughly $3e-11$, a challenge even for many Rubidium oscillators. There

is a push for even longer duration experiments (2 years) and higher sample rates (100 Hz), both of which drive even greater accuracy requirements. In practice clocks are typically used with a few orders of magnitude lower stability and their phase error is corrected in post processing by measuring the offset to a stable oscillator (typically GPS) at the beginning and end of an experiment to determine the average frequency error of the clock. This can work reasonably well, but it relies upon the frequency error of the clock being constant [3].

The Woods Hole Oceanographic Institution operates nearly 100 Ocean Bottom Seismographs as part of the U.S. National OBS Instrument Pool. Our instruments currently utilize one of two models of clocks. The majority have SISMTB Timebases manufactured by Seascan, Inc. of Falmouth, Massachusetts [4]. These are microprocessor compensated crystal oscillators requiring only 5 milliwatts of power. We have used them for 15 years in a total of over 275 clock-years combined on-bottom time. In addition to our extensive records from these deployments we also have a total of over 65 years of combined lab test-time monitoring the actual frequency error minute to minute in different temperatures. Based on these data, and backed up by ambient noise techniques [5] we can get a good idea how well the Seascan perform. While they have enabled a generation of seismologists to make many discoveries about the Earth, it is clear that the Seascan drift is not linear, typically leaving tens of milliseconds of residual timing error after applying standard linear corrections to a one-year-long experiment. This is an exemplary accomplishment for a clock drawing only 5 milliwatts, and until 2011 was better than virtually all clocks drawing less than a few watts.

When the CSAC was commercially released we were among the first customers, installing them in 15 OBS. Since then 6 new instruments with CSACs were commissioned in 2015. Between these 21 instruments our CSACs have seen a total of 50 clock-years of on-bottom time. We have recorded an additional 40 clock-years of data from the clocks in the lab to monitor their performance in real time. We have had a total of 40 CSACs come into our lab in several groups since the initial prototype in 2010.

Our experience with the first production run of 20 CSACs in 2011 was mostly very good, although as we did have four clocks which did not meet stability specifications within the first year. The remaining clocks from the original group

continue to meet expectations. However, there have been a number of issues with more recently purchased clocks. Most immediately troubling is that we have seen three clocks fail. A more pernicious problem is that Microsemi has relaxed the aging specification by an order of magnitude. A number of the more recent clocks that we have received fall in this expanded aging spec, which renders them far less suited for OBS work than the earlier batch. We will examine both of these issues here, and compare actual CSAC performance to Seascan performance.

II. CSAC TIMING PERFORMANCE

All of our CSACs are installed on custom PCBs that provide a number of features specific to our requirements [3]. This allows us to monitor the phase error of every CSAC which is typically logged once per hour as long as the system has access to a GPS derived clock signal in the lab or on the deck of a research vessel. We use these phase errors to calculate average frequency errors in order to monitor the clock performance in the lab, and to assess and correct timing errors during ocean bottom experiments.

In the following sections we will present plots of CSAC frequency error versus time elapsed where the former represents the average frequency error calculated over 6 hour intervals. When the phase readings are less frequent than once every 6 hours a flat section will be seen in the plot. We do not know the exact frequency error at any one time during this interval, but since it is calculated from the phase error at the end points it is truly the average frequency error over this interval. There are a few intervals where a tare occurred either due to power cycle of the clock or a “resync” of the clock without corresponding phase checks immediately before and after. We use the units Parts Per Billion (PPB) for convenience, i.e. a clock with a 1 PPB frequency error will accumulate 1 nanosecond phase error per second, or roughly 31.5 milliseconds per year. Note that since we are dealing with frequency errors less than 100 PPB, frequency error is interchangeable with period error for practical purposes.

The most important thing to look at in the following plots is not the absolute frequency error, but rather the change in frequency error. Standard procedure for OBS experiments is to measure the phase error immediately before deployment and after recovery in order to determine the average frequency error throughout the experiment. The timestamps on the seismic data collected can then be adjusted accordingly [3]. If the frequency error was constant between the beginning and end phase measurements then the residual phase error at any point will be zero. However no clock will have a completely constant frequency. The more the frequency changes, the larger the residual phase error, post correction, will be. Like most oscillators, the CSACs are given an aging specification which places bounds on the rate of frequency change. Originally the CSAC was specified with a 1 PPB per year aging rate, but it has since been relaxed to 10 PPB per year. Section IV will look at these effects quantitatively.

The CSACs incorporate a “Steering” function, allowing a user to enter a manual frequency correction [6]. We typically use this functionality roughly once per year to ensure that the

CSAC frequency error is as low as possible prior to each experiment. We record the magnitude and time of each frequency correction, so that we can add it back into the plots seen here to assess the actual frequency changes of the Cesium oscillator. Thus, although some clocks may show a several PPB error after a few years in these plots, the actual frequency error measured is much lower.

A. 14 Great Clocks from the First Batch

Out of 20 CSACs received in 2011, we have 14 CSACs in the lab which have been functioning well since being put into service. Fig. 1 shows the measured frequency errors of these 14 clocks over their life, covering 4.5 years. A few of these showed accelerated aging during the first year of operation, particularly the two traces which jump up to 2.5 and 4 PPB within the first few months. In discussions with Symmetricom we learned that it is not unexpected to see significant aging for a few weeks after power up, especially when the clock has been off for a long time. Because of this we have implemented keep-alive batteries and procedures to ensure that our clocks remain powered at all times, even during maintenance operations. If a clock must be powered down we take care to keep the interruption as brief as possible.

After the initial few months all 14 of these clocks performed as expected. The average magnitude of clock drift-rate for all experiments that utilized these clocks after discarding these two outliers from the first deployment is 9.1 milliseconds per year or 0.29 PPB. The clock frequencies were manually adjusted to calibrate out any error prior to each deployment by writing to the Steering register over the serial interface. For comparison, the average magnitude of clock drift-rate for other instruments on the same experiment with Seascan clocks was 333 milliseconds per year or 10.6 PPB. It is important to remember that these numbers do not necessarily guarantee better clock accuracy during the experiment since a linear drift correction is applied to all clocks, but it is still an indicator of a more stable clock. This is discussed in [3] and [5] and will be analyzed quantitatively in section IV.

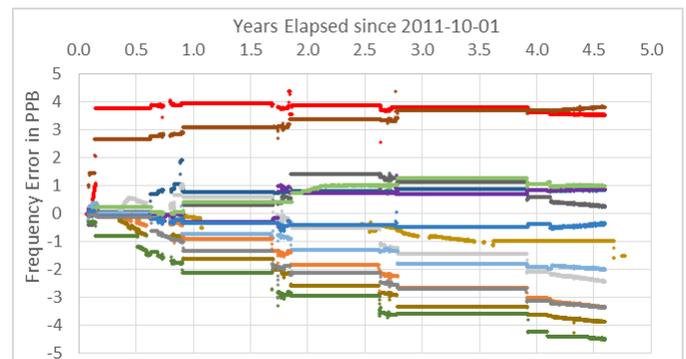


Fig. 1. Frequency errors for 14 CSACs manufactured in 2011 showing good performance

There is enough data to make a good estimate of aging rates for many of these clocks. Fig. 2 shows 6 representative clocks after the first year with linear trend lines plotted. The worst example here is aging at a rate of -0.75 PPB per year. This is

within the original Symmetricom aging rate specification of 1 PPB per year. Many clocks are aging significantly less than this. One of the most stable clocks here is shown in the red trace, which experienced a large shift within the first few months of operation, but since then has been aging at a rate of only -0.10 PPB per year.

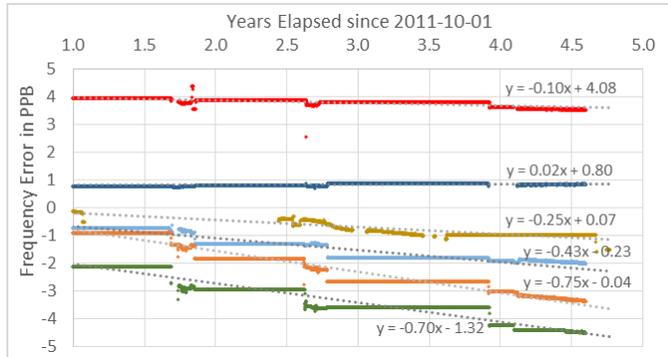


Fig. 2. Frequency errors for 6 CSACs manufactured in 2011 with trend lines showing roughly linear aging rates

B. 5 Good Clocks from later batches

Out of 20 clocks we have received manufactured in 2012 and later we have only five units that meet the original aging specification and are suitable for OBS use. The frequency errors for these clocks are shown in Fig. 3 over the span of about 2.5 years. With the exception of the purple trace these have all stayed within a range of less than 1 PPB, so they are well within the original specification. The purple trace was reasonably stable for the first year, but then shows a jump from -0.95 PPB during a year-long deployment to 0.25 PPB in the lab. We will be monitoring this clock closely in case it continues to age at this rate.

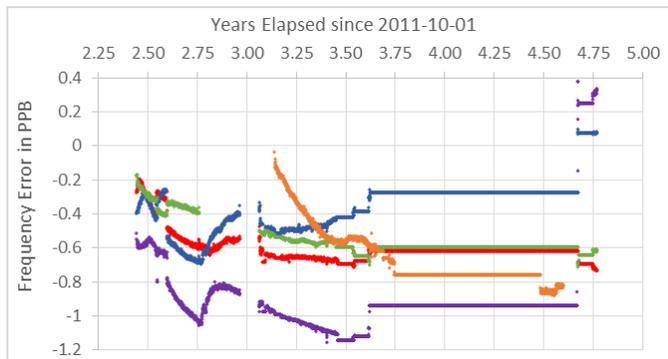


Fig. 3. Frequency errors for 5 clocks from 2012 and later with low aging rates

C. 4 Not-So-Great Clocks from the First Batch

The first batch of CSACs we received in 2011 did have a few outliers. Fig. 4 shows the frequency errors from these clocks. Three were deployed along with most of the clocks from Fig. 1. Two of these had an excessive phase error at the end of the deployment as did a couple of the other clocks that subsequently have behaved well. The third deployed clock,

blue trace, did not drift too badly on the deployment, but all three showed excessive aging rates well beyond 1 PPB per year in the lab after recovery. The fourth clock, yellow trace, remained on test in the lab and showed a very consistent aging rate of -4.51 PPB per year. In addition to the clear excessive aging rates, two of these clocks also showed accelerated vacuum loss as will be discussed later. None of these had reached the critical stage of vacuum failure yet, but were well on their way.

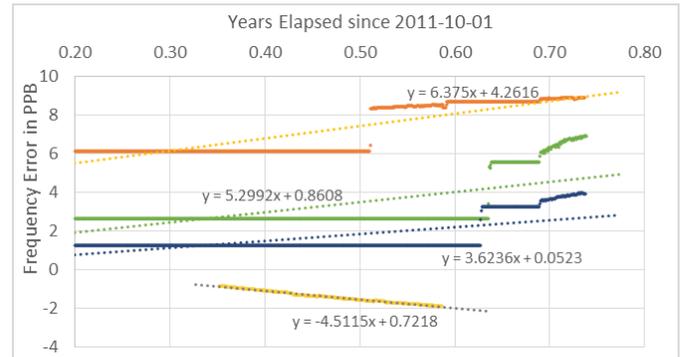


Fig. 4. Frequency errors for four CSACs returned in 2012 for excessive aging

D. 4 Clocks with consistent aging rate from later batches

Four of the clocks from later batches have demonstrated a relatively consistent aging rate that is well outside of the original 1 PPB per year specification. All of these are still within Microsemi's revised specification of 10 PPB per year, however they are not adequate for OBS use. As seen in Fig. 5 these have aging rate magnitudes from 2.62 to 6.63 PPB per year.

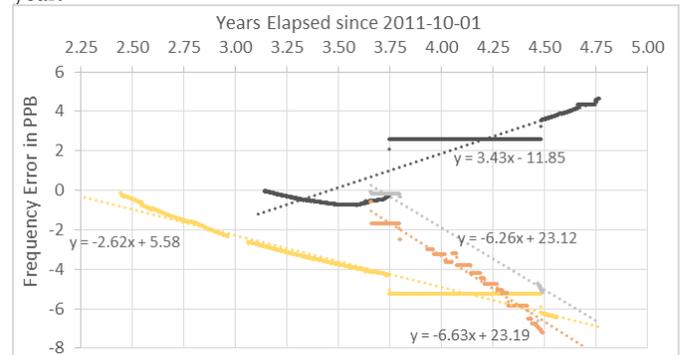


Fig. 5. Frequency errors for four CSACs with relatively consistent aging rates

E. 3 Clocks with inconsistent aging rates from later batches

A final 3 clocks from 2012 and later are aging unpredictably as shown in Fig. 6. The green trace was selected for a year-long deployment because it appeared to have settled out by around year 3.25, and seemed to be the best option of available clocks. However it is clear now that the apparent flat frequency was just the turning point, and it is now aging relatively quickly in the opposite direction. It may be aging at a

consistent rate now, but it is too soon to tell since we only have an average drift rate during the deployment.

The blue and orange traces were both running in the lab and showed significant variations in frequency before about 3.75 years in the figure. Unfortunately equipment failure led to the frequencies not being recorded over a 9 month interval after this point, but we do know the average frequency error over that interval as shown in the figure. The orange trace appears to now be aging rapidly at a relatively consistent rate, while the blue trace has leveled out. However there is not enough data to have any confidence in either, and both took much longer than expected to reach a steady state. These clocks may be adequate for many applications, but they likely will not maintain accurate time for an OBS.

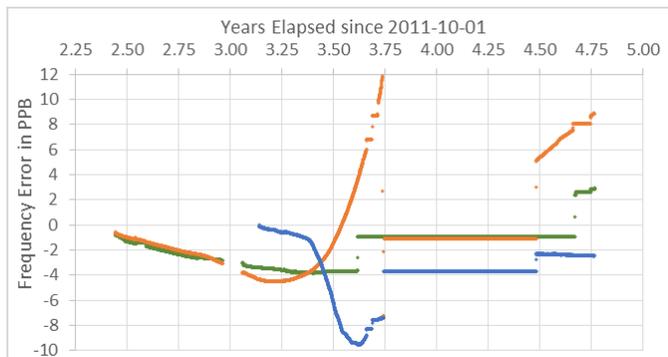


Fig. 6. Frequency errors for three CSACs with unpredictable aging rates

III. CSAC VACUUM FAILURES

Out of the 40 CSACs we have received 4 have experienced a hard failure. Three of these exhibited a vacuum failure in the lab. The fourth had unusual electronics problems in our lab, but experienced a vacuum failure en route to the manufacturer for analysis, and thus nothing more was learned about it.

A. Brief Technical Overview

The CSAC's hermetically sealed metal casing contains a second small sealed chamber containing the physics package. This is where the Cesium atoms which provide the stable frequency reference reside. In order to operate, the physics package must be raised to 70 degrees Celsius to vaporize the Cesium. The inner sealed chamber containing the physics package is evacuated to a high vacuum for thermal insulation. Thus the heater can generally maintain temperature regulation for only 10 to 15 milliwatts. (Several white papers are available from the manufacturer giving more technical details on the typical operation [7][8][9][10][11][12].)

In some clocks, particularly those manufactured in late 2012 and after, the vacuum may degrade quickly. This can be due to a weak seal in the chamber or outgassing from the materials used. As the vacuum deteriorates the power required to maintain the operating temperature increases. The internal heating element cannot produce more than around 35 milliwatts of heat, so eventually the physics package drops in temperature and the clock performance begins to degrade.

B. Warning Signs and Symptoms Observed

The first time we encountered a vacuum failure was after the first 6 month deployment we did with 13 CSACs. We contacted Symmetricom because a few of the clocks were not keeping time as well as we expected (Fig. 4). They helped us to determine that one of the clocks had nearly lost its vacuum. The heater power, which can be read from the CSAC status string [6], is typically around 10 to 15 milliwatts, but for this particular clock had started over 20 and was nearly 35 after the instrument was recovered. Fig. 7 shows the heater power of the failed clock (orange trace) along with a clock which was operating as expected for comparison (blue trace). Unfortunately this particular instrument had a glitch in the engineering data record so there are some significant gaps, but the trend can be seen. At this point we sent the clock back for a replacement. It was still regulating temperature, but it would not continue to do so much longer, and the operating power was significantly increased. Further, this clock had aged more than any of the other clocks in the first deployment, with a frequency error of nearly 8 PPB after less than 0.75 years, shown in the orange trace in Fig. 4.

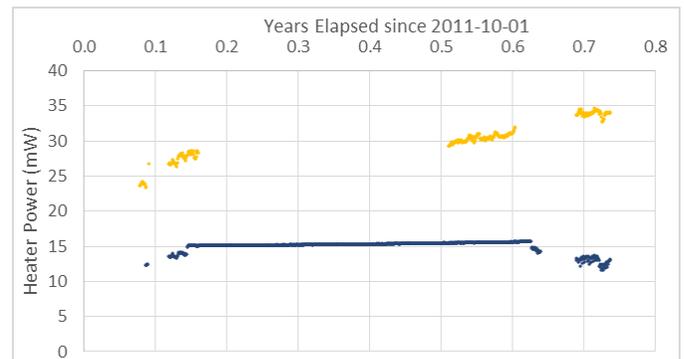


Fig. 7. Heater power of 1 failed CSAC and one good CSAC; 2011-2012

In 2014 we powered up several clocks in the lab for the first time. Some had been received in late 2012 as replacements for clocks returned, others had been purchased for new instruments. All had a serial number date of August 2012 or later. Three of these failed within the first year. Fig. 8 shows the heater power for these three clocks. The orange trace shows a relatively constant increase in heater power (notwithstanding the shift in heater power around 2.54-2.6 years due to temperature changes) until losing regulation at around 39 milliwatts. This steady increase may give the operator time to notice a failing clock and take it out of service. Fig. 9 shows that the frequency of this clock is not too bad up to the time the heater rails. The blue trace shows a relatively stable heater power until around 2.85 years when it suddenly jumps to the rail. A failure by a sudden jump like this would be difficult to detect prior to complete failure. Interestingly the frequency of this clock never strayed too far, though it is clearly not as stable after the heater rails. Finally, the gray trace shows a clock that did not experience vacuum failure in the lab, but rather had an electronics failure of some form. This was visible due to the loss of lock indicator as well as a variety of non-zero status codes in the serial status strings [6]. Fig. 9 shows that this failure was accompanied with wildly changing frequency. Also note that it exhibited a significant glitch with the

temperature change around 2.54 years, before it had started to exhibit significant symptoms otherwise. Unfortunately this clock experienced a vacuum failure while in transit to Microsemi, so they were unable to troubleshoot the electronics failure.

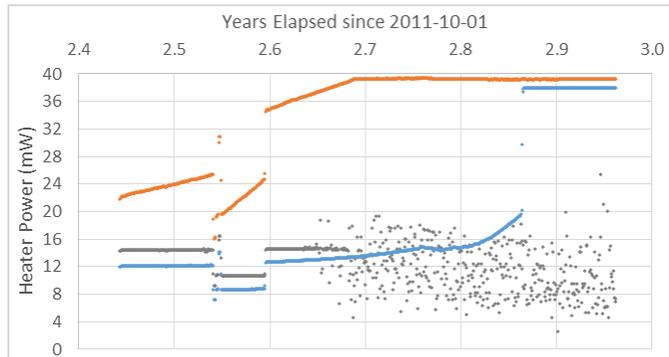


Fig. 8. Heater power of 3 CSACs which failed in 2014

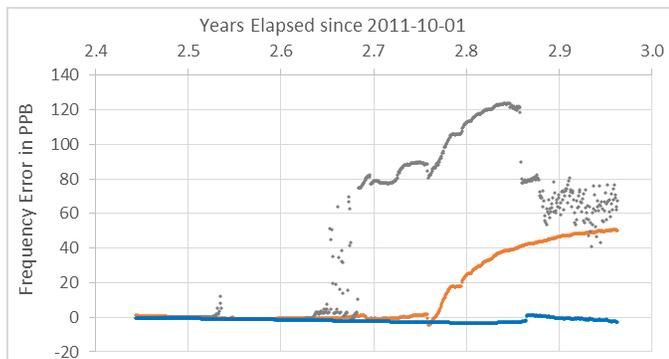


Fig. 9. Frequency error of 3 CSACs failed in 2014

The final vacuum failure we experienced occurred in early 2016, and was one of the original batch of CSACs from 2011. This clock had operated admirably throughout four long deployments, performing much like the clocks shown in Fig. 1. However the heater power was steadily increasing throughout its life. Fig. 10 shows both the frequency error and heater power from this clock over its entire service life. The steady sections of the heater power are where it is deployed and at a constant ocean bottom temperature. The dips in between these steady sections are where it is in the lab at room temperature and in transit. Here it does not need to work as hard to maintain temperature, but the fluctuations in ambient temperature lead to minor fluctuations in heater power. Throughout these short and long term variations there is a clear increasing trend. This was noted before the final deployment.

Around 4.1 years the heater power begins to increase more rapidly and at the same time the frequency begins to increase more rapidly. Just before 4.5 years the heater power hits the limit and the frequency shoots up wildly. Before this time the aging rate of the clock had been reasonably low. It had a significant change in frequency in the first weeks, but then was aging within the original 1 PPB per year specification. This clock has been removed from service now and replaced.

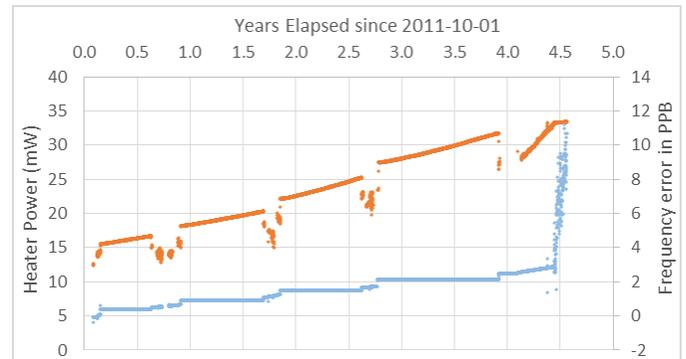


Fig. 10. Heater power (orange) and frequency error (blue, inferred rate of atomic clock after steering corrections) of CSAC in operation from 2011 until vacuum failure in 2016

C. Manufacturer's Recommendations

Microsemi has been aware of the vacuum failure issue since at least 2014 and has been actively pursuing a resolution. To ameliorate the problem in the meantime they have reduced the maximum rated operating temperature from 70 degrees Celsius to 35 degrees Celsius, as this reduces the likelihood of certain forms of vacuum loss. The CSACs will still operate at elevated temperatures, but with increasing probability of vacuum failure. It should be noted however that none of our clocks have spent time over 35 degrees Celsius and some still failed.

Microsemi claim to be close to a resolution. They have identified a number of contributing factors and put into place various fixes. Per our most recent correspondence with them they hope to have "production-pilot" units ready for sale late in 2016 with a restored operating temperature spec and reliability. These revised units will not address the change in specified aging rate however.

D. Heater Power as a Predictor of Lifetime

An important feature of the heater power indicator in the CSAC status string is that it can provide a reasonable estimate of remaining lifetime of the clock. Fig. 11 shows the heater power readings of 9 different clocks that are representative of the overall group of our clocks. The dark orange trace is the same clock seen in Fig. 10, for reference. The other clocks are all currently operating normally. As with Fig. 10, the periods of stable ocean-bottom temperature are easy to recognize by the stable reading and relative increase compared to the periods in the lab. It is clear from this figure that all of the heater powers are increasing with time, some more rapidly than others. The gray trace for example started just above 14 mW at the beginning of its first deployment on the ocean floor, but was above 22 mW at the end of its last deployment. The dark grey trace on the other hand started higher at nearly 15mW, but has not yet reached 17 mW. The purple trace, from a newer clock manufactured in 2014 started out very low, but has recently begun to increase rapidly.

Unfortunately this indicates that eventually all of these clocks will experience vacuum failure. However, disregarding the outliers and failures already seen, and assuming that the

heater power continues to increase at a consistent rate, these clocks should all last for the 100,000 hour MTBF rating given by Microsemi with only one or two exceptions.

It is clear that in any application where reliability is important the heater power must be monitored to identify impending failures. The temperature dependence of heater power and the anticipated future operating range must be considered in this analysis. Also, this technique will hopefully allow the end-user to verify that vacuum integrity has increased in new units once Microsemi begins selling CSACs with restored temperature range.

As useful as the heater power may be as an indicator of remaining lifetime it is also important to recall that some failures may not be predicted (e.g. electronics failure above), or that the heater power may suddenly begin increasing much more rapidly (e.g. blue trace in Fig. 8), shortening the amount of warning given.

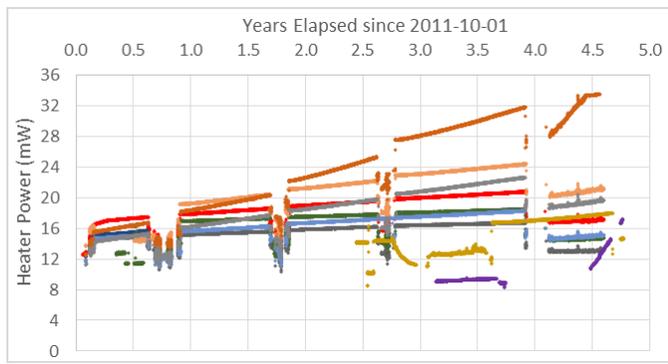


Fig. 11. Heater power for 9 representative CSACs from 2011 to 2016

IV. COMPARISON OF CSAC AND SEASCAN PHASE ERRORS

The data presented above are useful for comparing one CSAC to another in order to select the best unit for a particular deployment. However we can go further than this and use time-series of CSAC frequency errors collected in the lab to estimate the actual phase error that would accumulate over a certain interval. This same technique can be applied to Seascan Timebases, or any other clock in order to compare expected performance in the terms that matter for OBS.

A. Seascan phase error in lab test

We have been running a test on 39 Seascan Timebases in the lab using the same methodology as described in [3] for a longer period of time than any prior test. The clocks are housed in an insulated box inside a small temperature chamber which was stabilized at 20 degrees Celsius for several weeks prior to the test. The temperature was then changed to 2 degrees Celsius and held steady. We now have data covering an interval of roughly 11 months from this cold shock. Prior tests have been limited to 3 months or shorter so we have not been able to observe behavior over an equivalent timeframe to typical deployments.

As described in [3], the phase of the pulse per second output from each Seascan is measured roughly once per minute

to a resolution of 100 ns. There were several glitches during the test where the clocks all shifted by random large phases due to power interruptions. These glitches are trivially removed in post processing. Fig. 12 shows the frequency errors averaged over 6 hour intervals. One of the clocks has a large frequency error well above 100 PPB and rapidly changing, and has been left out of the following analysis, as it clearly has a fault. The remaining 38 clocks can be seen to shift suddenly by a certain amount at day 0 when the temperature shock occurs, and then undergo a slow ongoing change. This matches what we have seen in a number of previous experiments [3]. However we can now see that the majority of these clocks flatten out within the first 100 days or so. There are a few exceptions, and there is still some variation from month to month visible, but they are clearly trending towards stability.

In order to assess what effect the frequency errors seen here will have in a real world deployment, we have first calculated the average frequency error for each clock from day -0.25 to the end of the dataset and subtracted this static offset from the frequency error of each clock. This is analogous to the linear drift correction which is applied to real-world data sets – the average frequency error over the experiment will be 0, i.e. there will be no phase error at the beginning and at the end of the experiment. In this lab test we know the actual frequency error throughout the duration of the experiment, so by integrating this frequency error as shown in Fig. 13, we can get the actual phase error for each clock throughout the experiment. As expected the phase error is zero at the beginning and end of the simulated deployment, but deviates from zero during the duration of the experiment.

Many of the clocks have maximum phase errors of 30 to 70 milliseconds. These would make poor clocks for OBS deployments. However there are some that show much better behavior. Out of the 39 clocks in the test, 19 have a maximum absolute phase error throughout the test of less than 20 milliseconds, 9 have less than 10 milliseconds. There are 12 clocks with an average absolute phase error less than 5 milliseconds throughout the test. While this may seem like a small percentage of the total clocks in the test, it should be noted that the majority of the clocks in the test are ones that have been removed from service after previous lab testing revealed significant long term frequency shifts. Thus this group represents some of the worst clocks we have seen. The good ones in the group are new arrivals and a few back from recalibration that have not been assigned to instruments yet. It is reasonable to expect based on this and past testing [3] that the majority of the clocks currently in active service are providing performance around 10 millisecond maximum absolute phase error over a one year deployment. Assuming time can be found in the deployment schedule we hope to put some of the in-service clocks into longer-term tests to verify this.

B. CSAC phase error in lab test

Although we have not performed a similar long-term test on the CSACs we do have significant amounts of frequency error data over long periods from many clocks. We can apply similar techniques to this data to estimate the phase errors of year-long deployments with CSACs. In Fig. 14 we have

selected 330 day chunks of data from 9 representative CSACs. This includes the three clocks from Fig. 6 with inconsistent aging rates, two clocks from Fig. 5 with consistent but large aging rates, and four clocks from Fig. 3 with good small aging rates.

The temperature for these plots is typically roughly 20 degrees Celsius, so there is no temperature shock in them as in the Seascan tests, and the temperature is not as stable as in the Seascan tests. However past tests have shown that the CSACs are reasonably insensitive to temperature variations [3].

Fig. 15 shows the calculated phase errors from these 9 CSACs. The four with low aging rates have maximum absolute phase errors of between 0.33 milliseconds and 1.27 milliseconds with an average absolute phase error of just 0.33 milliseconds. The remainder have maximum absolute phase errors between 7.4 and 24 milliseconds. Two have an average absolute phase error greater than 10 milliseconds.

Performance may be improved for some of these clocks by changing the linear drift correction applied in post-processing to a second order drift correction. For the 4 clocks from Fig. 4 we see a consistent aging rate in the lab. If this remains relatively consistent in the field then we can apply a correction accordingly. In the field we can measure the frequency error as well as the phase error before each deployment and after each recovery. This gives us 4 data points, one more than necessary to fit a quadratic curve, so we may be able to make an assessment of how good the fit is likely to be. However this technique is likely to suffer from a number of shortcomings. First of all, the small variation in frequency due to temperature will affect the results, since we can only measure frequency on the deck of a ship, not at ocean bottom conditions. Second, we have seen in the lab that many of the clocks do not show a consistent aging rate, so a quadratic correction may not help, and may even make the phase error worse. Knowing when to apply a linear correction and when to apply a quadratic correction would be problematic. Finally, this complicates the post processing significantly, and requires a change to procedures which have been proven over hundreds of years of data sets.

V. CONCLUSION

Out of 20 CSACs we purchased in 2011 we still have 14 that are performing well for OBS work, one worked well for 4 years before signaling an impending failure, and we have no data on one more clock due to instrument loss. Based on the heater power trends we have reasonable expectations of continued service lives for the 14 remaining of between several years and decades. However, out of a total of 19 clocks received in 2012 and later, only 5 have proven suitable for OBS work. Three have experienced hard failures early in life, and 7 are aging too rapidly for use in OBS. The 4 remaining clocks have not been on test long enough to make definitive statements, but it appears that at best one out of these may be suitable for OBS work.

The Seascan clocks in contrast are not expected to reach the same level of performance as the best CSACs that we have tested. After removing many of the worst offenders from service we can get performance that is adequate, and at a

significantly reduced power requirement. It has proven to be a reliable workhorse, with very few outright failures, and minimal maintenance requirements.

For many applications the current CSACs will still be very good choices. Applications where GPS is typically available can benefit from CSACs as holdover oscillators. Shorter term underwater work and work where the level of accuracy is not as stringent as OBS may benefit as well. The current temperature limits may be restrictive and the overall reliability is concerning. Microsemi has acknowledged these issues and should have a solution this year. We hope that they will then focus on improving the aging specs so that they can again offer clocks with the performance we saw from our original batch. We would encourage other groups with similar timing needs to our own to contact Microsemi.

We are aware of no other solutions currently on the market providing the accuracy necessary for modern OBS at a power consumption low enough to operate autonomously for long periods. There are many miniature Rubidium oscillators that provide exceptional timing, but these all require over an order of magnitude more power than a Microsemi CSAC. No other company has a CSAC commercially available. Many crystal oscillators are available with very modest power requirements, and some even have temperature characteristics comparable to or better than the Seascan, but we are not aware of any which can match the Seascan's exceptional long-term aging. Hopefully this gap in the market will be filled again soon.

ACKNOWLEDGMENT

The authors would like to thank Pierre Tillier and Peter Jones of Seascan, Inc., and Steve Fossi, Doug Estrich, and Will Krzewick of Microsemi for their assistance over the years in seeking the best possible time for the least possible power.

REFERENCES

- [1] Microsemi Corporation, "Quantum SA.45s CSAC," Datasheet available from manufacturer.
- [2] S. Fossi, "The Chip Scale Atomic Clock and Satellite Navigation," Presented to the China Satellite Navigation Conference, May 17, 2013.
- [3] A. T. Gardner and J. A. Collins, "Advancements in high-performance timing for long term underwater experiments: A comparison of chip scale atomic clocks to traditional microprocessor-compensated crystal oscillators," 2012 Oceans, Hampton Roads, VA, 2012, pp. 1-8.
- [4] P. Tillier, "Seascan Inc. low-power high-accuracy time base module", Documentation available from manufacturer.
- [5] P. Gouédard, T. Seher, J. J. McGuire, J. A. Collins, and R. D. van der Hilst, "Correction of Ocean - Bottom Seismometer Instrumental Clock Errors Using Ambient Seismic Noise," Bulletin of the Seismological Society of America June 2014 104:1276-1288.
- [6] Microsemi Corporation, "SA.45s Chip-Scale Atomic Clock: 098-00055-000 User Guide, Rev C," Manual available from manufacturer.
- [7] R. Lutwak, "Principles of Atomic Clocks," 2011 Joint Conference of the IEEE International Frequency Control Symposium and European Frequency and Time Forum.
- [8] R. Lutwak, et. al., "The chip-scale atomic clock – low-power physics package", Proceedings of the 36th Annual Precise Time and Time Interval (PTTI) Systems and Applications Meeting, December 7- 9, 2004, Washington, DC, pp. 339-354.
- [9] R. Lutwak, D. Emmons, W. Riley, and R.M. Garvey, "The chip-scale atomic clock – coherent population trapping vs. conventional interrogation", Proceedings of the 34th Annual Precise Time and Time

Interval (PTTI) Systems and Applications Meeting, December 3-5, 2002, Reston, VA, pp. 539-550.

[10] R. Lutwak, et. al., "The chip-scale atomic clock – prototype evaluation", 36th Annual Precise Time and Time Interval (PTTI) Systems and Applications Meeting, November 26-29, 2007 Long Beach, CA, pp 269-290.

[11] R. Lutwak, et. al., "The miniature atomic clock – pre-production results", Proceedings of the 2007 Joint Meeting of the European Time

and Frequency Forum and the IEEE Frequency Control Symposium, May 29- June 1, 2007, Geneva, CH, pp 1327-1333.

[12] M. Mescher, R. Lutwak, and M. Varghese, "An ultra-low power physics package for a chip-scale atomic clock," Transducers '05, IEEE International Conference on Solid-State Sensors and Actuators, June 5-9, 2005, Seoul, Korea.

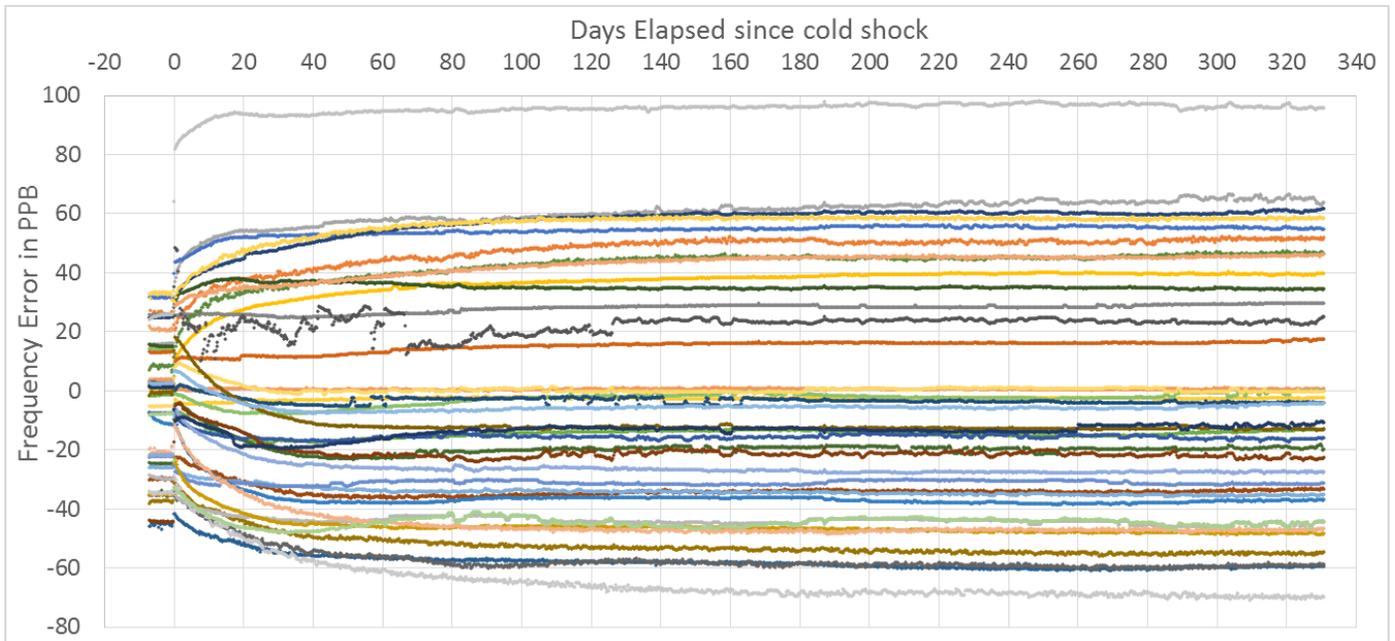


Fig. 12. Frequency errors for 38 Seascan Timebases in a simulated deployment.

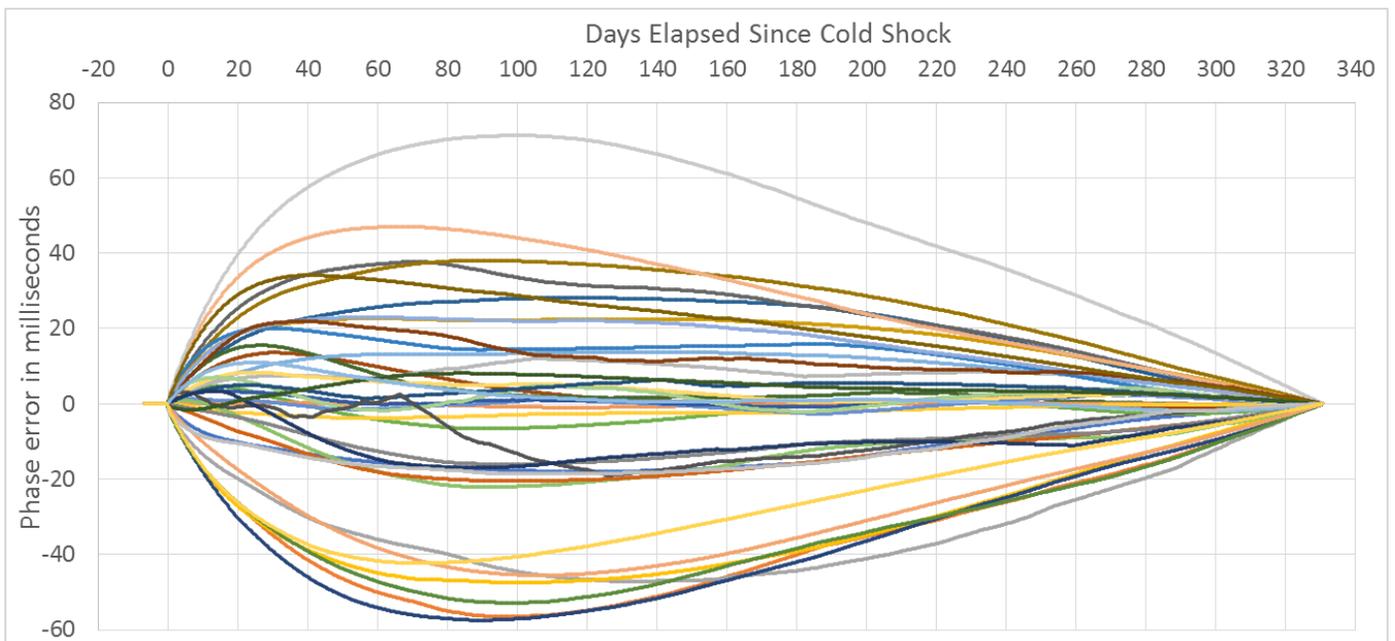


Fig. 13. Integrated phase errors from 38 Seascan Timebases in simulated deployment after linear drift correction has been applied.

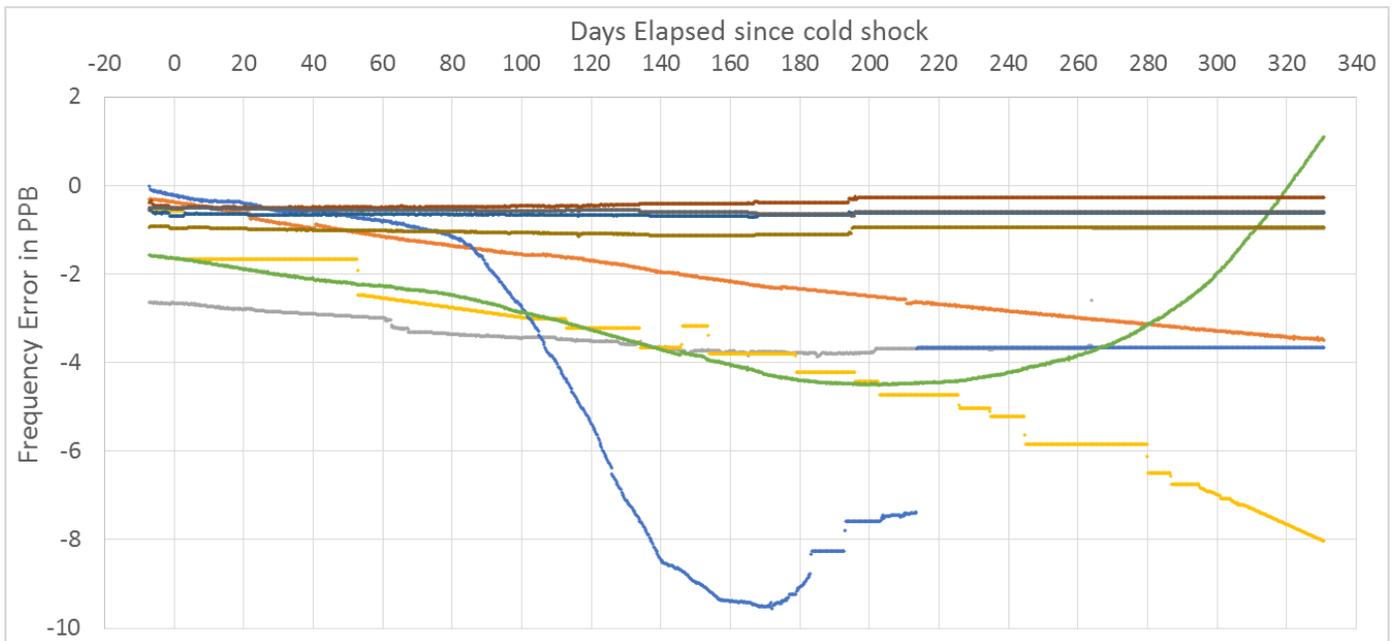


Fig. 14. Simulated deployment of 9 CSACs using the same methodology as the Seascan test above (without temperature shock)

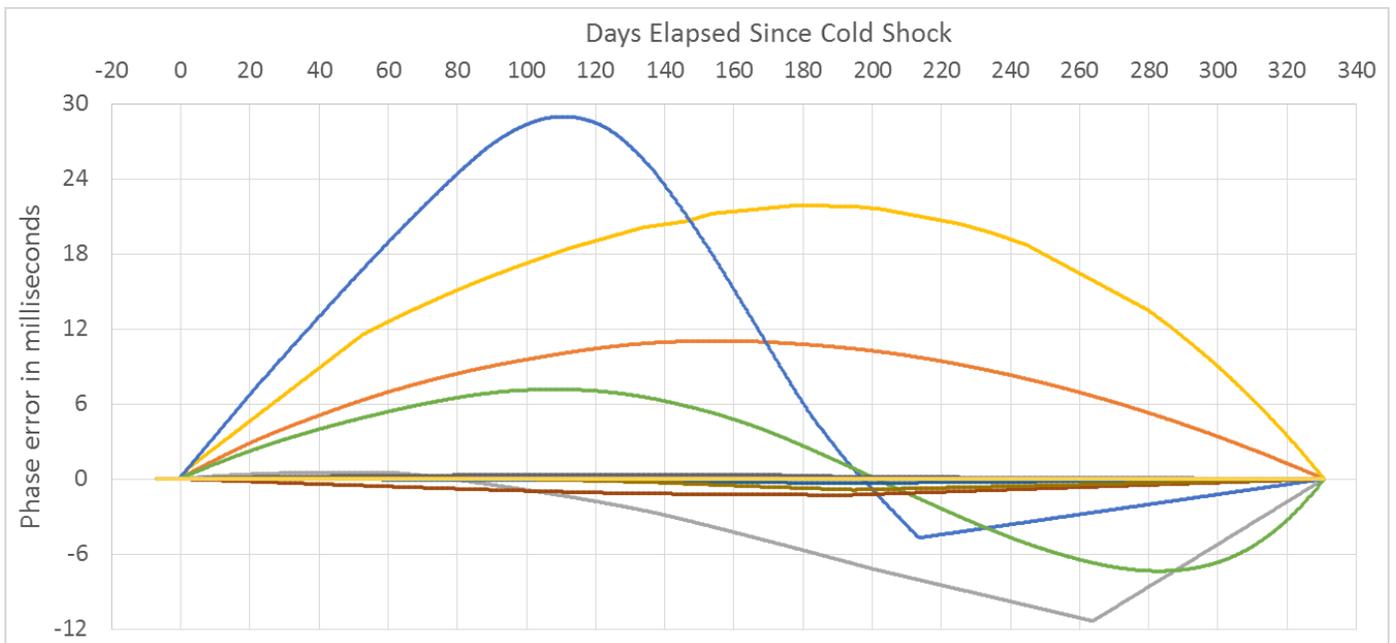


Fig. 15. Integrated phase errors from the 8 CSACs above.