HEWLETT PACKARD

Architecture, Design Aspects, and Performance of a New Cesium Beam Frequency Standard

Articles Reprint

Table of Contents

A NEW CESIUM BEAM FREQUENCY STANDARD PERFORMANCE DATA

John A. Kusters
James L. Johnson

Hewlett-Packard Co.
Santa Clara, California 95052-8059

Introduction

Recent papers have discussed the principles of a new architecture for cesium beam frequency standards.[1-4] The design goal of the new architecture was to develop an atomic frequency standard that is essentially independent of environmental effects. Reduction in Ramsey pulling due to cesium beam tube design effects was discussed in 1991[1,2]. With the addition of electronics that address power shift, C-field effects, and Rabi pulling, the new cesium standard shows the expected reduction in overall instabilities due to environmental factors and electronics errors. This paper reports actual performance data obtained from a number of standards designed incorporating the new architecture as measured by various measurement laboratories. The data substantiates the design theories, plans, and goals to produce a frequency standard whose sensitivity to environmental factors is greatly reduced when compared to existing standards.

Background

Previous cesium beam frequency standards generally exhibit a substantial response to environmental influences, such as temperature, pressure, humidity, and external magnetic fields[5-9]. An example of this is shown in Figure 1 where data collected over a year shows that frequency changes measured in several frequency standards maintained in a constant temperature environment, show a strong correlation with measured humidity.

Responses to environmental effects have become the major limitation on overall accuracy and on long-term stability in atomic frequency standards. This in turn requires metrology laboratories using these devices to design elaborate physical facilities to house the instruments in an attempt to minimize the environmental effects.

In addition, warm-up of a cesium frequency standard requires 24-36 hours of stabilization time to achieve final stability. This is a consequence of the typical cesium beam tube design, which conventionally consists of a copper Ramsey cavity weighing more than 1 kg. with a cesium oven at one end operating at 90-130°C, and a hot-wire ionizer at the other end operating in excess of 900°C. Typical thermal response times of a cesium tube are on the order of 6-10 hours, resulting in overall thermal equilibrium times of at least 24 hours.

Primarily through an expanded understanding of the underlying causes of the cesium standard's sensitivity to outside environmental factors, and extensive computer modeling of the atomic physics and of the electronic architecture, we became convinced that with the proper design, environmental factors could be greatly reduced. At the same time this should result in a more useful instrument that requires no special operating environment, and that should be fully operational meeting all specifications within 30 minutes of power-up. The latter condition is especially desirable for field transportable, fast set-up metrology and calibration facilities.

Consequently, a multi-facility, multi-year program was undertaken to develop a cesium standard based on the new architecture.

Environmental Testing

As a part of the investigation of the new architecture and as a qualification process for instruments developed during the project, the response of the new cesium beam frequency standards to external effects had to be accurately determined.

Reference Standard

The fundamental requirement for determining the response of any standard to an outside influence is to have an appropriate reference standard against which to measure. Given the expected thermal response time of the cesium beam tube and other components in the standard, full temperature equilibrium takes about 18 hours to achieve. This dictated that temperature changes in the environmental chamber should occur no more often than once every 24 hours. This in turn requires that the reference standard have sufficient stability over extended time periods to be able to measure the environmental effects. Early in the testing program, we determined that our house frequency standard, composed of several HP5061 cesium standards, did not have the needed accuracy or stability to be used as a reference standard. Given its sensitivity to temperature and humidity, the expected stability of the house standard over several days of measurement was on the order of a part in 10^{13}. During the first temperature run, no change in frequency of the new standard was observed within the allowable measurement errors of the house standard.

The first challenge was to create a reference standard that was sufficiently stable and had sufficient intrinsic accuracy to be used as a reference. As a result, three of the new standards were constructed and connected as an active ensemble as shown in Figure 2. In this configuration, the controller drives the three standards to operate at exactly the same frequency and phase, producing a frequency that is the mean frequency of the three standards, and where the absolute sum of the corrections applied to the individual standards are at a minimum. Operating in this manner, for identical standards, the overall accuracy and stability improve as the square root of the number of units in the ensemble, while the phase noise spectral density (in dB) improves as 10 log(number of units). This results in a reference standard whose expected accuracy and stability is a factor of 1.7 greater than that of a single unit.

Testing at facilities outside of Hewlett-Packard was done using various models of hydrogen masers. Up to sampling times in the vicinity of 10^5 seconds, this proved adequate. Beyond this, corrections had to be made to the data to account for the drift of the maser used as the reference frequency standard.

Temperature Sensitivity

During early temperature measurements, analysis of the data using correlation techniques showed that there was a frequency change response that occurred every 24 hours, the same period as the temperature cycling. The correlation also showed peaks with time periods less than 24 hours, indicating that some other process was occurring. Given that at this time we were looking for effects on the order of parts in 10^{14}, after a difficult and time-consuming process the phase detector used to measure the phase difference between the unit under test and the reference standard was found to be temperature sensitive and varying with the room air conditioning. After also environmentally controlling the phase detector, the upper bound of the temperature sensitivity of the standard was determined to be 1×10^{-15}/°C.

Further testing at JPL was performed to refine the data. This facility was used because of the excellent environmental chambers used, the existence of several hydrogen masers as reference standards, and the outstanding work done by this group in evaluating and understanding the environmental effects seen in our older standards[8,9]. Figure 3 is a plot of the Root Allan Variance of the standard under test. This is a conventional method of expressing oscillator stability as a function of averaging time[10]. In this test, temperature was cycled between 15° and 35°C. once per day. If there had been a significant response to temperature, this would have shown as a peak in the stability curve at a period corresponding to 24 hours. Given that there is none, the upper bound of the temperature sensitivity over the 20°C temperature range was determined to be less than 1×10^{-14}, or correspondingly, less than 5×10^{-16}/°C.

Humidity Sensitivity

Early testing for humidity effects was performed in our facility with no detectable results. An extended run was done using NRL facilities. Figure 4 is a plot of the relative humidity and observed frequency. During this run, the temperature was also being cycled from -17° to 57°C. exceeding the design specification at both ends. Figure 5 is a plot of the correlation function between relative humidity and frequency. Two peaks are shown, but they are too weak to draw any conclusion. NRL data indicates that both temperature and humidity effects are bounded by the measurement uncertainty of 1×10^{-14}.

Pressure Sensitivity

No sensitivities to changes in ambient pressure were seen within the limits of the reference standard. Tests were run over pressure ranges equivalent to 0 - 12,000 m. altitude. An interesting sidelight is that although most frequency standards are specified in terms

of altitude change, the new unit is sufficiently stable that relativistic effects dominate all other effects in a change in altitude.

Magnetic Field

No frequency shifts within the limits imposed by the reference standard (in this case, a hydrogen maser) were seen in a ±2 gauss field. Data was taken in both DC and 55 Hz magnetic fields in all three axes.[11]

Further Characterization

Time Domain Stability

Figure 6 is a plot of the time domain stability in the range of 0.01 to 20 seconds. Figure 3 is a similar plot extending the sampling time out to 10^6 seconds. Of interest here is that at roughly 10^5 seconds, the cesium standard uncertainty crosses the locus for the hydrogen maser used as the reference standard. Freedom from environmental effects results in a flicker-floor substantially below that of conventional cesium frequency standards.

Warm-up Time

Figure 7 is a plot of warm-up time and frequency offset from a hydrogen maser as measured by NIST. From this plot, measured on an instrument that was flown from Santa Clara to Denver, driven to NIST Boulder, then plugged in with no stabilization time, the warm-up time to meet specifications was 10 minutes, 16 seconds.

Phase Noise

Characteristics of the new standard in the frequency domain is shown in the phase noise plot in Figure 8. Apparent here is the improvement in phase noise especially in regions closer to the carrier, from 1 Hz to 1000 Hz offset.

Other Effects

Tests for all other effects, including EMI and radiated and conducted susceptibilities were conducted with no observable results within the measurement uncertainty of the reference standard. Tests also included verification to conformance to MIL-STD-461C, EC-92, and MIL-T-28800D.

Conclusions

The experimental data indicates that it is possible to design and construct frequency standards that are environmentally insensitive, at least to the level of 1 part in 10^{14}. A sound understanding of the physics of processes involved also leads to a rapid warm-up of the cesium instrument.

Acknowledgments

The authors acknowledge the technical direction and continued contributions of Drs. Len Cutler and Robin Giffard of HPL, and the joint HPL and Santa Clara R&D team for the development of the new standard. We also gratefully acknowledge the contributions and assistance of Dr. Richard Sydnor and associates at the Jet Propulsion Laboratory, California Institute of Technology., Dr. Joseph White and associates at the Naval Research Laboratory, Dr. Gernot Winkler and associates at the United States Naval Observatory, and David W. Allan and associates at the National Institute of Standards and Technology for making the majority of the highly precise measurements of the new standard contained in this report.

References

[1] Cutler, L.S., *et.al.*, "Frequency Pulling in Cesium Beam Frequency Standards due to ΔM=±1 (Sigma) Transitions," *Proceedings of the 45th Annual Symposium on Frequency Control*, pp. 544-553, IEEE Publication CH2965-2, 1991.

[2] Mueller, L.F., *et.al.*, "A New High-Performance Cesium Beam Tube Compensated for Ramsey Pulling," *Proceedings of the 45th Annual Symposium on Frequency Control*, pp. 554-559, IEEE Publication CH2965-2, 1991.

[3] Cutler, L.S. and Giffard, R.P., "Architecture and Algorithms for New Cesium Beam Frequency Standard Electronics," *Proceedings of the 1992 IEEE Frequency Control Symposium.*

[4] Karlquist, R., "A New RF Architecture for Cesium Frequency Standards," *Proceedings of the 1992 IEEE Frequency Control Symposium.*

[5] DeMarchi, A., "New Insights into Causes and Cures of Frequency Instability (Drift and Long-Term

Noise) in Cesium Beam Frequency Standards," *Proceedings of the 41st Symposium on Frequency Control*, pp. 54-58, NTIS Document No. ADA-216858, 1987

[6] DeMarchi, A., "Rabi Pulling and Long-Term Stability in Cesium Beam Frequency Standards," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, UFFC-34(6), pp. 598-601, November, 1987.

[7] Karuza, S.K., *et.al.*, "Determining Optimum C-Field Settings that Minimize Output Frequency Variations in Cesium Atomic Frequency Standards," *Proceedings of the 21st Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting*, pp. 385-400, 1989.

[8] Sydnor, R.L., "Environmental Testing at the Jet Propulsion Laboratory's Frequency Standards Laboratory," *Proceedings of the 43rd Annual Symposium on Frequency Control*, pp. 289-295, IEEE Publication 89CH2690-6, 1989.

[9] Sydnor, R.L., *et.al.*, "Environmental Tests of Cesium Beam Frequency Standards at the Frequency Standards Laboratory of the Jet Propulsion Laboratory,", *Proceedings of the 21st Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting*, pp. 409-420, 1989.

[10] Allan, D.W., "Statistics of Atomic Frequency Standards," *Proc. IEEE*, vol. 54, pp. 221-230, February, 1966.

[11] Private communication, Jet Propulsion Laboratory

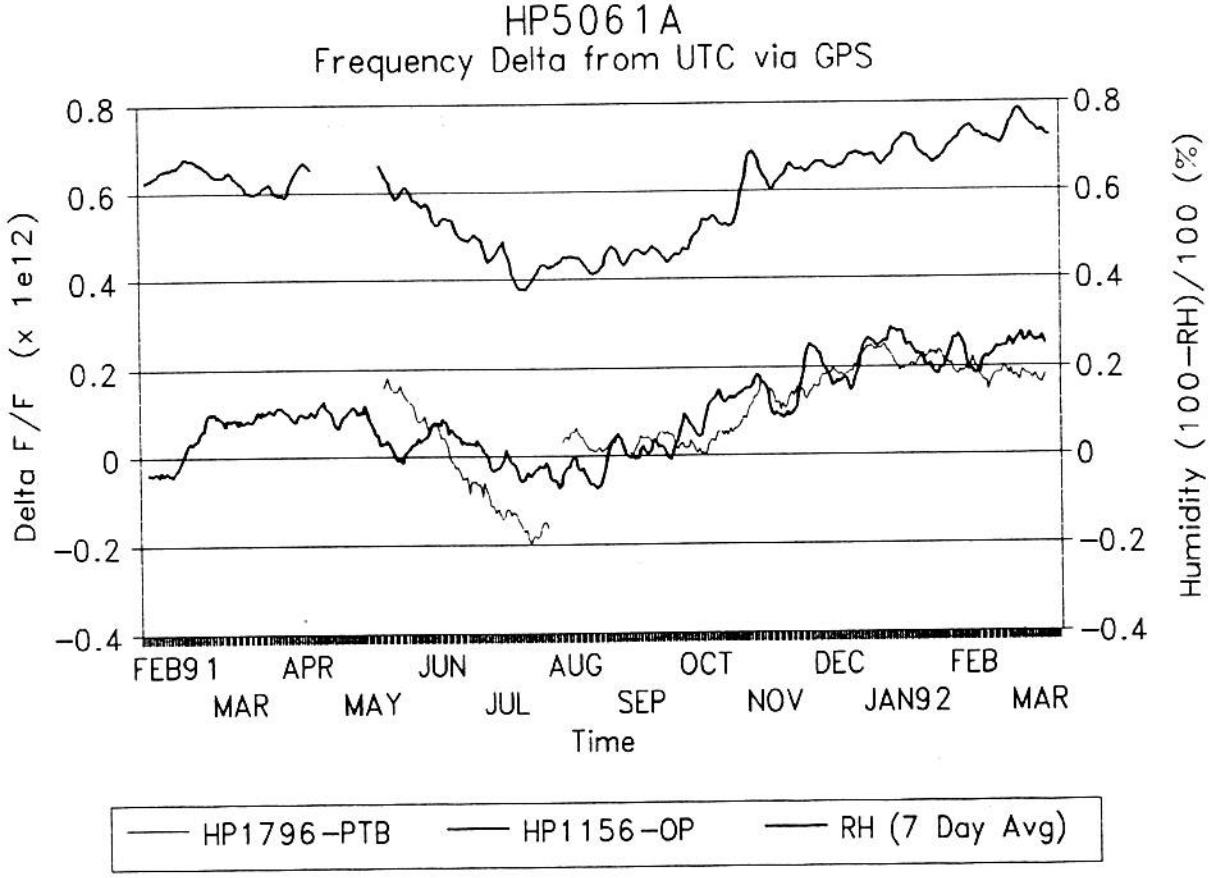

Figure 1: Upper trace - scaled relative humidity. Lower traces, - offset from UTC for two cesium standards as determined by common view GPS. (Data courtesy of HP Geneva Time & Frequency Laboratory)

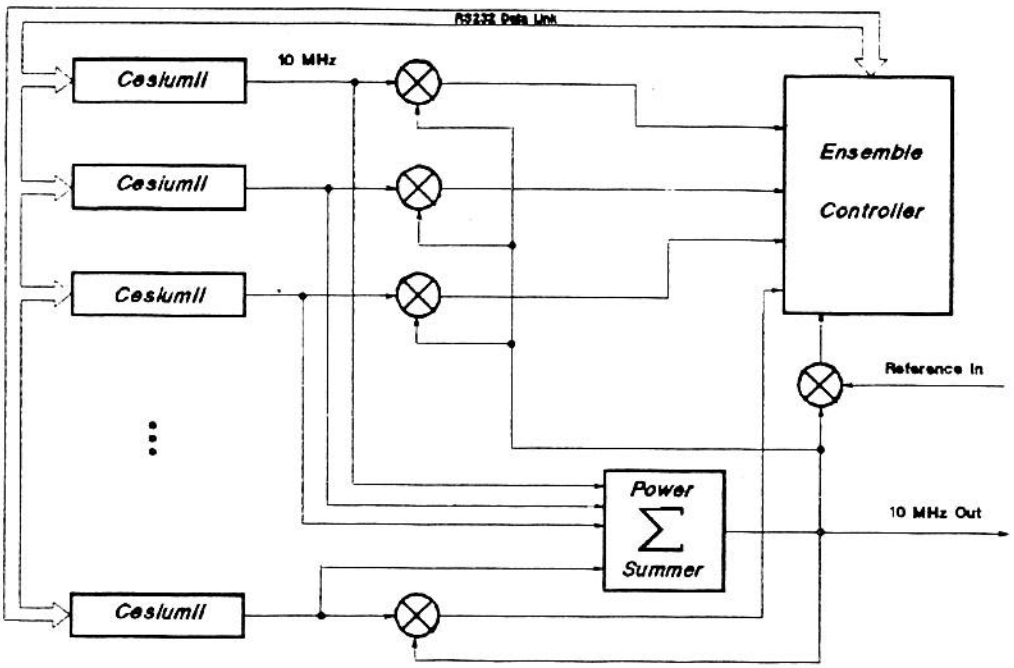

REAL-TIME ENSEMBLE OF FREQUENCY STANDARDS

Figure 2: Block Diagram, real-time ensemble of cesium standards.

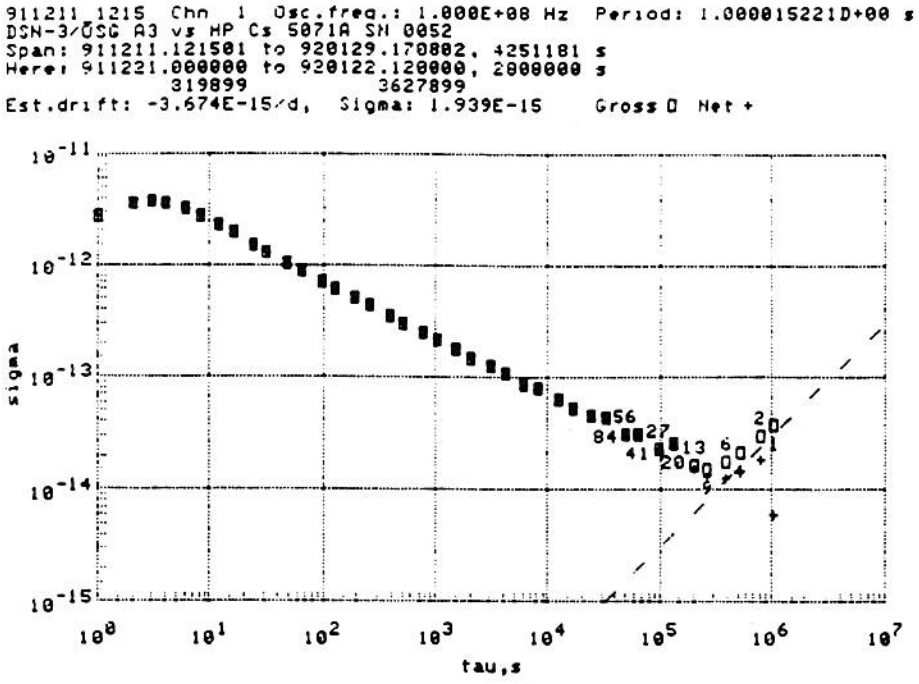

Figure 3: Root Allan Variance as determined by comparison to a hydrogen maser. Dashed line in lower right corner indicates expected instability in the hydrogen maser.
(Data courtesy of Jet Propulsion Lab, California Institute of Technology)

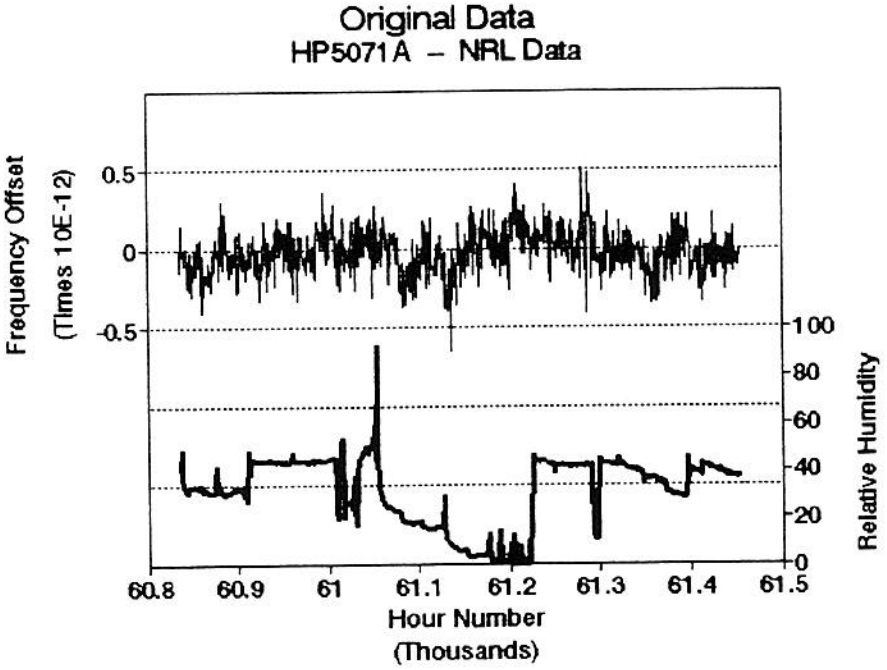

Figure 4: Upper trace - frequency offset from UTC.
Lower trace - measured relative humidity.
(Data courtesy of Naval Research Laboratory)

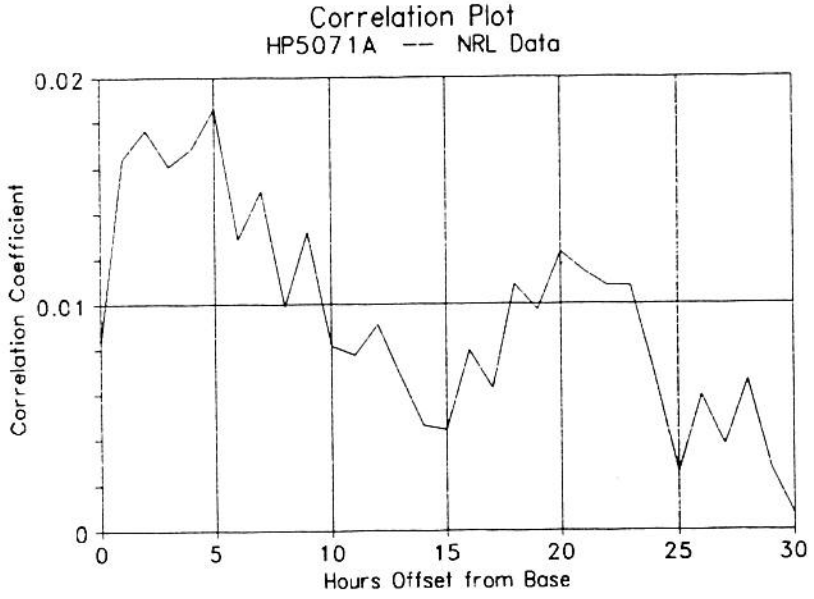

Figure 5: Correlation of data in Figure 4. Extremely weak peaks are observed at 5 and 20 hours offset.

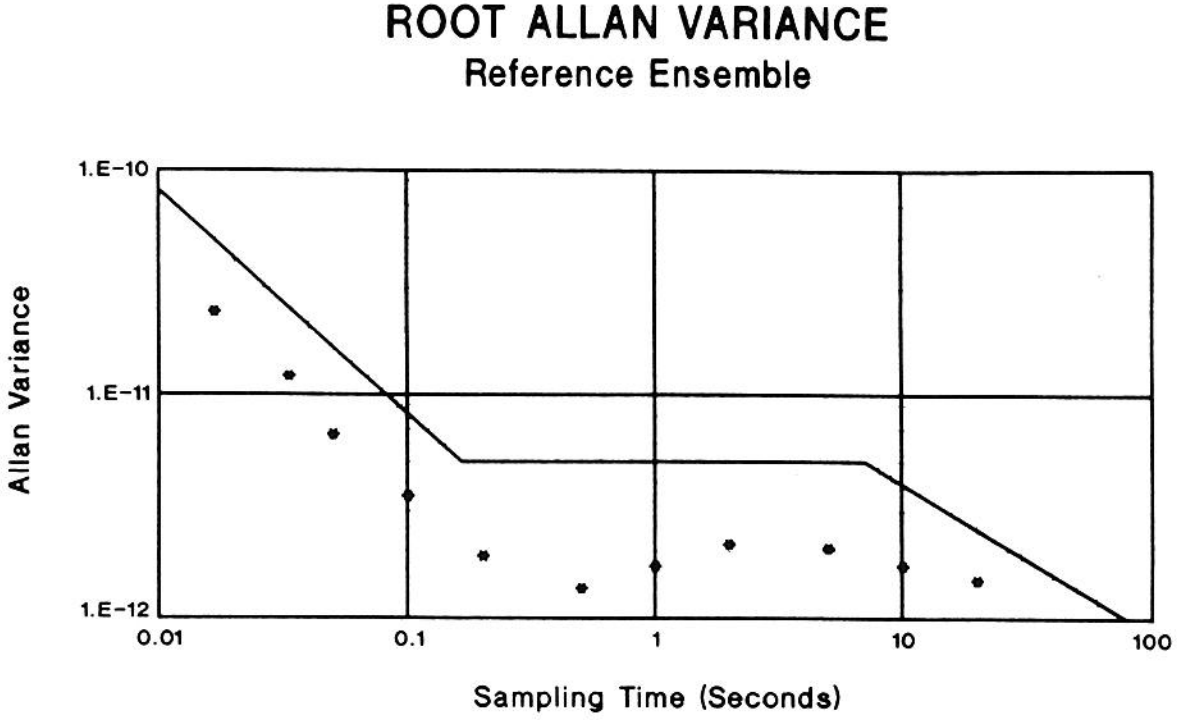

Figure 6: Root Allan Variance data for short sampling times.

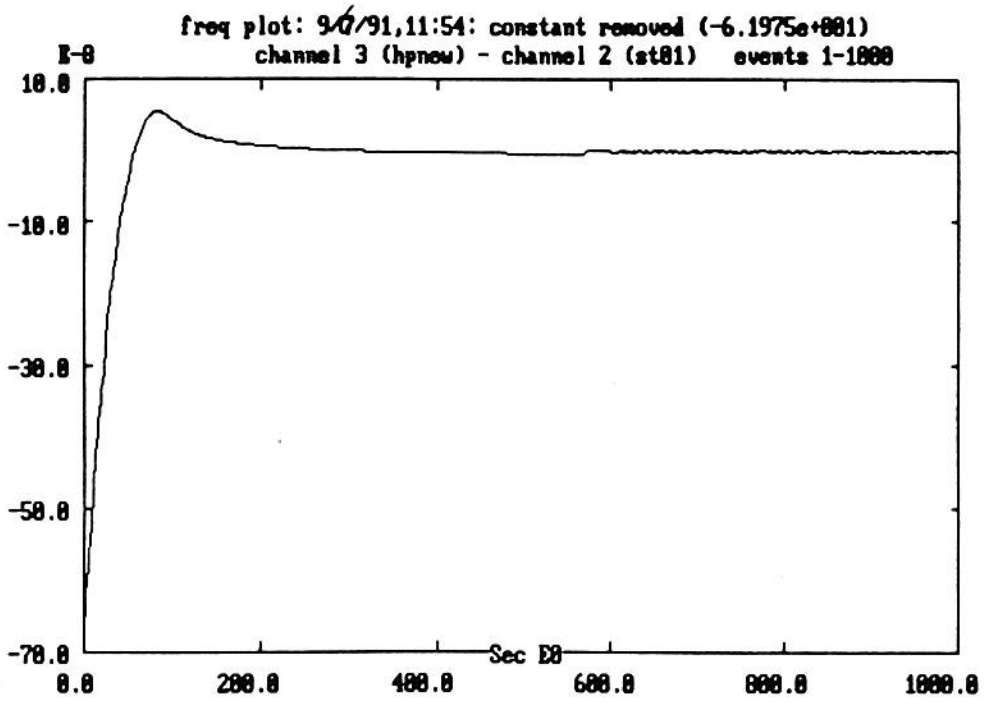

Figure 7: Measured frequency offset - from turn-on at NIST to 1000 seconds. Data indicates that frequency was stable at 616 seconds.

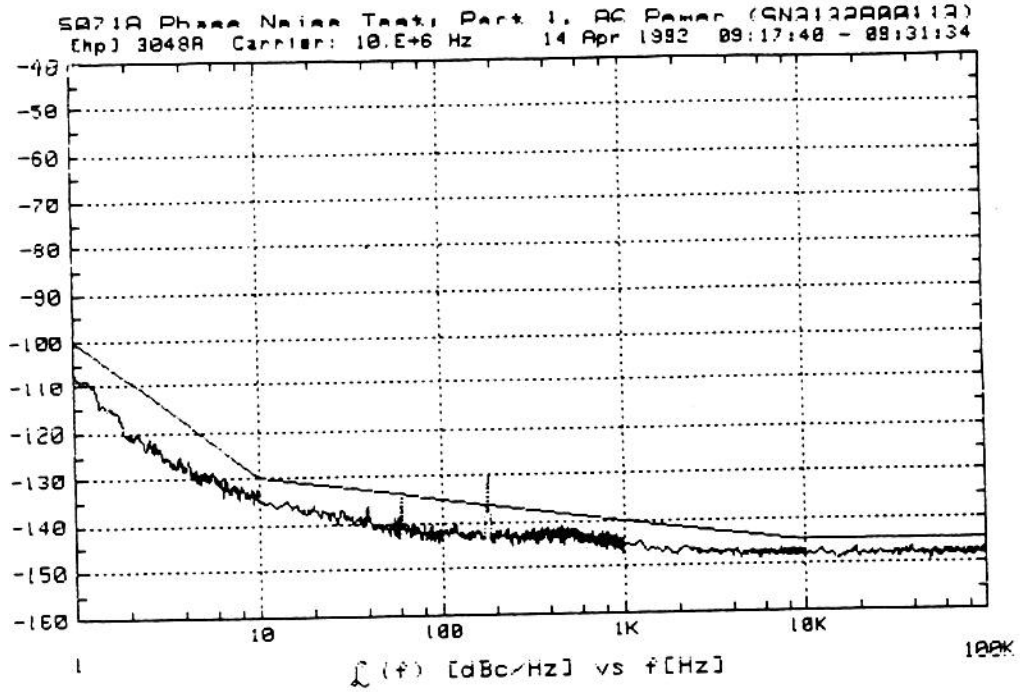

Figure 8: Phase Noise as measured on an HP3048A system

ARCHITECTURE AND ALGORITHMS FOR NEW CESIUM BEAM FREQUENCY STANDARD ELECTRONICS

Leonard S. Cutler and Robin P. Giffard
Hewlett-Packard Laboratories
P. O. Box 10350, Palo Alto, CA 94303-0867

Abstract-The main contributions to frequency error and long term instability in conventional cesium beam frequency standards are Ramsey and Rabi pulling, cavity phase shifts and cavity mistuning when there are microwave magnetic field amplitude variations in the microwave cavity. C-field instability is another large contributor. Changes in effective amplitude level or C-field with time, temperature, humidity, magnetic field, etc. are the main sources of frequency drift and instability with time or environmental changes. Careful design of magnetically deflected beam tubes can greatly reduce Ramsey pulling, and properly designed optically pumped tubes have virtually no Ramsey or Rabi pulling but the remaining frequency shifts with amplitude level and C-field changes can still be large. Finally, microwave spectral asymmetry and even order distortion of the line-center-finding modulation are also very important sources of frequency shift.

A description is given of a new electronics design that addresses microwave amplitude shift, C-field, Rabi pulling, microwave spectral purity, and modulation distortion. The design is microprocessor based with automatic lock-up and long term operation without operator intervention. The basic architecture and algorithms are covered as well as the design considerations necessary to limit error contributions due to electronics effects to the level of a few parts in 10^{14}. Frequency standards using the new electronics have demonstrated the expected stability and accuracy improvements [1].

Introduction

It is well known that one of the main contributions to frequency error and long term instability in cesium beam atomic frequency standards can be amplitude level changes in the microwave cavity magnetic field. When Ramsey pulling, Rabi pulling, cavity phase shifts, and/or cavity detuning are present, the frequency produced by the tube/electronics combination can depend on the microwave amplitude [2-5]. Then if the amplitude changes with time and/or environmental influences such as temperature or humidity, corresponding frequency drifts or instability can occur. In some cases the variation of frequency with amplitude is an oscillatory function of C-field value and can vanish at certain C-field values [6].

Asymmetry in the microwave spectrum can also cause frequency shift [7,8]. Asymmetry can arise from even order distortion in the intentional line-center-finding modulation or from the presence of unwanted sidebands produced in the frequency synthesis process used to obtain the desired microwave frequency. Also, if there is amplitude modulation that is coherent with the intentional line-center-finding phase or frequency modulation, spectrum asymmetry can occur. A simple example of this is sinusoidal frequency modulation with coherent amplitude modulation. If the relative phase of the two modulations is not exactly $\pi/2$ then the spectrum will be asymmetric. Coherent amplitude modulation can occur, for example, with mistuning of the tube's microwave cavity or a filter associated with it. There will then be a first order dependence of amplitude on frequency so that as the frequency changes during the modulation cycle there will be a corresponding amplitude change. Such modulation will be either in-phase or in antiphase with the intentional frequency modulation and thus lead to an asymmetric spectrum. This is the real origin of frequency pulling due to cavity mistuning.

There is, however, a saving feature. Due to the non-linear behavior of the beam tube's response, there is always a particular microwave magnetic field amplitude for which the frequency pulling due to coherent amplitude modulation vanishes. As will be discussed, this optimum amplitude depends on the type of intentional frequency or phase modulation. If the amplitude were to be stabilized at that level then pulling due to cavity mistuning would essentially vanish and, since the amplitude is stabilized, all the other amplitude dependent frequency shifts would either not be present or would be stabilized.

Note that since the cavity Q can change, constancy of microwave power does not imply constancy of magnetic field amplitude. The field amplitude is what is important and what must be measured.

Another main contributor is C-field instability. Clearly, if the C-field changes for any reason there will be a corresponding frequency shift. The main causes of

C-field changes are instabilities in the C-field power supply, expansion or contraction of the actual tube structure induced by temperature change, etc., and changes in the magnitude or direction of the external environmental magnetic field that are imperfectly removed by the tube's magnetic shields.

With all this in mind, the design of a new electronics package was undertaken with an error budget goal of less than 1 x 10^{-14} frequency shift for the effects mentioned above. Digital circuity was to be used as much as possible along with a microprocessor, and all important parameters were to be servo-controlled to remove the need for operator intervention and adjustments. A further goal was to make the standard capable of being completely controlled and monitored remotely through an RS232 interface. The architecture and algorithms of the design are described in the following sections.

Design and Algorithms

Modulation Type

The first choice to be made was the type of modulation to be used to find the line center. A modified type of slow square-wave frequency modulation [9,10] was selected. The microwave probe frequency is switched sequentially between various points on the resonance line. After each frequency change the tube output current shows a transient whose length is roughly the average transit time of atoms through the cavity smeared out by the atomic velocity distribution, the spread of the travel times to the detector, and the spread of cesium sticking times at the detector. After the transient has settled, the tube output current is averaged for a relatively long interval while the probe frequency is held constant. For the tubes used, the duration of the transient is 2 to 3 ms, and the frequency is changed each 12.3 ms, corresponding to a frequency of 40.7 Hz for pure square-wave modulation. There are several advantages gained by using this type of modulation which will now be discussed.

First, slow square wave frequency modulation allows the microwave amplitude to be controlled to produce the maximum signal from the beam tube at symmetrical points on either side of the line, the inflection points being preferred [14]. The microwave amplitude necessary to maximize the signal on the sides of the transition is exactly that amplitude for which the effects due to coherent amplitude modulation vanish so shifts due to cavity detuning are reduced to a negligible value. Detailed computer modeling verified these conclusions. Since the amplitude is now controlled to keep the microwave magnetic field in the transition regions of the tube constant, all the other amplitude dependent shifts are stabilized to a high degree giving excellent long-term stability as well as low sensitivity to temperature and humidity.

While there is also clearly an optimum amplitude [9] for other types of modulation such as sinewave phase or frequency modulation, or square-wave phase modulation, the optimum value can not be determined from tube signal measurements. For these types of modulation, computer simulation and experiments show, for example, that maximizing the current at the Ramsey peak, with or without modulation, or maximizing the recovered second harmonic does not give the optimum amplitude. Slow square wave frequency modulation is superior in this respect. Using the tube signal itself to determine and stabilize the proper amplitude is clearly optimum. Any other detection means would have to be calibrated and could drift with time or fluctuate with ambient conditions.

The second reason for the choice is that it is not too difficult to generate the required frequency modulation pattern with adequately low even-order distortion by using a direct digital synthesizer. The requirement is -120 dB with respect to the fundamental component to cause less than 1 x 10^{-14} frequency shift. The details of the frequency synthesizer are covered in a companion paper [11].

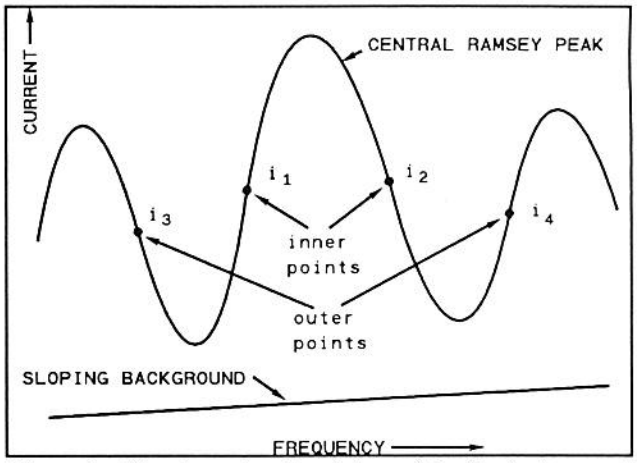

Figure 1. The four signal points used in the background slope removal technique.

The third reason is that it gives the best error-signal to noise ratio and thus the best asymptotic short term frequency stability for a given tube. Also, the error

signal is very easy to handle and can be converted at an early stage to the digital domain for processing.

Digital Synthesizer

One side benefit of using a digital frequency synthesizer is that the modulation can be extended easily to look at other points with opposite slope on the Ramsey pattern such as shown in Fig. 1. This is one of the modifications to the pure square-wave frequency modulation. Combining the error signals from the "inner" and "outer" pairs of points in the proper proportion and with the right sign gives an error signal that is free of any contribution from linear slope with frequency [12,14].

The total frequency error signal used for the case shown in Fig. 1 is given by:

$$i_{f\,err} = i_1 - i_2 - \frac{F_{p\,inner}}{F_{p\,outer}} (i_3 - i_4) \qquad (1)$$

where $F_{p\,inner}$ and $F_{p\,outer}$ are the frequency modulation deviations for the inner and outer points and i_j is the beam tube current at the point j. Such a background slope is produced by unbalanced neighboring $\Delta M=0$ transitions, the source of Rabi pulling. Using this algorithm greatly reduces the Rabi pulling. It must be mentioned that the optimum amplitude for the outer points differs from that for the inner points so a separate amplitude servo loop is required. A correction for this difference is applied in the algorithm that removes the background slope. Measurements made on the tubes used with this electronics show very little Rabi pulling. The calculated pulling without correction is somewhat less than 1×10^{-13}.

A major benefit of the digital synthesizer is that its frequency is programmable in precisely known steps thus allowing the cesium standard to be offset or steered by exactly known amounts either by means of its keyboard or remotely through the RS232 port. This is extremely useful for controlling the members of an ensemble of clocks.

Other benefits of the digital synthesizer arise from the use of its frequency agility. During part the modulation sequence it is set to the peak of the central line allowing a measure of tube signal size for setting the gain of the main frequency servo loop. It is also sequenced to the sides of the upper adjacent zeeman transition line. This is used for servo control of the C-field to keep that frequency constant. This is extremely important since it effectively removes C-field variations (except for any effects of field inhomogeneities which must be carefully considered in the design of the tube).

Frequency Architecture

The next main choice was the frequency architecture. To meet the goal of 1×10^{-14} requires that any unbalance in sideband pairs or any individual sidebands in the tube microwave probe signal be lower than about -70 dBc if they fall anywhere in the range of $\pm$ 130 Khz from the M=0 to M=0 clock transition [7,8]. This takes into account all the pulling effects including a sideband directly exciting any of the field dependent $\Delta M=0$ transitions at a point of maximum slope. Sidebands outside this range have a pulling effect that decreases inversely with the frequency separation and so become less important [7,8]. To ease the requirements on the synthesizer with regard to sideband levels, its output is translated to the microwave frequency in two stages with no frequency multiplication giving the additional benefit of no degradation in its frequency resolution. The first stage of the translation scheme uses a moderately narrow band phase-locked quartz crystal oscillator with the synthesizer output as one of its reference signals. The output of this oscillator at 87 MHz is used as the one of the references for the second stage of the translation, a phase-locked microwave dielectric resonator oscillator that provides the actual probe signal to the beam tube. The synthesizer frequency is thus effectively added to a 9192.5 MHz reference to obtain the tube probe signal. Both the 9192.5 Mhz signal and the synthesizer signal are derived from the same reference so they are coherent and the synthesized probe frequency is exactly known with respect to the reference. The two phase-locked oscillators performing the frequency translation provide symmetrization and filtering of undesired sidebands from the synthesizer and the other reference signals. With these techniques, the -70 dBc requirement for unbalanced sidebands is easily met without the need for a high cavity Q or any other microwave filtering. Considerably more detail is given in a companion paper describing the RF architecture [13].

The microwave adaptor that completes the cavity external to the tube and couples the microwave probe signal into the tube was designed to have minimal asymmetry about the plane normal to the E field in the waveguide. This should help reduce the phase shift associated with the unwanted mode excited by these asymmetries and thus reduce frequency shift.

The basic frequency source for the microwave chain is a clean, stable 10 MHz crystal oscillator which is

frequency controlled by the main cesium standard frequency servo loop. The output of this oscillator is not used for the output signal of the cesium standard for the following reasons. The frequency that is actually being controlled and stabilized is that of the microwave input to the cesium beam tube. To get to the tube input from the 10 MHz oscillator requires a lot of frequency multiplication with possibly poor phase shift versus temperature or humidity characteristics. Thus the phase of the 10 MHz oscillator with respect to the stabilized microwave frequency could change with time as the temperature or humidity changes. This would lead to transient frequency shifts since frequency is proportional to the time derivative of phase. To keep the user's output signals from the cesium standard tied as tightly as possible to the cesium tube input, a signal well up the frequency multiplier chain at 320 MHz is digitally divided by 4 to get to 80 MHz which is then effectively divided by 16 or 8 to get the actual 5 or 10 MHz output signals. The phase stability behavior of the digital division is considerably better than frequency multiplication [13].

Frequency and Microwave Amplitude Servo Loops

After amplification, the signal from the beam tube is converted to a digital signal and all further processing is done digitally up to the output DACs (Digital-to-Analog Converter). The frequency loop is completed by feeding the frequency control DAC output through a small amount of analog filtering to the 10 MHz crystal oscillator electronic frequency control input.

The error signal is processed using the background linear frequency slope removal algorithm mentioned earlier. The sense of the frequency modulation is periodically reversed and a form of second differencing is performed in order to obtain a frequency error signal that is free from the effects of signal and background linear drift with time. The resulting digital error signal is then processed with a software digital filter that performs a double integration with some damping to provide loop stability. Double integration removes the offset that would occur with linear drift of the quartz oscillator with time if only single integration were used. A digital servo loop monitors and keeps the current at the center of the line constant by controlling electron multiplier voltage and amplifier gain. The size of the actual signal is determined by the difference in current between line center and the sides of the line. After conversion to the digital domain, the signal size is kept constant by the software. This keeps the overall frequency servo loop gain constant. The closed loop time constant is several seconds.

The signal from the beam tube also contains the error information necessary to optimize the microwave amplitudes for the inner and outer interrogation points. Separate software signal processing is done to develop the two control words for the amplitude control DAC and these are sequenced into the DAC in synchronism with the frequency modulation so that the correct amplitude is always applied to the tube cavity. The amplitude modulation signal necessary to determine the optimum amplitudes is also generated by the DAC. The amplitude servo loops are closed by feeding the output of the DAC to a carefully designed amplitude modulator at the output of the microwave chain.

Microprocessor System

The microprocessor system controls almost all functions of the standard. It does all the digital signal processing and filtering and controls the cesium tube oven temperature and the C-field current. It also provides the interface between the operator and the instrument by means of a small keypad located on the front panel as well as the remote interface through the RS232 serial port. The SCPI (Standard Commands for Programmable Instruments) command set is used for remote control and communication.

Many subsystems provide monitor signals to the microprocessor and these signals are scanned frequently to see if they are in the proper range. The servo loops implemented in software are also continuously monitored. The basic status of the standard is indicated by a pair of lights controlled by the microprocessor on the front panel. One of these is a green continuous operation light. Most of the conditions and levels are available as data on a front panel LCD dot-matrix display or through the RS232 port. Conditions that require attention but do not cause improper operation are flagged as warnings. Anything that is interpreted as improper operation leads to a fatal error state. Either a fatal error state or microprocessor failure prevents the continuous operation light from being on. Warnings and fatal errors are documented in a log produced by the microprocessor and kept in non-volatile memory so the user can tell what occurred at some later time.

A TTL open collector status output is provided on the rear panel and the state for which it pulls down can be user programmed. Consequently, a number of standards can be wire-ORed together on one line giving the user a simple way to tell that at least one of the standards is in its programmed pull-down state.

Start-up is automatic and completely under microprocessor control. After initial warmup of ovens,

etc., levels and gains are set and a frequency scan is done to locate the line center and to set the C-field. If everything is normal, the servo loops are then closed and the system is in full normal operation.

All the information about the cesium tube that is necessary for proper operation is contained in a ROM attached to the tube. The microprocessor reads and uses this information to set up the proper operating conditions. This feature makes changing tubes extremely simple. There are no adjustments or measurements needed.

1 Pulse-per-Second circuitry

The circuitry for the 1 PPS (Pulse Per Second) is driven from an 80 Mhz signal derived from the same divider that runs standard's RF output circuitry thus making the 1 PPS coherent with the 5 or 10 MHz outputs. The timing or delay of the output pulse can be set manually or synchronized automatically to an external pulse with a resolution of 50 ns.

Power Supplies

The main internal supplies that are distributed to all the modules are +5, +12, and -12 volts DC. These are generated from high efficiency regulated switching supplies which are very carefully shielded and filtered. The input to these supplies can come from an internal supply powered by the AC line, or an external DC supply, or an internal battery. Power selection and internal battery charging are automatic and under control of a separate power steering logic circuit.

Many of the modules must have extremely clean power sources to avoid frequency pulling effects or sideband generation. Since the main distributed voltages are not perfectly clean most of the modules have additional internal post-regulator/filters or power conditioners. The worst offending signal on the power sources is at 81.4 Hz, the "second harmonic" of the basic interrogation frequency (the situation is complicated because the modulation is not just a simple square wave). Any extraneous frequency modulation of the microwave probe at this frequency can cause large frequency errors. One generator of this frequency on the power supplies is the microprocessor itself since it is processing chunks of data at that rate. Very careful filtering is done to keep this effect under control. Tests of complete standards include measures of both the susceptibility to this frequency on the supplies as well as the amount of signal actually there thus allowing the worst-case pulling to be calculated.

High voltage for the beam tube's electron multiplier and ion pump are developed by low power fly-back circuitry which is synchronized at a submultiple of the basic standard frequency. As mentioned earlier, the electron multiplier voltage, and consequently gain, is controlled by the microprocessor.

Software

The software for the whole standard was written in C. It is fairly complex since it includes all algorithms for the servos, signal processing, start-up, run, and general control. It also contains all the code that manages the interfaces, SCPI commands, diagnostics and logging.

Block Diagram

Fig. 2 is the basic block diagram of the cesium standard. Most of the features have already been described. Not shown are any of the monitoring connections or power supply details.

Summary

A new design for cesium beam atomic frequency standard electronics addresses most of the causes of frequency instability and error that are associated with circuity. The error budget for the various contributions is 1×10^{-14}. A microprocessor controls essentially everything so that start-up and operation is automatic without operator intervention. Complete remote control of the standard through an RS232 port, including precise frequency steering or offsets, makes it a valuable and versatile system component. Demonstrated performance with time and varying environmental conditions is excellent [1]

Acknowledgements

John Sadler contributed to a number of the algorithms, particularly the start-up algorithm, and wrote most of the fairly large amount of software. Rick Karlquist contributed to the frequency architecture, particularly to the standard's unique RF output circuits. Jim Johnson is credited with the efficient, small size, and low noise power supplies as well as the power steering logic.

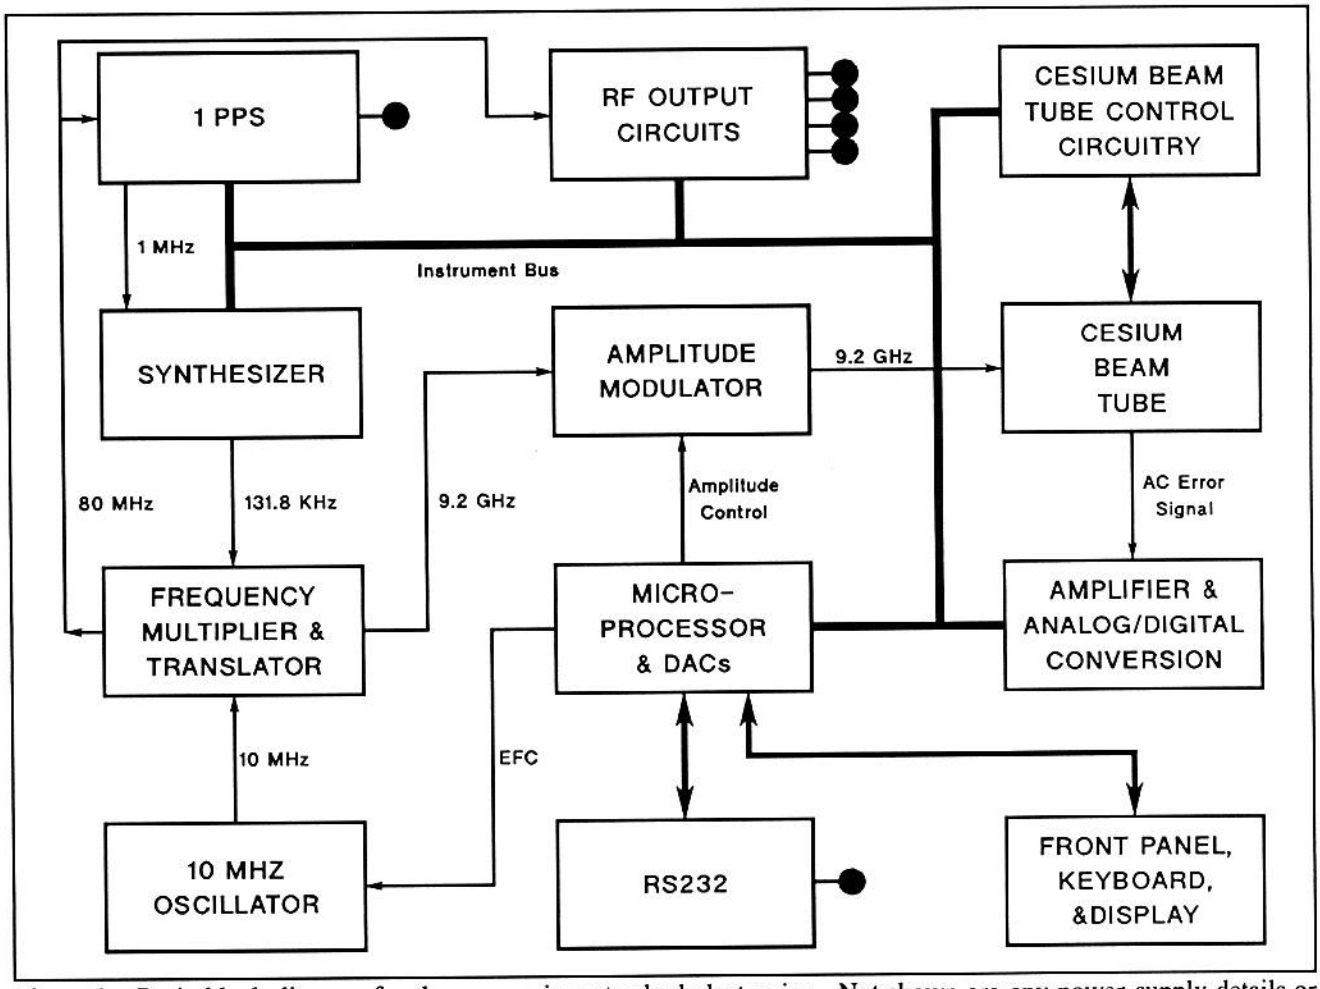

Figure 2. Basic block diagram for the new cesium standard electronics. Not shown are any power supply details or any of the monitoring connections.

References

[1] J. Kusters and J. Johnson, "A new Cesium beam standard: performance data", to be published in Proceedings of the 1992 IEEE Frequency Control Symposium, IEEE 1992.

[2] A. de Marchi, "Understanding Environmental sensitivity and ageing of Cesium beam frequency standards",in Proceedings of the 1st European Forum of Time and Frequency, Besancon, France, 18-20 March 1982.

[3] A. de Marchi, G.D. Rovera and A. Premoli, "Pulling by neighboring transitions and its effects on the performance of Caesium-beam frequency standards", Metrologia 20,37-47, 1984.

[4] L.S. Cutler, C.A. Flory, R.P. Giffard and A de Marchi,"Frequency pulling by hyperfine sigma transitions in Cesium beam atomic frequency standards", J. Appl. Phys., 69(5), pp. 2780-2792, March 1991.

[5] L.S. Cutler, C.A. Flory, R.P. Giffard and A de Marchi,"Frequency pulling in Cesium beam frequency standards due to delta M= ± 1 (sigma) transitions", in Proceedings of the 45th Annual Symposium of frequency Control, pp. 544-553, IEEE 1991.

[6] A. de Marchi, "Rabi pulling and long-term stability in Cesium beam frequency standards", IEEE Transactions on Ferroelectrics and frequency control, UFFC-34, 6, pp. 598-601, November 1987.

[7] J.H. Shirley, "Some causes of resonant frequency shifts in atomic beam machines", J. Appl. Phys., 34, pp 783-791, April 1963.

[8] C. Audoin, M. Jardino, L.S.Cutler, and R.F. Lacey, "Frequency offset due to spectral impurities in Cesium-beam frequency standards" IEEE transactions on Instrumentation and Measurement, Vol IM-27 No.4, pp.325-329, December 1978.

[9] A. de Marchi, G.D. Rovera and A. Premoli, "Effects of servo loop modulation in atomic beam frequency standards employing a Ramsey cavity", IEEE Transactions on Ferroelectrics and Frequency Control, UFFC-34, 6, pp. 582-591, November 1987.

[10] Y. Nakadan and Y. Koga, "A squarewave F.M. servo system with digital signal processing for cesium frequency standards", in Proceedings of the 36th Annual Frequency Control Symposium, pp. 223-229, IEEE 1982.

[11] R.P. Giffard and L.S. Cutler, "A low frequency, high-resolution digital synthesizer", to be published in Proceedings of the 1992 IEEE Frequency Control Symposium, IEEE 1992.

[12] A. Bauch, K. Dorenwendt and T Heindorff, "The PTB's atomic frequency standards CS2 and CSX: frequency shifts by pulling due to neighboring transitions", Metrologia 24, pp. 199-203, 1987.

[13] R. Karlquist, "A new RF architecture for Cesium frequency standards", to be published in Proceedings of the 1992 IEEE Frequency Control Symposium, IEEE 1992.

[14] U.S. Patent pending.

Abstract-A high resolution digital frequency synthesizer was designed to provide modulation and fine frequency interpolation in a new Cesium frequency standard. Requirements to be satisfied include resolution, spectral purity, phase-continuous frequency modulation, simplicity and low power. In the following paper we describe how these requirements were satisfied with a Direct Digital Frequency Synthesizer (DDFS). This type of frequency synthesizer was originally described by Tierney, Rader, and Gold [1], and has since been widely used and discussed [2,6]. The application to a hydrogen maser has been described by Matisson and Coyle [7].

Design Requirements

In the new Cesium standard, the microwave probe frequency has to be tuned to both the F=3, Mf=0 to F=4, Mf=0 "main" transition near 9192.632 MHz, and the F=3, Mf=1 to F=4, Mf=1 "Zeeman" line. At the chosen magnetic field, this requires the synthesizer to have a range of about 50 kHz. The necessary resolution is fixed at better than 92 microhertz by the requirement that the standard be steerable with a settability of 1 part in 10^{14} or less. To carry out effective square-wave frequency modulation, the synthesis chain must change frequency and settle to an accuracy of 1 part in 10^{14} in a time of about 2 ms, the duration of the transient generated by the tube when the frequency changes. For frequency changes of less than 1 kHz, the frequency should change without any phase discontinuity. Finally, to avoid offensive frequency pulling, the synthesizer should be free of sidebands close to the resonance line to a level of about -70 dBc. The synthesizer output has to be free of sidebands coherent with the modulation process and its harmonics to about -120 dBc. Because of the need for fast phase-continuous frequency switching, the requirements for this synthesizer are somewhat different from those relevant to hydrogen maser design [7].

A Direct Digital Frequency synthesizer was chosen to obtain the necessary resolution. In order not to degrade phase noise, the output frequency of the synthesizer is translated without multiplication to the probe frequency. To minimize power consumption and cost, the DDFS operates at the lowest satisfactory clock frequency. The overall synthesis in the cesium standard is defined by the following equation:

$$Fp = 29 \times 320 \text{ MHz} - ((35/32) \times 80 \text{ MHz} - Fs), \quad (1)$$

where Fp is the probe frequency, and Fs is the frequency of the digital synthesizer. The required range of the synthesizer is thus 125 kHz to 175 kHz. The architecture and overall synthesis scheme of the new Cesium standard are discussed in other papers [8,9].

Direct Digital Synthesis

The principal of the DDFS is well-known [1-3]. Its performance depends on the operation of the digital phase accumulator shown in Fig 1. The adder A causes the contents of the phase register R to increase by the value Na of the addend at each clock cycle. If the adder is binary and the phase register contains P bits, after N clock cycles the phase register will contain the number $N \cdot Na$ MODULO 2^P. The output of the phase accumulator is a sampled saw-tooth wave with a repetition period of $2^P/(Fc \cdot Na)$. It is well known that if the exact regularly spaced phase samples are used to calculate the inputs Vi to an ideal digital-to-analog converter (DAC) using the equation:

$$Vi = Const \times \sin\left[\frac{2\pi Ai}{2^P}\right], \quad (2)$$

the DAC output signal after an ideal anti-alias filter will be a pure sinusoid at a frequency $Fs = Fc \cdot Na/2^P$, completely free of non-harmonic distortion. The Nyquist condition: $Fs < Fc/2$, must be satisfied.

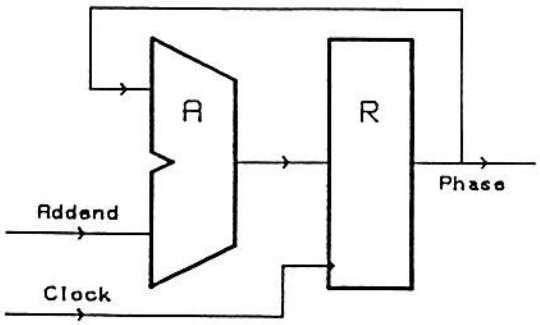

Fig 1. Digital phase accumulator consisting of a binary full-adder A and a clocked register R.

In practice, the conversion of linear phase to sinusoidal DAC input values is usually carried out by a "look-up table" or "waveform-map" consisting of a static read-only memory (ROM), avoiding the need for real-time calculation. Since the amount of memory required to encode the entire width of the phase accumulator would usually be prohibitive, only the Q most significant bits of the accumulator output are generally used to calculate the sine-wave samples. The truncation of P-Q bits of the phase leads to errors in the computation, and the output waveform cannot be a pure sine-wave, although its average frequency is not changed. The errors cause spurious frequencies known as "phase truncation sidebands" to appear [4,5].

It is well known in the field of digital signal processing that the finite resolution of the DAC used to construct the waveform introduces sidebands due to quantization error. Non-ideal static and dynamic DAC effects lead to further degradation.

We will now consider how a synthesizer was designed to meet the requirements of the frequency standard.

Design

It is shown above that the DDFS resolution df is related to the clock frequency and the phase register length by $df=Fc/2^P$, so that register length and adder width are minimized by choosing the lowest clock frequency that will allow satisfactory anti-alias filtering. The requirements of the instrument are satisfied by a 34-bit register and a 1 MHz clock frequency. The clock frequency is more than 4 times the highest output frequency required, allowing the two highest bits of the adder to be implemented as a synchronous counter, and making the anti-alias filter less critical.

The choice of the DAC resolution is dictated by the -70 dB spectral purity requirement. The quantization noise power in an ideal uniform DAC is easily shown to be of the order of $S^2/12$, where S is the step size. For a DAC having D bits, the maximum sine-wave output power is $2^{(2D-3)}S^2$, so that the total noise power is theoretically -1.76-6.02 D dBc [10]. In practice this power is distributed over a number of frequencies which are m/n intermodulation products of the clock and synthesizer frequencies. The power in any given sideband can be minimized to some extent by choice of the ROM program [5]. A 12-bit DAC was chosen, and a computer program was used to calculate the level of the dominant spurious sideband and choose the ROM coding. Allowing for static errors, an ideal 11-bit DAC was simulated, and it was found that the level of the unwanted sidebands should be below -90 dBc.

In order to obtain sideband levels approaching these values in real DACs, it is normally necessary to suppress glitches either by following the DAC with a sample-and-hold circuit, or by blanking its output while it settles after each data transition. Large-scale non-linearity in the DAC transfer function is expected to cause harmonic distortion, which is not important in this application.

The magnitude and spectral distribution of truncation sidebands has been explored theoretically by Nicholas and Samueli [4] and by Nicolas, Samueli and Kim [5]. A very large number of sidebands is produced in general, with a complex distribution of frequencies. The amplitude of the sidebands is found to depend on the width of the phase accumulator P, and the number of bits Q input to the waveform map. A small modification of

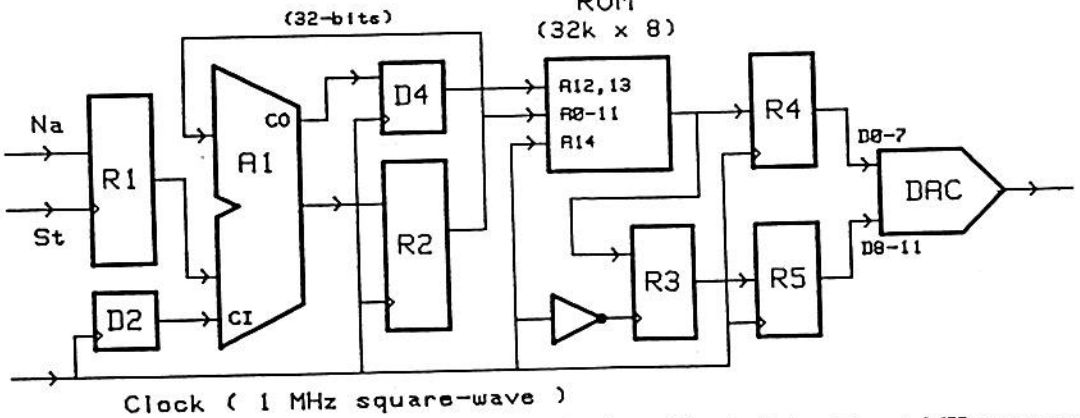

Fig 2. Block diagram of the Direct Digital Frequency synthesizer. The clock input is a 1 MHz square wave. The operation of the circuit is explained in the text. The DAC de-glitcher, and the anti-alias filter are not shown.

the accumulator reduces the worst possible sideband amplitude by about 4 dB. It was decided to use an address width of 14 bits in the synthesizer since this was a convenient match to a 256-k ROM with two bytes for each phase value. With an effective P value of 35, and $Q=14$, worst-case truncation sidebands are calculated to be below -84 dBc. We expect that as a result of the frequency switching sequence used in our application, the chance of a spurious sideband coherent with the modulation process is negligible.

Hardware Design

The hardware design of the synthesizer is shown in fig. 2. The 32 bit addend word *Na* is transferred to the register R1 by a hardware-derived timing strobe. The adder A1 is a 32-bit full binary adder whose output is latched into the 32-bit phase register R2 at each clock cycle. The carry output of the adder forms the enable of a 2-bit synchronous binary counter D4 whose output forms the two most significant bits of phase. This configuration is equivalent to a 34-bit phase accumulator with a 32 bit input word, limiting the output frequency to *Fc*/4. As suggested by Nicholas and Samueli, the carry input of the adder is at half the clock frequency. The upper 14 bits of the phase word are fed to address bits 0 through 13 of the ROM, and bit 14 alternates at the clock frequency. Registers R3-R5 align the upper and lower byte outputs of the ROM in time, and ensure that the DAC inputs all change simultaneously. To obtain good phase noise it is essential that the clock input to R4 and R5 has very low jitter. The DAC is a 12-bit, bipolar device with symmetric current outputs. De-glitching is implemented by gating off the DAC output for 200 nS after each data change with a balanced current switch. A

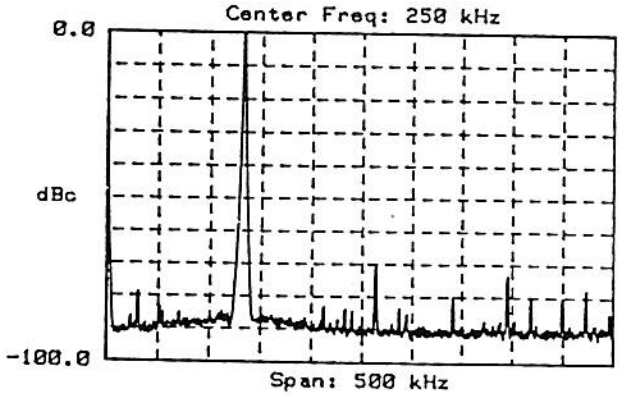

Fig 3. Output spectrum of the DDFS up to the Nyquist limit for a programmed output frequency of 131770 Hz. The lines in the spectrum are listed in Table I.

Table I

Frequency (Hz)	Level (dBc)	Origin
263540	-72.5	$2Fs$
395310	-74.1	$3Fs$
472920	-78.2	Fc-$4Fs$
30710	-78.4	$23Fs$-$3Fc$
341150	-79.2	Fc-$5Fs$
449470	-79.7	$11Fs$-Fc
418760	-80.1	$2Fc$-$12Fs$
51450	-83.0	$(2^{14}+1)Fs$-$2159Fc$

three-pole anti-alias filter with a cut-off of about 300 kHz is used. The power consumption of the complete synthesizer, implemented in the HCMOS logic family, is 760 mW. With no large-scale integration, the circuit is made on a single-sided surface-mount circuit board 10 x 19 cm in size.

Performance

Figure 3 shows the output spectrum of the synthesizer, programmed for a frequency of 131770 Hz. The plot is from zero frequency to 500kHz, the Nyquist limit. The output is taken before the anti-alias filter. Table I shows the amplitudes and numerical structure of the major spurious lines in the spectrum. The most prominent lines apart from the desired output are the second and third harmonics. A few m/n intermodulation products are visible at a level of about -80 dBc and below. An unexpected line which is an extremely high order mixing product with $m=1+2^{14}$ is seen at -83 dBc.

In our application, the critical frequency range for spurious signals is the region of the main resonance line, corresponding to frequencies between 120 and 145 kHz. The required synthesizer output frequencies are between 131.0 and 132.5 kHz. A computer program was used to find which orders of intermodulation would generate significant spurious sidebands. The output spectrum was then examined in detail to find the amplitudes of these products. Figure 4 shows how the sideband amplitude varies with order n up to 100. The lowest significant order is 37, at which the amplitude has fallen to about -90 dBc.

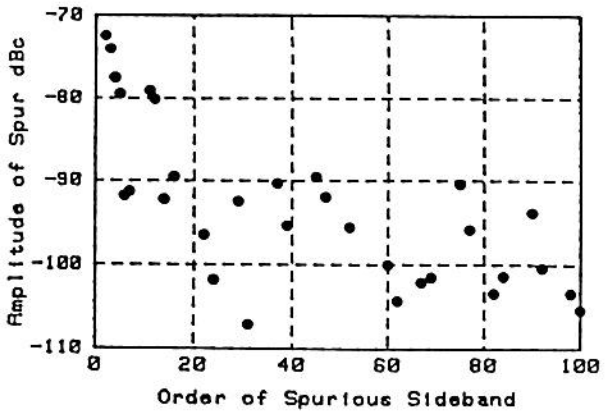

Fig 4. The observed variation of the amplitude of spurious products with their intermodulation order n. The frequency of these products is given by ABS($m \cdot Fc - n \cdot Fs$) where m and n are integers. m runs form 0 to 13.

Figure 5 shows the frequency region around a 131770 Hz output signal with greater resolution. A pair of spurious sidebands at $Fs \pm (38Fs - 5Fc)$ is present at a level of about -90 dBc. At the level of about -93 dBc no truncation sidebands are seen for this addend value.

Figure 6 shows, for comparison, the same frequency range as Figure 5 but with the DAC de-glitching disabled. A large number of sidebands have appeared with amplitudes up to -66 dBc. Careful examination shows that these are m/n products. Clearly the de-glitching is useful.

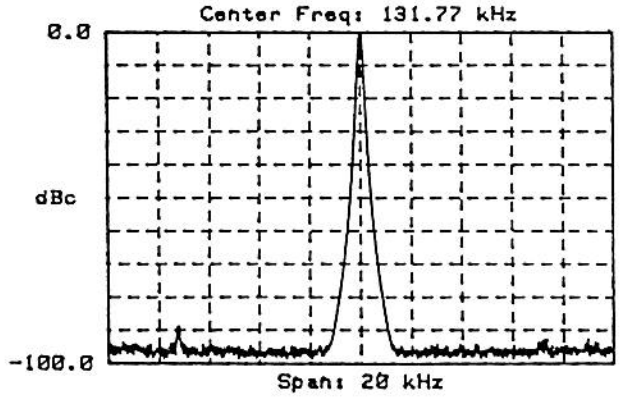

Fig 5. Output spectrum within ± 10 Khz of a programmed output of 131770 Hz. The effective noise bandwidth of the spectrum analyzer is 100 Hz. The frequencies of the spurious sidebands are given by $5Fc - 37Fs$ and $38Fs - 5Fc$.

Summary

A DDFS consisting of a 34-bit phase accumulator feeding a 14-bit wide address map and a 12-bit de-glitched DAC has been evaluated. For output frequencies in the range 131.0 to 132.5 kHz, which are required to excite the main Cesium resonance line, the level of spurious sidebands is less than -85 dBc. This is compatible with the expected performance of the DAC. No truncation sidebands have been seen down to -90 dBc. A set of sidebands with frequencies given by $2154Fc - (2^{14}-1)Fs$ and $(2^{14}+1)Fs - 2154Fc$ are detectable with an amplitude of -83 dBc. The origin of these sidebands is not understood, and they do not occur at significant output frequencies.

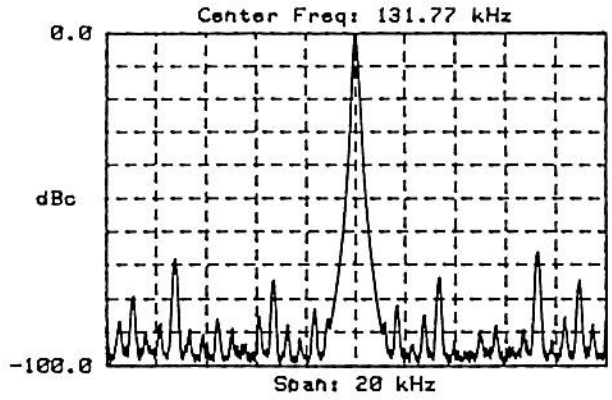

Fig 6. Same as Fig 5, with DAC de-glitcher disabled. The spurious sidebands visible are m/n intermodulation products with values of m and n between 37 and 260.

References

[1] J. Tierney, C.M.Rader, and B. Gold, "A digital frequency synthesizer" IEEE Trans. Audio Electroacoust., Vol AU-19, pp. 48-56, Mar. 1971.

[2] A.L. Bramble, "Direct digital frequency synthesis", in Proceedings of the 35th Annual Symposium on Frequency control, pp 406-414, IEEE, 1981.

[3] V. Reinhardt, K. Gould, K. McNab, and M. Bustamante, " A short survey of frequency synthesizer techniques", in Proceedings of the 40th Annual Frequency Control Symposium, pp 355-365, IEEE 1986.

[4] H.T. Nicholas, III and H. Samueli, "An analysis of the output spectrum of direct digital frequency synthesizers in the presence of phase-accumulator

truncation", in Proceedings of the 41st Annual Frequency Control Symposium, pp 495-502, IEEE 1987.

[5] H.T. Nicholas, III, H. Samueli, and B. Kim, "The optimization of direct digital frequency synthesizer performance in the presence of finite word length effects", in Proceedings of the 42nd Annual Frequency Control Symposium, pp 357-363, IEEE, 1988.

[6] J.F. Garvey, and Daniel Babitch, "An exact spectral analysis of a number controlled oscillator based synthesizer", in Proceedings of the 44th Annual Frequency Control Symposium, pp 511-521, IEEE, 1990.

[7] E.M.Matisson and L.M.Coyle, "Phase noise in direct digital synthesizers", in Proceedings of the 42nd Annual Frequency Control Symposium, pp 352-356, IEEE, 1988.

[8] L.S. Cutler and R.P. Giffard, "Architecture and algorithms for a new cesium beam frequency standard electronics", to be published in Proceedings of the 1992 IEEE Frequency Control Symposium, IEEE 1992.

[9] R. Karlquist, "A new RF architecture for Cesium frequency standards", to be published in Proceedings of the 1992 IEEE Frequency Control Symposium, IEEE 1992.

[10] "Digital Signal Processing", by A.V. Oppenheim and R.W. Shafer, Prentice-Hall, Englewood Cliffs, 1974.

A NEW RF ARCHITECTURE FOR CESIUM FREQUENCY STANDARDS

Richard K. Karlquist
Hewlett-Packard Company, Santa Clara Division
5301 Stevens Creek Blvd., MS 52U/7, Santa Clara, CA 95052

ABSTRACT

Recent advances in cesium beam tube technology, modulation and signal processing, and control system techniques have decreased the errors due to these sources to a level comparable to the errors due to RF chains typically used in atomic frequency standards. Improved RF chains are necessary in future standards to obtain optimum performance. However, the implementation of conventional architectures has already been pushed to its practical limits, requiring substantial changes to achieve the necessary results.

A novel RF chain is described which overcomes many of the limitations of previous RF chains making possible cesium standards with unprecedented performance. The new architecture greatly reduces the four main problems in existing RF chains: excessive spurious sidebands, temperature-induced phase instability, inadequate isolation of internal and external subsystems, and significant modulator distortion.

A new cesium standard realizes higher performance by utilizing amplitude modulation in addition to frequency modulation and controls the RF level with a servo loop based on atomic interactions. The RF chain described here implements these techniques. It also has sufficient available drive power for the beam tube to allow considerably lower Ramsey cavity Q.

The new RF chain starts with a surface acoustic wave (SAW) oscillator (locked to a conventional quartz flywheel oscillator) which is split into two signals. One signal drives a relatively short frequency multiplier chain ending at the tube. The other signal drives a direct digital synthesizer (DDS) that provides the user's outputs. The frequencies of the outputs are programmable and have 120 dB of output to output isolation.

The microwave output section uses a 9192 MHz dielectric resonator oscillator (DRO), which is phase locked to the low frequency signals. Frequency modulation is applied at this point by adding a programmable frequency offset to the phase locked loop. A linear amplitude modulator has been added to support the new servo systems. This results in a spectrum with extremely low spurious sidebands.

Introduction

The RF architecture used in commercial cesium frequency standards has been relatively unchanged over the last decade or two. However, recent advances in understanding of cesium beam tubes (CBTs) and computer controlled servo systems have created a need for a corresponding improvement in RF architectures. New servo techniques using stepped frequency and amplitude modulation greatly decrease errors due to incorrect RF and C-field amplitude, CBT cavity detuning, and effects such as Rabi pulling. [1] This has created a need for a new RF architecture to implement the new modulation scheme and have reduced error contribution commensurate with the CBT improvements. The principal error mechanisms in the RF chain include insufficient spectral purity, phase instability over temperature and modulation distortion.

Requirements

Error budget

Current commercial CBTs are accurate to at best 1×10^{-13} fractional frequency error. In order that the RF chain contribute negligible additional error, it is reasonable to limit the principal errors in the electronics to an order of magnitude below this, or 1×10^{-14}, which will be assumed in the discussion below.

Modulation

The new servo techiques employ precise steps of frequency and amplitude to interrogate the central and adjacent lobes of the principal Ramsey line, plus the Zeeman resonance approximately 40 kHz. away. The frequency needs to settle to better than 1×10^{-14} in 3 ms to obtain maximum measurement time utilization. 3 ms is available without contributing additional dead time because it takes that long for transients related to CBT atomic transit time to die out. [1]. The amplitude steps require about 2 dB of amplitude modulation capability and compensation of CBT cavity detuning loss requires about 4 dB additional range (based on a Q of 900 over 0 to 50 degrees C).

The frequency and amplitude stepping imposes a settling time requirement that arises from the need to minimize any dead time in the frequency measurement process. Following a step change in the rf frequency, the output of a

typical CBT shows a transient lasting about 3.5 ms caused by transit time effects in the cavity and detector. This transient is allowed to settle out before the tube output is measured to determine the frequency error. The microwave signal phase must have settled accurately in about 1.5 ms, however, because the average time taken for an atom to pass from the entrance of the cavity to the detector is about 2ms. Detailed calculation shows that if the slowest time-constant associated with phase-locked loops in the frequency modulation path is shorter than 150 μseconds, the error due to settling will be negligible. Accurate settling for these purposes is defined as frequency settling to within 1×10^{-14} and amplitude settling to within .01 dB. [2].

Spectral purity

Any sideband power on the interrogating signal that is not balanced will cause the servo to move off the center of the atomic resonance. This applies to both phase noise and discrete spurious sidebands (spurs.) If the sidebands are perfectly balanced, they will cause equal and opposite pulling and cancel out. However, it is poor engineering practice to depend on perfect balance because there are many AM-to-PM or PM-to-AM conversion processes that can unbalance the sidebands. Mistuning of the CBT cavity (or auxiliary filters) induces PM to AM conversion by slope detection. These effects worsen at increasing offset from the carrier and with increasing loaded cavity Q. It is impractical to keep the cavity tuning under servo control because present tuning diodes have insufficient Q and/or variation of capacitance.

For error budget purposes, it is convenient to divide the spectral purity requirements into three categories according to the offset from the carrier: the spurs outside the Zeeman band, the spurs in the Zeeman band, and close-in spurs related to the modulation processes. For spurs outside the Zeeman band (more than 130 kHz offset from the carrier) the allowable unbalanced sideband level for an error contribution of 1×10^{-14} varies linearly with offset frequency from -37 dBc. at 130 kHz to -20 dBc at 10 MHz offset. However, it is good engineering practice to keep all spurious sidebands down to -40 dBc to prevent them from competing with the carrier for power. For spurs inside the Zeeman band, an upper bound has been established that corresponds to about -70 dBc for a pulling error of 1×10^{-14} [3,4].

For spurs that are coherent with the rate of measurements (about 82 Hz in the new design to be described below), the effect of coherent modulation will cause errors that correlate from measurement to measurement instead of averaging to zero, which is the case for 120 Hz spurs. A mechanism by which the RF chain could be modulated at the measurement rate is the presence of 82 Hz noise on the power supply. This can be caused by periodic switching of bus drivers, multiplexers, etc. by the microprocessor. The effect is software dependent and can be minimized by proper coding. In the worst case, a fractional frequency deviation of 1×10^{-14} at the measurement rate will translate to a 1×10^{-14} error for the standard. For a 9.2 GHz carrier, this corresponds to an absolute peak deviation of 92 μHz, so the modulation index (β) for incidental 82 Hz modulation must be held to about 1×10^{-6} . This corresponds to an 82 Hz bright line at a maximum level of -126 dBc. Note that since the 10 MHz quartz flywheel oscillator is ultimately multiplied by a factor of 928, any 82 Hz spurs on the 10 MHz signal will increase by 20 Log 928 = 59 dB so they must not exceed a level of -185 dBc at 10 MHz.

Phase stability

When the standard is required to operate in environments with temperature transients, phase instabilities in the RF chain can result in frequency errors in the standard. The are also self-induced temperature transients during warm-up. These transients last longer than the time it takes for modern commercial cesium standards to warm up and operate to high accuracy. [5] This error mechanism arises because angular frequency, by definition, is the time rate of change (derivative) of phase. If phase is a non-constant function of temperature, and temperature is changing with time, then phase will change with time. This represents a frequency offset. If the standard takes 1 hour to reach thermal equilibrium resulting in an internal temperature rise of 15 degrees, then there is a temperature ramp of about .004 degrees per second (taking a linear approximation.) A 10 MHz output circuit with a temperature coefficient of phase of 150 μradians per degree C would drift at a rate of 0.6 μradians per second. Since there are 2π radians per cycle, this is equivalent to a frequency offset of 0.1 μcycles per second, or 0.1 μHz, which is a fractional frequency error of 1×10^{-14}. The phase stability requirement of 150 μradians per degree C at 10 MHz can also be expressed as a delay stability of 2.5 ps per degree C, which is independent of frequency. This would be 125 ps over 0 to 50 degrees C.

Conventional architecture

In order to provide a proper context for the description of the new architecture, a brief review of prior art will be presented. Figure 1 shows the block diagram of a representative example (Hewlett-Packard model 5061B.) A 10 MHz quartz oscillator is divided by two to obtain 5 MHz to which phase modulation is applied. This is multiplied to 90 MHz with transistor multipliers which then drive a step-recovery-diode (SRD) comb generator to produce harmonics of 90 MHz. The 102nd harmonic at 9180 MHz is used as a carrier that is phase modulated at a rate of about 12.6 MHz to produce a sideband at about 9192.6 MHz. A filter attenuates the carrier and spurious sidebands

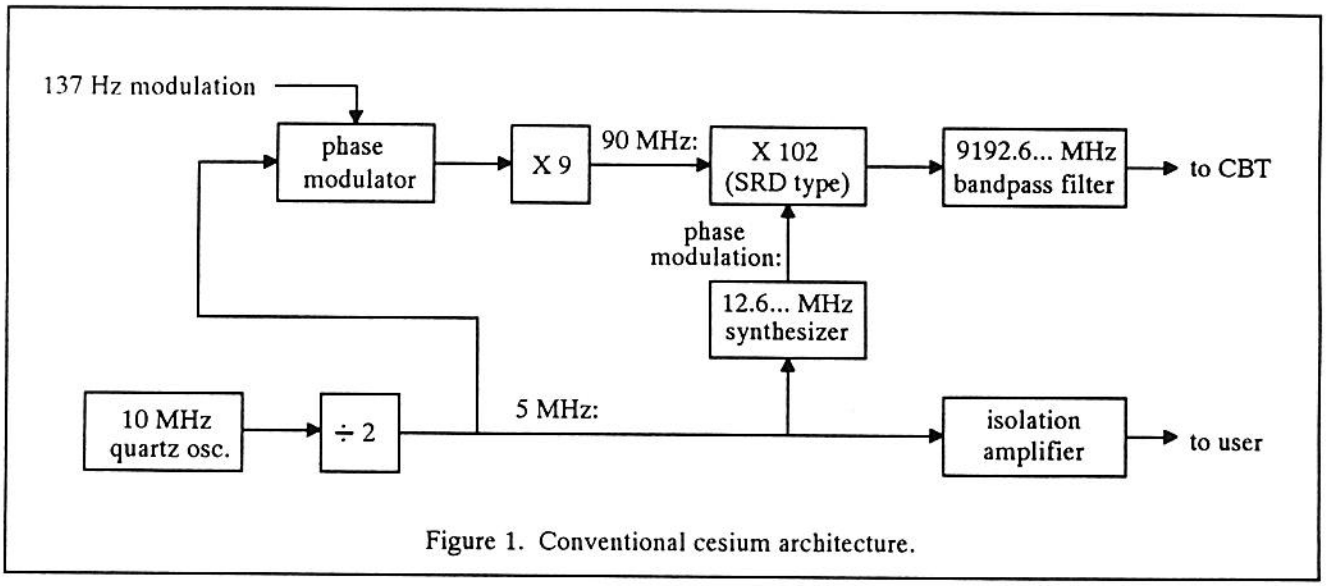

Figure 1. Conventional cesium architecture.

leaving just the 9192.6... MHz sideband to excite the cesium resonance in the CBT. The sideband is generated by modulating the SRD bias voltage with a high resolution synthesizer locked to the quartz oscillator. The center of the resonance is found by detecting the response of the CBT to the audio FM with a sychronous demodulator. The quartz oscillator signal is buffered and supplied as the user's output.

Most cesium standards and some non-cesium standards are similar to this. The deficiencies of this approach are difficulties with spurious sidebands, insufficient power output, no convenient or accurate way to do amplitude modulation, excessive phase modulation distortion, poor phase stability with temperature, and undesireable manufacturability and reliability characteristics.

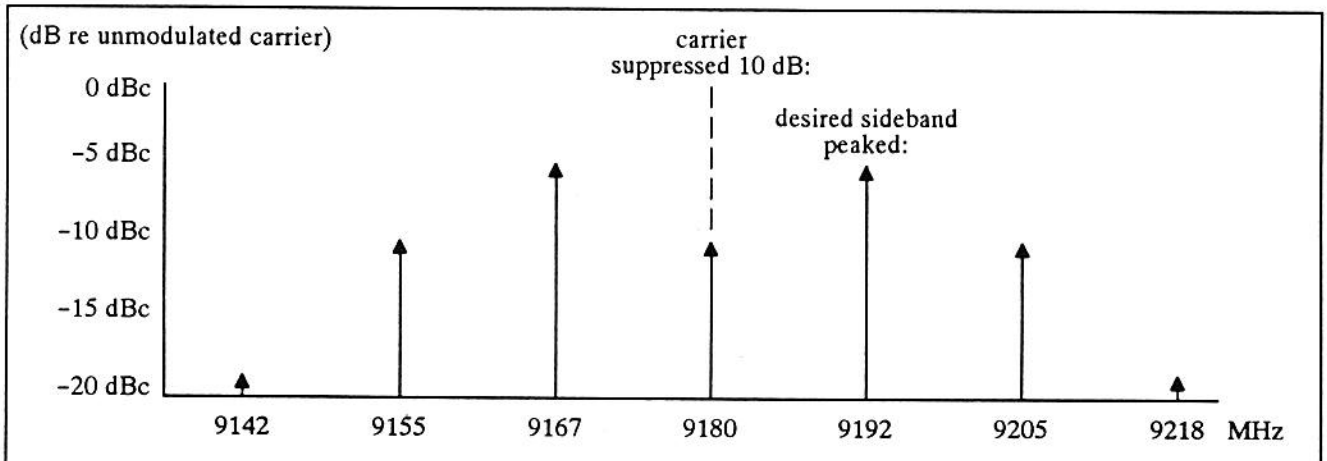

Figure 2. Typical microwave spectrum before filtering (conventional architecture)

Microwave subsystem

Figure 2 shows the microwave spectrum after the SRD multiplier and before the microwave bandpass filter. 0 dB. refers to the unmodulated carrier level. The modulation index (β) is set to the value (1.85) that maximizes the first sidebands (at 9192.6... and 9167.4... MHz.) Under this condition, the desired sideband is down 5 dB from the unmodulated carrier. The carrier is suppressed 10 dB below its unmodulated level. This leaves the carrier down only 5 dB relative to the desired sideband. The second harmonic sidebands (25.3 MHz from the carrier) are also down only 5 dB relative to the desired sideband. The carrier and second harmonic sideband must be suppressed an additional 14 dB in order to achieve an error contribution of 1×10^{-14}. This requires a high-Q auxiliary filter made of

invar in combination with a high-Q CBT cavity. The auxiliary filter adds to the complexity and the high-Q CBT cavity requires temperature compensation.

Manufacturability

The conventional architecture dictates extensive use of waveguide construction. This leads to practical problems with corrosion, contact resistance, and difficult-to-make adjustments. Also, there are many opportunities for RF to leak out and the atmosphere to leak in. the It is difficult to prevent stray 9.2 GHz fields from impinging on the drift region in the CBT and cause errors due to extraneous quantum transitions. Atmospheric leakage can cause sensitivity to humidity and exascerbate corrosion problems in areas with air pollution.

Modulation

The architecture of Figure 1 does not include amplitude modulation capability, because conventional servo systems did not use it. Adding it would require insertion of a waveguide-type PIN diode modulator between the SRD multiplier and the bandpass filter. This is inconvenient and runs the risk of thermal transients in the PIN diodes causing phase transients that result in frequency errors. The analog phase modulator would have to have an impractically low second harmonic distortion of -120 dBc in order to hold its error contribution below 1×10^{-14} with a conventional servo system. [1]

Phase stability

The conventional architecture suffers from poor phase coherence between the signal at the CBT and the user's output. This is due to phase shifts in the initial multiplier stages and the long chain of isolation amplifiers necessary to prevent output loading from affecting the accuracy of the standard. The initial doubler from 5 to 10 MHz must have considerable filtering to attenuate spurious output signals at 5 and 15 MHz before driving the next stage, because these would result in 5 MHz sidebands on the CBT drive signal. In order to keep these sidebands down to -40 dBc after multiplication to 9.2 GHz, it is necessary keep the 5 and 15 MHz spurious signals down to -90 dBc. Phase drift in the filter results in a frequency shift when the ambient temperature is changing. Additional phase drift occurs in the path from the quartz oscillator to the CBT.

New RF chain

The new architecture is shown in Figure 3. [6] There are four main areas of difference. The 9192.6 MHz signal for the CBT is generated by a dielectric-resonator oscillator (DRO) running directly at that frequency. The signal path now splits between the CBT chain and the output chain at the 320 MHz point rather than the 10 MHz point, and a direct-digital-sythesizer (DDS) is used to generate the user's outputs. The modulation type and technique is different to accommodate the needs of the new servo loops.

The new architecture, like the conventional one, retains a 10 MHz quartz oscillator that is controlled by the servo loop. However, it is more convenient to think of the signal flow originating with the 640 MHz surface-acoustic-wave (SAW) oscillator which drives a frequency divider to generate 320 MHz. (A SAW oscillator frequency of 640 MHz rather than 320 MHz was chosen because 640 MHz SAW resonators were readily available, not for any fundamental reason.) The 320 MHz signal splits into two paths, one leading to the CBT and the other leading to the user's output. To excite the CBT, the 320 MHz signal is multiplied by 29 to produce a reference at 9280 MHz. This is used in an offset phase locked loop to tune the DRO to 9192.6... MHz, which is the difference frequency between the 9280 MHz reference and the upconverted agile synthesizer (near 87.4 MHz). The output of the DRO is amplitude modulated by the servo loop before driving the CBT. To produce the user's output, the 320 MHz signal is also divided by 4 to clock a DDS at 80 MHz. This produces sine wave outputs at 5 and 10 MHz after only minimal filtering. The short term stability of the SAW oscillator is inadequate for an atomic frequency standard, hence it is locked to the 64th harmonic of the 10 MHz quartz oscillator to clean up close-in phase noise.

Microwave subsystem

A detailed block diagram of the microwave subsystem is shown in Figure 4. The 320 MHz is multiplied using an SRD and the 29th harmonic is extracted by a 9280 MHz bandpass filter. This is mixed with the DRO frequency to produce an 87 MHz intermediate-frequency (IF). A limiting IF amplifier is used to get a reliable input for the phase detector that compares the IF frequency with the 87.4 MHz reference from the agile synthesizer. A PLL with 1 MHz bandwidth tunes the DRO to the correct frequency. The output of the DRO is amplitude modulated by the servo loop using a mixer before being sent to the CBT.

The DRO has a power output of over 10 mW, which allows it to drive the local-oscillator (LO) ports of the two mixers. Since the SRD drives the RF port of a mixer, only a small amount of SRD output power is needed. Only -20 to -10 dBm is needed to overcome the IF noise floor. Hence a simple, non-critical, low-efficiency SRD circuit can be used. The limiting IF can tolerate wide variations in SRD output. Also, the 9280 MHz filter is not critical; only minimal filtering of the comb spectrum is necessary to keep from overloading the mixer with the combined power of dozens of harmonics. Nearby spurious signals

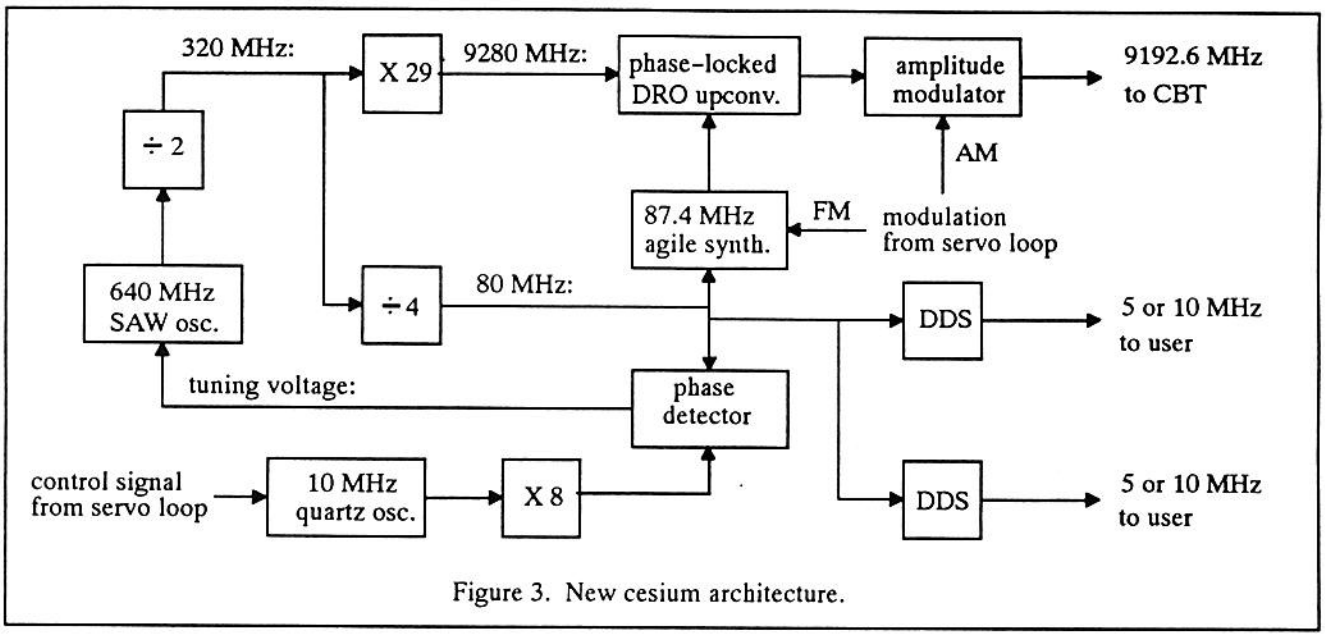

Figure 3. New cesium architecture.

such as the 28th and 30th harmonic can be tolerated because they are outside the bandwidth of the PLL. Insertion loss of the filter is also non-critical because it can be made up with extra gain in the IF amplifier. Hence it is acceptable to implement the filter with coupled-line microstrip construction.

The DRO also drives the LO port of the modulation mixer, which is a conventional single-balanced diode mixer with a DC-coupled IF port. Since the DRO is driving the diodes hard enough to act as switches, there is little AM-to-PM conversion or thermal effects in the diodes. There is an inherently linear up-conversion process from the IF port to the RF port so the same mixer can serve as the stepped amplitude modulator and also the average power controller (to compensate for RF amplitude changes due to CBT cavity drift) without interaction. Also, if the IF port is driven with a 50 ohm source impedance, then the RF port will automatically have very close to a 50 ohm source impedance as well. This contrasts to the extreme difficulty of maintaining low VSWR at the ports of a PIN diode modulator.

Another feature of the DRO is that sidebands above 1 MHz on the 320 MHz or 87.4 MHz inputs are attenuated by the PLL. This relaxes the spec for 10 MHz sidebands on the 320 MHz signal and 21.9 MHz sidebands on the 87.4 MHz signal. This coupled with the fact that the

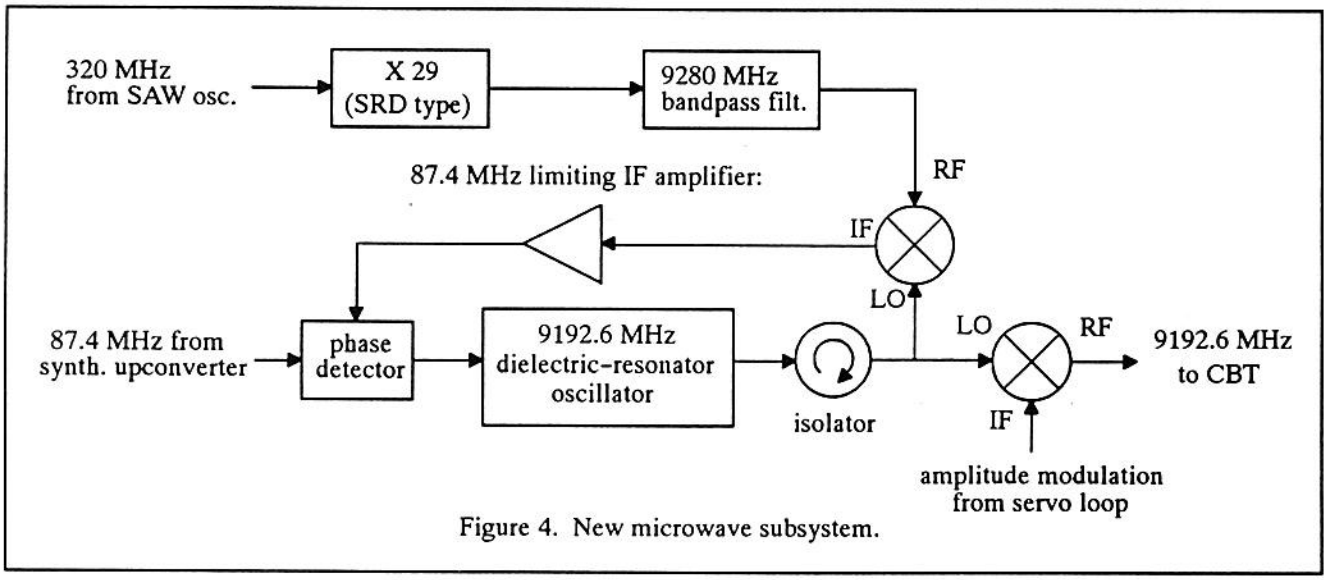

Figure 4. New microwave subsystem.

DRO frequency is used to excite the atoms directly (rather than with a sideband) eliminates spectral clutter.

Saw oscillator

In designing this part of the architecture, a number of alternatives were investigated. It would be possible to multiply directly to 320 MHz using five cascaded frequency doublers. However, good engineering practice would dictate that all spurious sidebands (such as 10 MHz) be held to 40 dBc after the 320 MHz is multiplied by 29, which requires 70 dBc at 320 MHz. The 10 MHz sidebands on the output of the first doubler would be enhanced by 24 dB by the subsequent four doublers hence requiring 95 dBc at 20 MHz. This would require extensive filtering and multiple shielded chambers.

Starting with a signal at 320 MHz reduces parts count and simplifies shielding to two chambers: one for the SAW oscillator and one for everything else.The noise floor of the SAW oscillator is lower than the multiplied-up noise floor of the quartz oscillator. Any spurious sidebands on the quartz oscillator more than a few kHz. away from the carrier, such as 200 kHz from the switching power supply, are stripped off by the PLL. Any drift in the phase of the quartz oscillator with respect to the SAW oscillator will affect the CBT drive signal and the user's output equally and hence drop out of the error budget as a common mode error.

Output subsystem requirements

A fresh look was taken at the output amplifier problem and a new set of requirements generated that were commensurate with the high performance in the rest of the standard. It was decided that two isolated high performance outputs were necessary. One could be permanently connected to the user's distribution system. The other one would then be available for testing, development work, phase comparison, etc. In order to make this scheme useful, the two outputs should be isolated by at least 100 dB so that connecting test equipment to one does not appreciably affect the phase of the other output that is being distributed. Also, the phase of the two outputs should track very closely. The outputs should be extremely well behaved to minimize interactions with the distribution system: the harmonic distortion should be below 1% (-40 dBc), the source impedance should be 50 ohms, not just at the output frequency, but harmonics as well. This will assure that any reflected signals from the user's equipment including distortion generated by phase detector inputs will be absorbed by the output rather than reflected back to the user. This is a common problem in existing hardware. The output circuit should have enough dynamic range to avoid voltage clipping when driving an open circuit and current clipping when driving a short circuit. A short circuit load would arise if an electrical quarter wave of coax (about 5 meters at 10 MHz) were connected from the output to a high impedance load. Finally, as mentioned above, the delay (phase) drift due to temperature should ideally be held to 2.5 ps per degree C or 125 ps for an ambient change of 50 degrees, to keep the phase drift error below 1×10^{-14}.

Output subsystem design

Figure 5 shows a block diagram of an output subsystem that meets these requirements. The 80 MHz clock signal is divided by 8 or 16 to produce a 10 or 5 MHz. digital ramp. The ramp drives a DAC to produce a 4-times oversampled sine wave as shown in Figure 6. If the DAC were perfect, there would only be outputs at 5, 35, 45, 75, 85, etc. MHz (for a 10 MHz output). In practice, suppression to about -50 dBc is reliably obtained. A first order hold after the DAC reduces the amplitude of the 70 and 90 MHz signals to -17 and -19 dBc respectively before filtering. (See Figure 7) This allows much less filtering than would be necessary for a square wave, reducing phase drift in the output filters. Extremely high speed, temperature-compensated logic with 500 ps propagation delay is used to keep delay drift below 125 ps over temperature in the digital sections. When the output is at 5 MHz, the first harmonic requiring filtering is at 35 MHz. This allows a common reconstruction filter to be used for either frequency by designing it to pass frequencies below 10 MHz and cut frequencies above 35 MHz . Hence the output frequency can be programmed to be either 5 or 10 MHz simply be changing the divide ratio. A dissipative filter is used to provide a 50 ohm output impedance at the output frequencies and all harmonics. A continuously-counting auxiliary divide-by-16 is used to assure phase coherency between the two outputs even if the output frequencies are reprogrammed. The digital nature of the frequency dividers and DACs provides for high reverse isolation with low complexity.

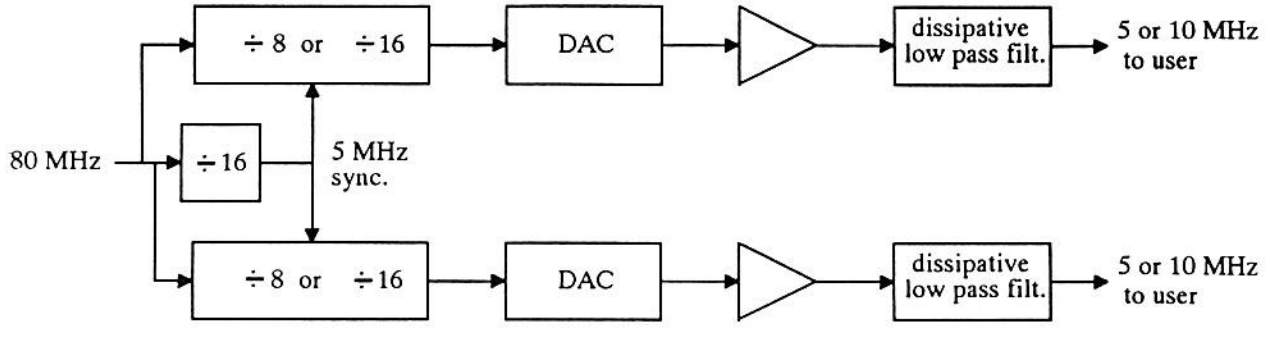

Figure 5. New output subsystem.

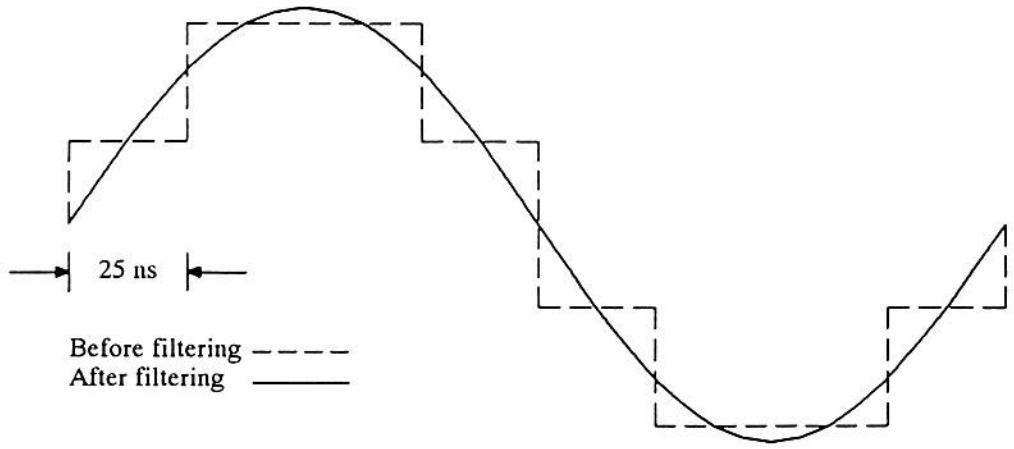

Figure 6. Output waveform before and after reconstruction filter (5 MHz output).

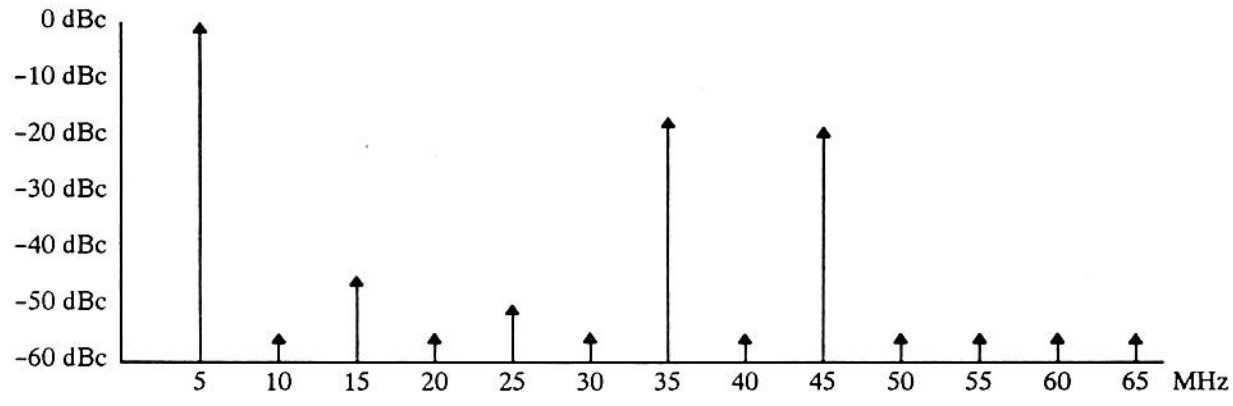

Figure 7. Output spectrum before reconstruction filter (5 MHz output).

Synthesizer upconverter

A high resolution agile synthesizer [2] with an output from 125 to 175 kHz (nominal frequency, about 131 kHz) is used to generate the frequency modulation needed by the servo loop. [1] The synthesizer upconverter raises the synthesizer frequency to around 87.4 MHz, where it is upconverted again by the microwave subsystem described above. In both cases, it is not multiplied, which would require additional resolution by the DDS.

Figure 8 shows the block diagram of the synthesizer upconverter. A 21.837 MHz VCXO with a tuning range of at least .1% is quadrupled to 87.368 MHz and sampled at 2.5 MHz by an SRD sampler. This has the effect of mixing it with the 35th harmonic of 2.5 MHz (87.5 MHz) to produce an IF frequency of 125 to 175 kHz. A phase detector compares the IF with the DDS signal to phase lock the VCXO.

Because the VCXO loop has only 3 msec. to settle to within 92 μHertz, a rather wide bandwidth of 10 kHz is necessary. The most severe requirement occurs when the DDS jumps about 40 kHz after interrogating the Zeeman line for the C-field servo. This requires very careful analysis of loop acquisition and settling time characteristics. The maximum loop bandwidth is limited by spurious crystal responses located a few tenths of a percent above series resonance. These degrade loop stability by adding extraneous gain and phase shift to the VCXO tuning voltage input.

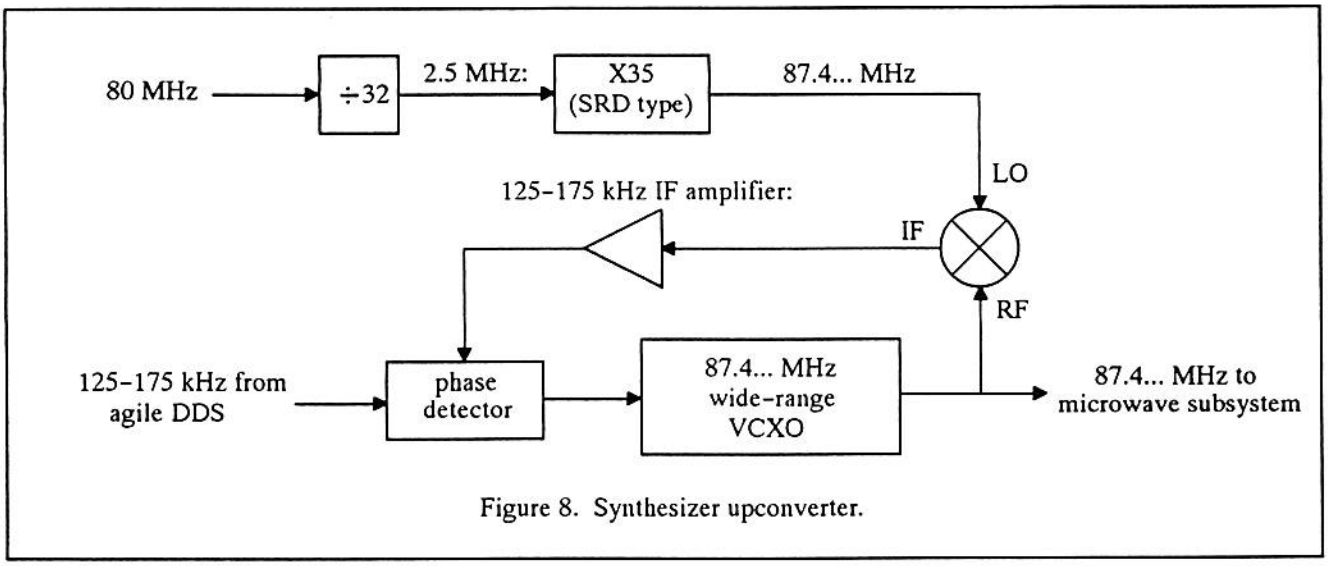

Figure 8. Synthesizer upconverter.

Summary

A design has been shown which meets the objective of keeping errors due to RF problems such as spectral impurities, modulation distortion, and phase instability below a budget of 1×10^{-14}. The resulting RF chain is more manufacturable and repeatable than previous ones. Demonstrated performance with time and varying environmental conditions is excellent [5].

ACKNOWLEDGEMENTS

L. Cutler suggested the microwave mixer technique for applying amplitude modulation. R. Giffard contributed to the design of the power conditioning and the microwave system design. They also acted as technical advisers. J. Kusters assisted with writing and presenting this paper. T. Parisek did the mechanical design of the RF shielding. J. V.Brzeski characterized some of the microwave hardware and contributed to the synthesizer upconverter.

REFERENCES

[1]: L.S. Cutler and R.P.Giffard, "Architecture and algorithms for a new cesium beam frequency standard electronics", to be published in Proceedings of the 1992 IEEE Frequency Control Symposium, IEEE 1992.

[2]: R.P.Giffard and L.S. Cutler, "A low-frequency, high resolution digital synthesizer", to be published in Proceedings of the 1992 IEEE Frequency Control Symposium, IEEE 1992.

[3]: J.H. Shirley, "Some causes of resonant frequency shifts in atomic beam machines", J. Appl. Phys., 34, pp. 783-791, April 1963.

[4]: C. Audoin, M. Jardino, L.S. Cutler, and R.F. Lacey, "Frequency offset due to spectral impurities in cesium-beam frequency standards" IEEE Transactions on Instrumentation and Measurement, Vol IM-27 No. 4, pp. 325-329, December 1978.

[5]: J. Kusters and J. Johnson, "A new Cesium beam standard: performance data", to be published in Proceedings of the 1992 IEEE Frequency Control Symposium, IEEE 1992.

[6]: R. Karlquist, U.S. Patent Pending, filing date Aug. 1, 1991.