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

134

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

135

Figure 1. Conventional cesium architecture.
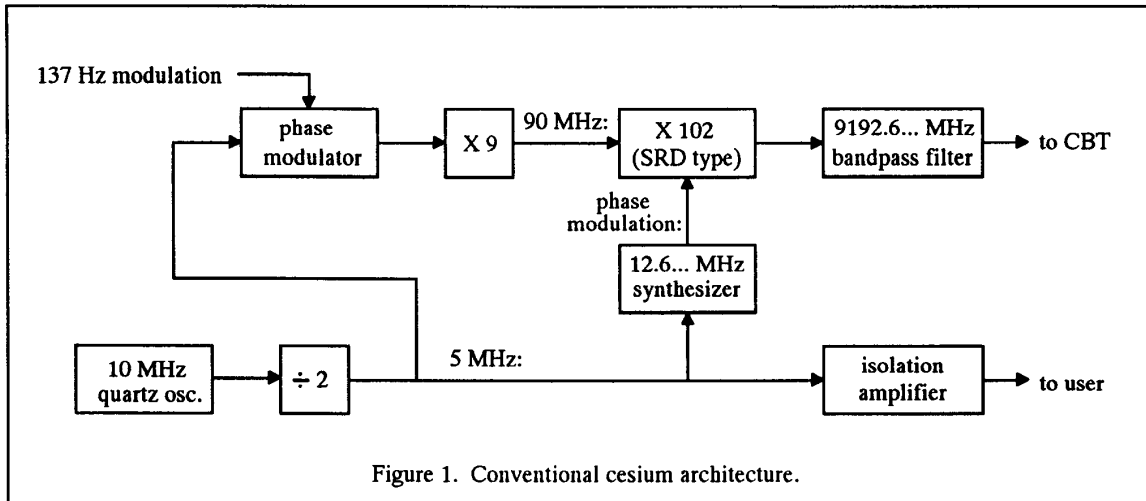
leaving just the 9192.6... MHz sideband to excite the cesium resonance in the CBT. The sideband is generated by modulating the SRD bias voltage with a high resolution synthesizer locked to the quartz oscillator. The center of the resonance is found by detecting the response of the CBT to the audio FM 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)
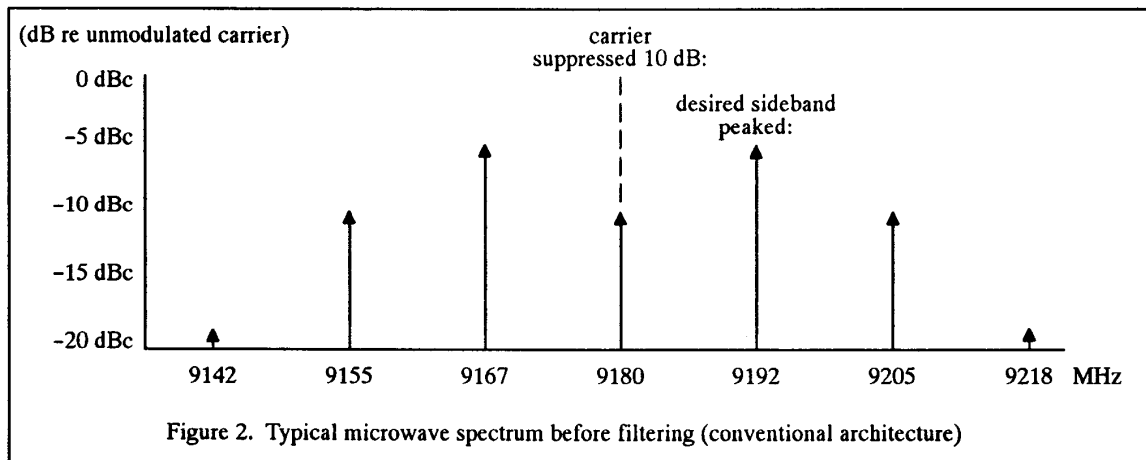
Microwave subsystem

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

136

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

137

Figure 3. New cesium architecture.
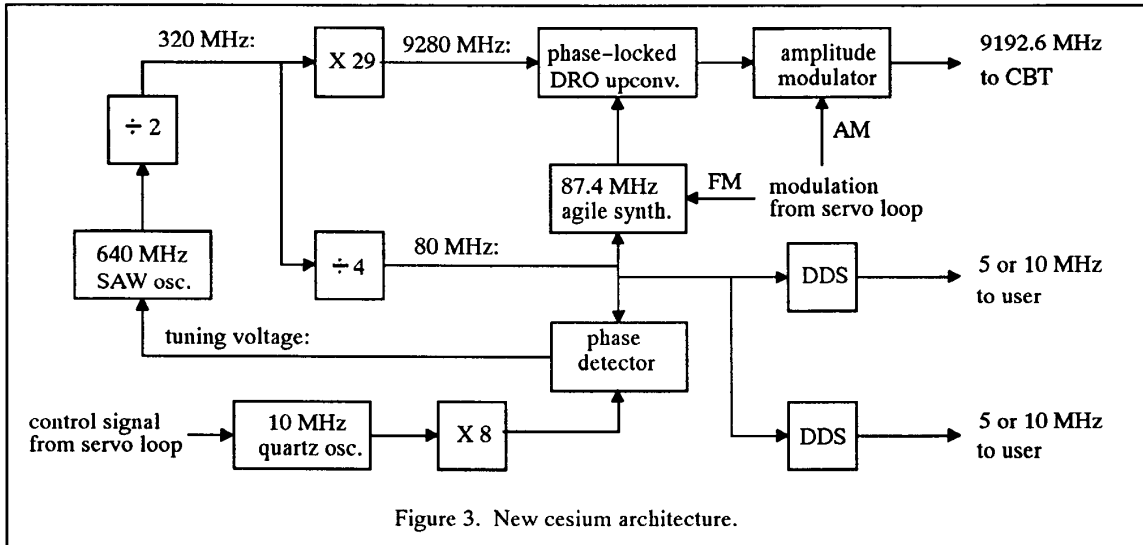
such as the 28th and 30th harmonic can be tolerated because they are outside the bandwidth of the PLL. Insertion loss of the filter is also non-critical because it can be made up with extra gain in the IF amplifier. Hence it is acceptable to implement the filter with coupled-line microstrip construction.

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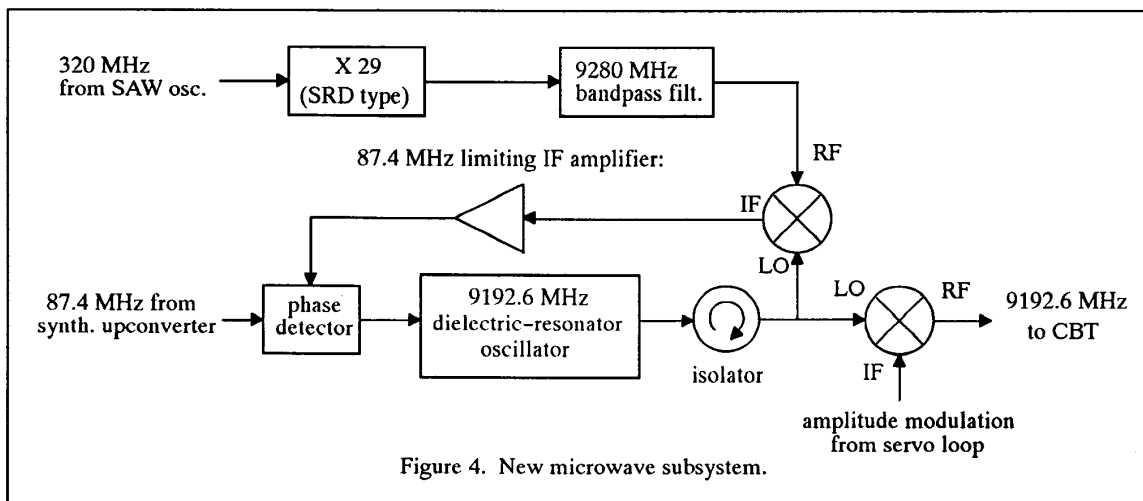
DRO frequency is used to excite the atoms directly (rather than with a sideband) eliminates spectral clutter.

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.
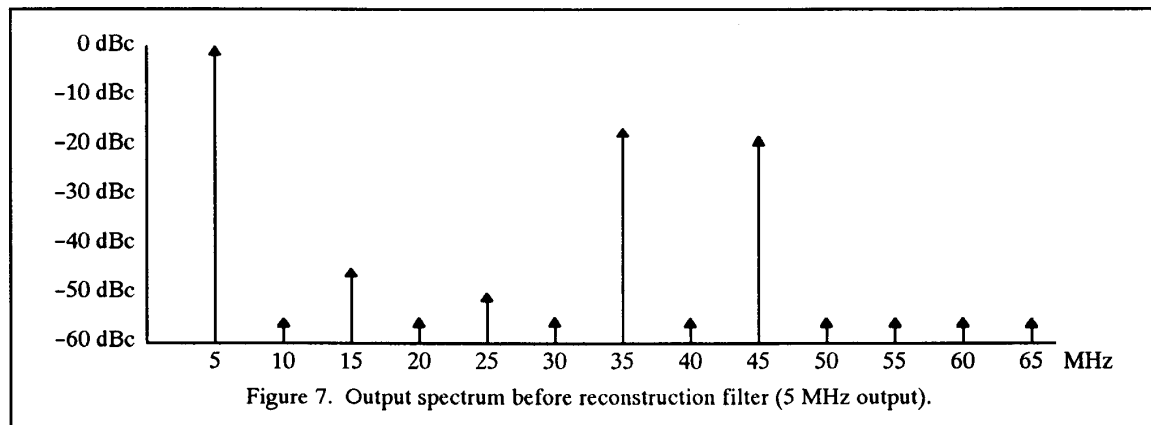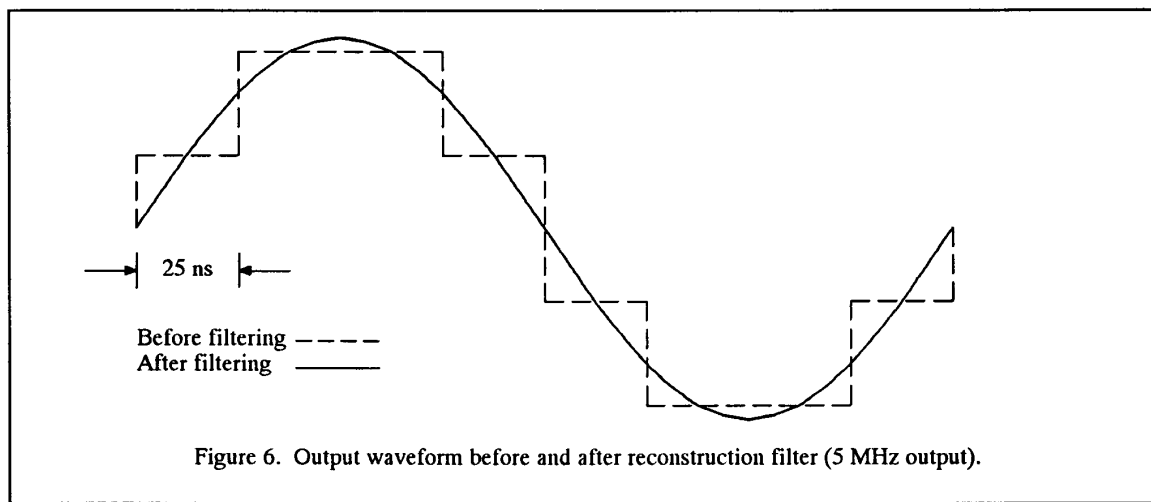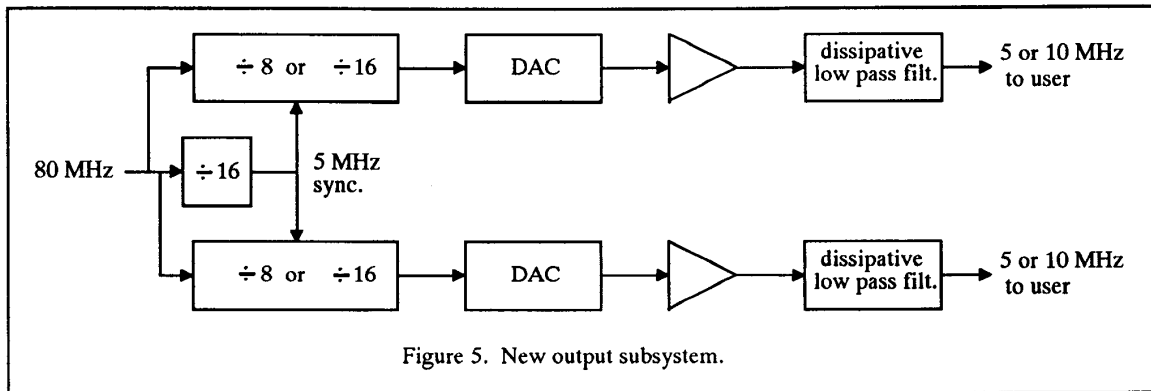
Output subsystem requirements

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 imped-

ance 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.

139

Figure 5. New output subsystem.

25 ns

Before filtering — — — —
After filtering ————

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.
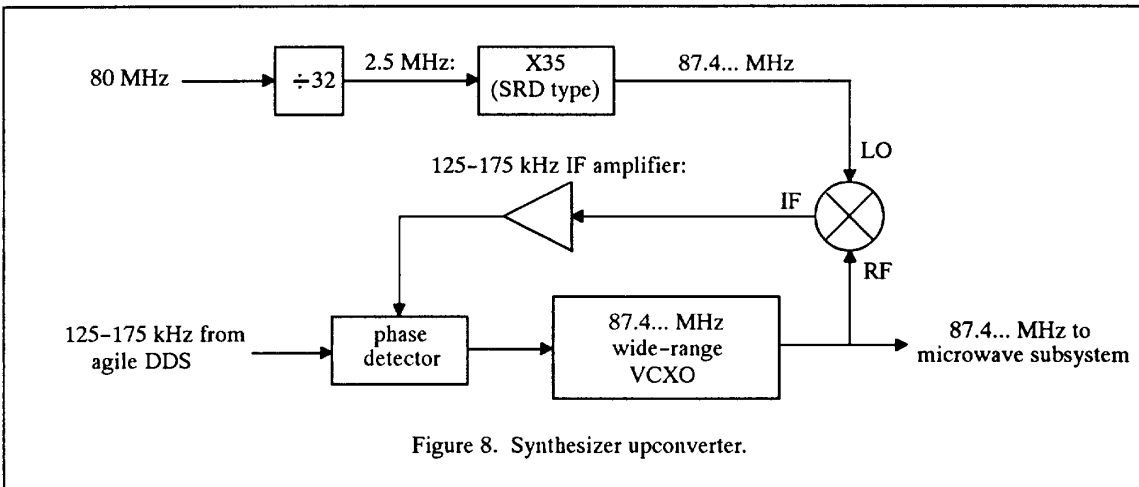

Figure 8. Synthesizer upconverter.

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].

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.