MICROCOMPUTER–COMPENSATED CRYSTAL OSCILLATOR FOR LOW–POWER CLOCKS

S. S. Schodowski, R. L. Filler,
J. A. Messina, V. J. Rosati and J. R. Vig
U. S. Army Electronics Technology and Devices Laboratory
Fort Monmouth, New Jersey 07703–5000

Abstract

Temperature–compensated crystal oscillators (TCXO) are widely used in applications requiring precision frequency control or timing, especially when power availability is limited and the temperature range is wide. The best overall accuracy in wide temperature range (e.g., −55°C to +85°C) TCXO has remained at about 1 ppm for the past 30 years because of thermal hysteresis, the trim effect, and thermometry inaccuracies.

In this paper a microcomputer–compensated crystal oscillator (MCXO) is described. This MCXO employs external compensation, i.e., there is no frequency "pulling". It thus permits the use of a low–hysteresis, low deviation–sensitivity SC–cut crystal in a non–trimmed oscillator. Resonator self–temperature–sensing, using a dual–mode oscillator, virtually eliminates thermometry–related errors. As a consequence, all basic TCXO limitations are overcome or significantly reduced. The MCXO incorporating this technology is capable of providing at least 10– to 100–times improvement in overall frequency accuracy when compared to conventional TCXO. Milliseconds–per–day timekeeping accuracy, with an input power of less than 40 milliwatts, is now available.

INTRODUCTION

Temperature–compensated crystal oscillators (TCXO) are widely used in applications requiring precision frequency control or timing, especially when power availability is limited and the temperature range is wide. The best overall accuracy in wide temperature range (e.g., −55°C to +85°C) TCXO has remained at about 1 ppm for the past 30 years because of thermal hysteresis, the trim effect, and thermometry inaccuracies.

Conventional TCXO are based on a method of "pulling" the crystal frequency, typically via a varactor, to counteract the crystal's frequency *vs* temperature (f *vs* T) deviation. This approach requires a high frequency–deviation–sensitivity resonator, i.e., a fundamental mode AT–cut, which usually possesses thermal hysteresis in the 0.2 ppm to 1.0 ppm range. The high deviation–sensitivity, however, also makes the TCXO susceptible to significant frequency error resulting from the thermal hysteresis and aging of circuit components. TCXO trim effect is the degradation (rotation) of the compensated f *vs* T characteristic when the output frequency is adjusted to correct for frequency aging. This

effect, which cannot easily be corrected by ordinary circuit means, typically produces a 0.1 to 0.5 ppm error for a frequency adjustment range of ±3 ppm. Temperature sensing in a TCXO relies on one or more thermistors placed close to the resonator. This method suffers from inaccuracies due to thermal lag (thermal path differences and unequal crystal and thermistor thermal time constants), thermal gradients across the crystal and between the crystal and thermistor (thermistor location and self–heating), and thermistor instability.

In this paper we describe a microcomputer–compensated crystal oscillator (MCXO). This MCXO employs external compensation, i.e., there is no frequency "pulling". It thus permits the use of a low– hysteresis, low deviation–sensitivity SC–cut crystal in a non– trimmed oscillator. Resonator self–temperature–sensing, using a dual–mode oscillator, virtually eliminates thermometry–related errors. As a consequence, all basic TCXO limitations are overcome or significantly reduced. The MCXO incorporating this technology is capable of providing at least 10– to 100–times improvement in overall frequency accuracy when compared to conventional TCXO. Milliseconds–per–day timekeeping accuracy, with an input power of less than 40 milliwatts, is now available. More detailed descriptions of the MCXO can be found in references 1 to 6.

RESONATOR SELF–TEMPERATURE–SENSING

Simultaneous excitation of a resonator in two modes provides a means for the resonator to sense its own temperature. Previously, the only means available for resonator self temperature–sensing was the b–mode of an SC–cut crystal when used in a dual b– and c–mode oscillator. However, the b–mode is not well–behaved over a wide temperature range due to a high incidence of frequency and resistance anomalies, i.e., activity dips.[1] In the MCXO, a different resonator self–temperature–sensing method is being applied.[2–4] This method uses a pair of c–modes of an SC–cut resonator in a dual c–mode oscillator. By employing an SC–cut, the c–modes are almost guaranteed to be free of activity dips over the full operating temperature range.[1]

Self–temperature–sensing is based on the "harmonic effect" as illustrated in Figure 1. The harmonic effect is the change in the apparent orientation angle between harmonically related modes of a resonator. As shown, the effect progressively diminishes with increasing harmonic number M. For SC–cut resonators, the effect primarily manifests itself as a change in the first–order temperature coefficient of the polynomial

$$\frac{\Delta f_m}{f_m} = a_M \Delta T + b_M \Delta T^2 + c_M \Delta T^3 \qquad (1)$$

where the normalized frequency, $\Delta f_M / f_M$, and the difference temperature, ΔT, are referenced at, e.g., 25°C, and a_M, b_M, and c_M, are the first–, second–, and third–order temperature coefficients of frequency at harmonic number M.

THE THERMOMETRIC BEAT FREQUENCY

In the illustration of Figure 2, the fundamental–mode frequency, f_1, is multiplied by three and mixed with the third overtone frequency, f_3, to obtain a beat frequency f_β, i.e.,

$$f_\beta = 3f_1 - f_3 \tag{2}$$

Alternatively, another beat frequency can be defined as

$$f'_\beta = f_1 - f_3/3. \tag{3}$$

Using the f *vs* T dependence of the two modes as given by Equation (1), the beat frequency may be described in normalized form as a function of temperature

$$\frac{\Delta f_\beta(T)}{f_\beta} = \frac{3a_1 - na_3}{3-n}\Delta T - \frac{3b_1 - nb_3}{3-n}\Delta T^2 - \frac{3c_1 - nc_3}{3-n}\Delta T^3, \tag{4}$$

where a_1, b_1, c_1, and a_3, b_3, c_3 are the temperature coefficients of frequency for the fundamental mode and third-overtone, respectively; n is the ratio of the frequencies for the harmonic pair at the reference temperature, i.e., $n = f_3/f_1$. A noninteger n is conveniently available, since the frequencies of harmonic modes of actual resonators are rarely exact multiples of each other. Values of n ranging from 2.94 to 2.98 are typical for a variety of SC-cut designs.[1] Equation (4) then yields the temperature coefficients for the beat frequency. The first-order term dominates the higher-order terms for at least 100° above and below 25°C. Therefore, when a pair of modes is simultaneously excited in a dual-mode oscillator and combined a thermometric beat frequency is obtained, which can be used in an MCXO for compensating either of the c-mode frequencies.

A typical resonator's fundamental-mode and third-overtone frequencies are shown plotted against temperature in Figure 3. The apparent angle-of-cut change is evidenced by the shift of turning point from +5°C to approximately +45°C.

Figure 4 shows the temperature dependence of the beat frequency for the same SC-cut crystal unit. The f *vs* T curve is monotonic and nearly linear. For example, for a −55°C to +85°C temperature range, the maximum deviation from the best-fit straight line is about two percent of the total frequency change. The high degree of linearity permits approximating the slope by its first order temperature coefficient, which is -96 ppm/°C. Since n is usually in the range of 2.94 to 2.98, the temperature sensitivity ranges from -60 ppm/°C to -150 ppm/°C.

MCXO RESONATOR

Dual c-mode SC-cut resonators that are suitable for self temperature-sensing and MCXO application have been successfully developed.[1] The major MCXO resonator requirements are:

a. f *vs* T characteristic that is free of significant anomaly,
b. minimum f *vs* T slope over the operating temperature range,
c. high Q, and
d. small hysteresis.

Plano-convex resonator designs of the proper contour were shown to be capable of meeting all requirements. Figure 5 shows the f *vs* T characteristics for an optimum-design resonator for a −55°C to +85°C operating temperature range. An optimum resonator f *vs* T characteristic is the one for

which the maximum f *vs* T slopes of the two modes are minimized. This condition is satisfied when the absolute value of the slope of the fundamental mode at the upper temperature extreme equals the slope of the third–overtone at the lower temperature extreme. For the −55°C to +85°C range, the absolute values of the maximum slopes are 2.7 ppm per °C, which occur when the turnover temperature of the third–overtone is at about +20°C. The exact condition for optimum f *vs* T is that the maximum f *vs* f_β slopes be minimized; however since f_β *vs* T is nearly linear, the difference between the two conditions is small.

FREQUENCY *vs* TEMPERATURE *vs* ANGLE OF CUT CHARACTERISTICS

It has been shown that the necessary angle–of–cut tolerances can be readily achieved with conventional cutting technology. For example, for the optimum −55°C to −85°C MCXO resonator (lower turnover temperature ≈ +20°C), a cutting error of 5 minutes results in an 8°C shift in the lower turnover temperature. For such an MCXO resonator, the rate of change in f *vs* T slope with lower turnover temperature is 0.043 ppm/°C^2 at −55°C. Therefore, a -5 minute error increases the lower turnover temperature to +28°C, which increases the slope at −55°C from the optimum 2.7 ppm/°C to 3.0 ppm/°C. A +5 minute error causes a smaller maximum slope change at −85°C, since the rate of change in f *vs* T slope is 0.029 ppm/°C^2 at −85°C. Therefore, in many MCXO applications, the angle–of–cut tolerance requirement may be loose enough for x–ray orientation and angle–correction to be unnecessary.

MCXO AND TCXO RESONATOR COMPARISON

Table I shows a comparison between resonators for the MCXO and for precision analog TCXO. Not only is the angle–of–cut tolerance looser for MCXO resonators, but so are the blank frequency and plating tolerances. For example, in one MCXO implementation,[3] pulse deletion is used to generate an accurate time–corrected pulse train. The resonator target frequency during plating is chosen to ensure that the resonator frequency is above the nominal clock frequency at all temperatures. Therefore, there is no need to specify a tight plating tolerance. In fact, the accuracy achievable with "rough" plating may be sufficient; i.e., no frequency adjustment should be necessary if the rough plating is reasonably well controlled.

The hysteresis and aging are lower for MCXO resonators than for TCXO resonators because third–overtone SC–cut resonators are inherently more stable than the fundamental–mode AT–cut resonators usually used in wide–temperature–range TCXO. Another possible factor may be that the interface between the rough plating and the fine plating can be eliminated in MCXO resonators.

Hysteresis is the major limitation on the f *vs* T stability that is achievable with the MCXO. Although the majority of resonators exhibited hysteresis at the low 10^{-8} level, some were in the 10^{-9} range. This shows that MCXO's with an f *vs.* T stability in the 10^{-9}, or better, range is a reasonable goal for the future. For that goal to be met, further research is needed to gain a better understanding of the mechanisms responsible for hysteresis.

THE DUAL–HARMONIC–MODE CRYSTAL OSCILLATOR (DHMXO)

Figure 6 shows one example of a dual–mode oscillator capable of generating the two c–mode frequencies. This DHMXO can be described as two separate oscillators sharing a crystal that operates at series resonance. A single section resembles that of the Butler or bridged–tee oscillator, and falls into the family of crystal oscillators characterized by the crystal current being essentially equal to the transistor emitter current. In this configuration, the crystal is grounded to facilitate design. In each gain loop, emitter degeneration provides negative feedback that is sufficient to reduce net loop gain below unity for all frequencies except the desired c–mode frequency, where the crystal branch impedance is a minimum. This condition is assured by the series tuned networks C_s, L_s, and C_s', L_s', which also serve to decouple the adjacent mode signal. Using this circuit, the coupled frequencies can be attenuated more than 50 dB below the primary output signal level. Positive feedback at each c–mode is provided by the phase–shifting PI–networks consisting of C_1, C_2, and L_3 and C_1', C_2', and L_3'.

Simultaneous excitation of the c–modes can also be obtained with the single–gain–loop Colpitts–type DHMXO shown in Figure 7. Suppression of undesired SC–cut b–mode frequencies is accomplished by the network C_2, L_2, C_2^1, L_2^1, which is designed to appear capacitive, and which provides the correct phase shift for oscillation only at the two c–mode frequencies. Inherent nonlinearities in the transistor provide frequency multiplication and mixing such that the beat frequency is available at the collector output following filtering of the higher frequency components. Either the fundamental mode or third–overtone frequencies can be extracted at the emitter using a frequency selective amplifier. The single–gain–loop DHMXO offers lower parts count and lower input power than the double–gain–loop circuit. However, it does trade off the flexibility of independent control of crystal–current for each mode, and it limits stability coefficient optimization. These limitations can prevent the single–gain–loop DHMXO from achieving the high level of performance of the double–gain–loop design. However, use of the single–gain loop may be warranted in applications permitting somewhat lower temperature–sensing accuracy.

IMPLEMENTATION OF THE DUAL C–MODE METHOD

Implementation of the dual c–mode method is shown by the MCXO simplified block diagram of Figure 8. In this illustration, the thermometric beat frequency is the one defined in Equation 3. f_β' is divided by M to obtain a convenient value for gating a reciprocal counter with the third–overtone frequency as the input. The reciprocal counter produces a number, $N1$, that represents the temperature of the crystal. The divisor M serves to scale the $N1$ counter range, consistent with the required temperature–sensing resolution and the number of counter bits. This number, or temperature word, is read by the microcomputer and is used to compute a second number, $N2$, that programs the correction circuitry. Correction algorithms may consist of evaluation of a polynomial, or a look–up table using coefficients or table entries, respectively, that were stored in memory during the calibration process.

The authors' laboratory is sponsoring two contracts, using two different approaches, for implementation of the MCXO, at Frequency Electronics, Inc. (FEI), and at General Technical Services (GTS).

THE FEI MCXO

The basic MCXO in development at FEI employs pulse deletion.[3] The SC–cut resonator frequency, f_1 or f_3, is selected so that, at all temperatures, it is always greater than the required output frequency. This ensures that pulses can always be deleted from the pulse train to provide the required number of output pulses in each interval. If the calculated correction includes a fraction of a pulse, the remainder is stored in memory, and an additional pulse will be deleted when the remainder, which is adjusted in subsequent cycles, has reached or exceeded a full pulse.

Figure 9 is a more detailed digital control block diagram. There are two inputs, f_β ($\approx$ 157 kHz) and f_1 ($\approx$ 3.3+ MHz). At the start of each measurement interval both the 12 stage f_β counter and the 18 stage f_1 counter are reset to zero and started simultaneously. The microcomputer is set to detect 36 overflows of the f_β counter to get a 0.94 second gate time ($36 \times 2^{12}/157$ kIIz = 0.94 s). The f_1 counter then contains the number N_c which is the remainder of $f_1/2^{18}$, i.e., f_1 modulo 2^{18}. The crystal resonator is manufactured to have a f *vs* T characteristic such that, over the intended temperature range, the number of overflows is always the same. For example, the resonator frequency can be between $12 \times 2^{18} = 3.145728$ MIIz and $13 \times 2^{18} = 3.407872$ MHz which is a $\Delta f/f$ of 262 kHz (80,000 ppm, not a major manufacturing dilemma). The microcomputer uses N_c to compute N_d, the number of pulses that must be deleted during the next measurement interval. The variable–divide counter is loaded with N_d/f_1 so that a pulse is deleted each time the variable–divide counter overflows. The net result is a pulse train with average frequency f_o and missing pulses uniformly spaced in time.

Start–up time, i.e., the time interval between application of power and the first valid correction (pulse deletion), is about 3.5 seconds. The output signal, f_o, is a time–corrected pulse train that can be divided to produce a 1 pps time reference, or can be used directly to drive a clock. Because of the noise characteristics created by the pulse deletion process, additional signal processing is necessary to provide rf output useful for frequency–control applications. This can be accomplished, for example, by imparting the MCXO accuracy to a low–noise, low–cost, voltage controlled crystal oscillator via a phase–locked–loop (PLL). A prototype MCXO, with a phase–locked VCXO operating at 5 MHz, showed a phase noise of -100 dBc at 100 Hz and -140 dBc at 100 kHz.

Construction of the basic MCXO consists of an SC–cut lateral field crystal unit, custom hybrid dual–mode oscillator, CMOS microcomputer, and a 2,550–gate application specific IC (ASIC). The power consumed in each of the major assemblies for a compensated f_1 output and one–second update is:

Dual–mode oscillator	8 mW
Microcomputer	18 mW
ASIC	15 mW
	41 mW

By extending the time between updates, e.g., to 60 seconds, the average input power can be made to approach the idle power of the MCXO, which is about 25 mW.

An even lower average input power, approaching 5 mW, can be obtained by combining the MCXO with a low–power clock.[6] In this approach, illustrated in Figure 10, the MCXO is turned off, except when calibrating the low power clock. If the duty cycle is sufficiently low, the total input power will be almost independent of MCXO power. As in the previous approach, temperature variations between calibrations will affect the frequency, reducing the accuracy of the clock. Thus, there will be a trade–off between input power and accuracy, as shown in Figure 11. The low–power clock can reduce the

input power to 5 mW, but accuracy continues to decrease, reaching 75 milliseconds in 24 hours.

The FEI MCXO uses a Motorola MC1468705G2 microprocessor that has 2K of EPROM and 112 bytes of RAM.

THE GTS MCXO

In the MCXO developed at GTS[4], f_3 is used to generate a correction frequency with a direct digital synthesizer. The correction frequency is added to the third overtone frequency in a PLL to produce a temperature–stable 10 Mhz output. In an alternative time only mode, f_3 is used to generate a constant 4 kHz, which is divided to produce a 1 pps clock output.

Referring to Figure 12, and assuming that the MCXO is set to the "frequency" mode, operation is as follows:

The temperature–dependent beat frequency, f_β', is the output of the digital mixer, with f_1 the fundamental–mode frequency and f_3 the 3rd overtone frequency. The mixer is configured such that

$$f_\beta' = f_1/2 - f_3/6 = f_\beta/2 \tag{5}$$

The output of the mixer, f_β'', is divided by 2^{16} to produce a gate width of approximately 0.16 seconds for the 24 bit counter. The counter input is $f_3/2$, i.e., nominally 5 MHz. The number of f_β'' counts measured in the 0.16 seconds, N_1, is equivalent to the period of f_β. N_1 is used by the microcomputer to determine the necessary correction from the stored calibration data. The microcomputer computes the number N_2 which sets the direct digital synthesizer (DDS). The DDS is driven by $f_3/32$ ($\approx$ 312 kHz) and gives an output frequency of between 200 and 900 Hz. This frequency, Δf, is compared with the difference between a (10 Mhz, AT-cut) VCXO and f_3 in the PLL. The error signal controls the VCXO so that the difference frequency is that computed by the microcomputer, i.e., the output is 10 MHz.

The third–overtone oscillator can operate equally well at an offset of up to -2500 Hz, so the plating target frequency of the SC–cut resonator can be 10 Mhz -1600 Hz $\pm$ 0.5 ppm. Therefore, the resonator for the GTS MCXO can also be only rough–plated.

Since no pulse deletion is used, one would expect a clean spectrum. However, under the condition where the number N_2 is other than an integer power of 2, the DDS is likely to produce a periodic phase modulation that manifests itself as spurious sidebands on Δf. By judicious selection of the PLL bandwidth (approximately 10 Hz), the spurious sidebands will be suppressed and the VCXO output phase noise will be only slightly degraded in the region outside the loop bandwidth. Within the loop bandwidth the close–in noise is dominated by the short–term stability of the third overtone frequency. The phase noise for the current design (f_1 = 10 MHz) falls from -100 dBc at 10 Hz to -160 dBc at 50 kHz. The noise floor depends on the VCXO signal–to–noise ratio, i.e., a low noise floor requires a high crystal drive current. In the present VCXO, crystal current is about 1.5 mA.

To understand the operation of the "time" mode, refer again to Figure 12, with the switch set to "T". In this mode, the PLL, including the VCXO, is disabled to conserve power. Further, the third–overtone oscillator is activated only periodically, to update N_1. In the time mode, the DDS is driven by $f_1/16$ (211 kHz, nominally), and N_2 is calculated to deliver a constant 4 kHz output. The 4 kHz is divided by 4,000 to give 1 pps. The long–term accuracy capability of the MCXO in the "time"

mode is nearly the same as in the "frequency" mode, i.e., about $\pm 1 \times 10^{-8}$. However because of the resolution of the DDS in the "time" mode, there is up to 4.7 microseconds of jitter. (The jitter in the "frequency" mode is less than 50 nanoseconds.[7])

The GTS MCXO currently (December, 1989) uses an Intel 87C51FA microprocessor that has 8Kbytes of EPROM and 256 bytes of RAM on the chip. The internal EPROM holds the program, and the RAM is used as a scratchpad. A 128 byte EEPROM (an ICT93C46) is used to store the polynomials for the fundamental and third overtone frequencies. The aging–correction coefficient is also stored in the EEPROM, making aging correction nonvolatile.

Start–up time in either operating mode is approximately 2 seconds. In addition to the initial temperature measurement and correction computation periods, the start–up time in the frequency mode includes the acquisition time of the PLL, and in the clock mode, the period delay in providing the first 1 pps time tick in the divide by 4000 output stage.

MCXO PERFORMANCE TESTS

Four configurations of the MCXO[5] were tested:

1. A packaged FEI MCXO consisting of a dual–mode oscillator, a microcomputer, and a gate array.[3] We shall refer to this device as the MCXO, and report results for five samples.

2. A dual–mode oscillator design using a lateral–field resonator and hybrid circuit construction. The lateral field resonator was designed to eliminate the requirement for the b–mode suppression circuit in the oscillator. We will refer to this device as the LFR dual–mode oscillator, which is a component in the MCXO described in Configuration (1).

3. A dual–mode oscillator design using a thickness–field resonator.[4] We will refer to this oscillator as the TFR dual–mode oscillator. We tested two samples of the TFR dual–mode oscillator. Each sample was tested with two resonators, one each from two manufacturers.

4. The dual–mode oscillator from configuration (3) with internal microprocessor compensation. Testing was performed by GTS.[8]

F *vs* T PERFORMANCE

TEMPERATURE PROFILE

Details of the experimental setup and measurement methods can be found in reference no. 5.

MCXO F *vs* T

Figure 13 shows the results for a typical MCXO.

DUAL–MODE OSCILLATOR F *vs* T

It can be seen from Fig. 13 that the compensation can be improved with a simple "rotation," that is, a change in the first–order temperature coefficient in the compensation algorithm. The necessary

frequencies are available from the MCXO, so that compensation of the devices could be simulated on a desktop computer.

To perform the "simulated" compensation, we made a least–square fit of the set of data (f_β, f_1) to a 6th–order polynomial.

$$F(f_{\beta i}) = \sum_{p=0}^{6} C_p (f_{\beta i}) p \tag{6}$$

We then computed the residuals R_i

$$R_i = f_{1i} - F(f_{\beta i}). \tag{7}$$

A plot of the residuals *vs* temperature is equivalent to an MCXO f *vs* T stability plot, assuming a perfect technique for applying the calculated correction, e.g., pulse deletion or numerical synthesis.

LFR DUAL–MODE OSCILLATOR F *vs* T

Figure 14 shows a typical plot of the residuals from a 6th order polynomial for the 5 LFR dual–mode oscillators. The algorithm does not take into account thermal history, therefore, it cannot eliminate hysteresis.

Even including the hysteresis, the f *vs* T stability is $\pm 2 \times 10^{-8}$. The probable cause for the difference between the MCXO results and the simulated compensation is that the MCXO was compensated midway in the manufacturing process. Subsequent manufacturing steps caused a change in the required compensation coefficients, possibly by changing circuit strays. In the future, improved manufacturing techniques should allow realization of the simulated compensation.

TFR DUAL–MODE OSCILLATOR F *vs* T

Figure 15 shows a plot of the residuals of a TFR dual–mode oscillator. The excellent results shown indicate that it is possible to obtain very low hysteresis.

AGING

All five MCXO's were placed into a temperature chamber and put through the sequence depicted in Figure 16. The duration of the aging segments was approximately 30 days. Figure 17 shows aging data for a typical MCXO. The dashed line is a linear approximation of the last several days of data. The value of the slope of the dashed line is given in the figure.

F *vs* T REPEATABILITY

Figure 18 is the superposition of 18 f *vs* T runs on MCXO 16723 taken over the course of 13 months. Two different temperature ranges (−55°C to +85°C and −46°C to +68°C), two step sizes (2°C and

3°C), and three soak times (4, 12, and 30 minutes) were used. The performance of the MCXO appears to be independent of temperature range, soak time, and step size. Most important, there is no evidence of any f *vs* T aging, which would occur, e.g., if the aging of the two modes differed significantly.

SUMMARY AND CONCLUSIONS

A performance summary is given in Table II.

Whereas conventional TCXO accuracy has remained at about 1 ppm for the past 30 years, the MCXO now allows a temperature stability of 2×10^{-8} and an overall accuracy of $< 5 \times 10^{-8}$ for a −55°C to +85°C temperature range.

A thermometry method of using two c-modes in a dual-mode oscillator has been described. Its potential is now being realized in the development of high–performance MCXO having stabilities 10 to 100 times greater than that of analog TCXO. The method overcomes the limitations of existing temperature–compensation methods, is easy to implement, is microprocessor compatible, and makes possible a new generation of inexpensive, high–accuracy, low–power, temperature–compensated crystal oscillators.

REFERENCES

1. R. Filler and J. Vig, "Resonators for the Microcomputer Compensated Crystal Oscillator," Proc. 43rd Annual Symposium on Frequency Control, IEEE publication no. 89CH2690–6, 1989.

2. S. Schodowski, "Resonator Self–Temperature–Sensing Using a Dual–Harmonic–Mode Crystal Oscillator," Proc. 43rd Annual Symposium on Frequency Control, IEEE publication 89CH2690–6, 1989.

3. M. Bloch, M. Meirs and J. Ho, "The Microcomputer Compensated Crystal Oscillator (MCXO)," Proc. 43rd Annual Symposium on Frequency Control, IEEE publication 89CH2690–6, 1989.

4. A. Benjaminson and S. Stallings, "A Microprocessor–Compensated Crystal Oscillator Using an SC–Cut Third Overtone Resonator," Proc. 43rd Annual Symposium on Frequency Control, IEEE publication 89CH2690–6, 1989.

5. R. Filler, J. Messina and V. Rosati, "Frequency–Temperature and Aging Performance of Microcomputer–Compensated Crystal Oscillators," Proc. 43rd Annual Symposium on Frequency Control, IEEE publication no. 89CH2690–6, 1989.

6. M. Bloch, M. Meirs, J. Ho, J. Vig and S. Schodowski, "Low Power Timekeeping," Proc. 43rd Annual Symposium on Frequency Control, IEEE publication no. 89CH2690–6, 1989.

7. A. Benjaminson, private communication, October, 1989.

8. A. Benjaminson, private communication, December, 1989.

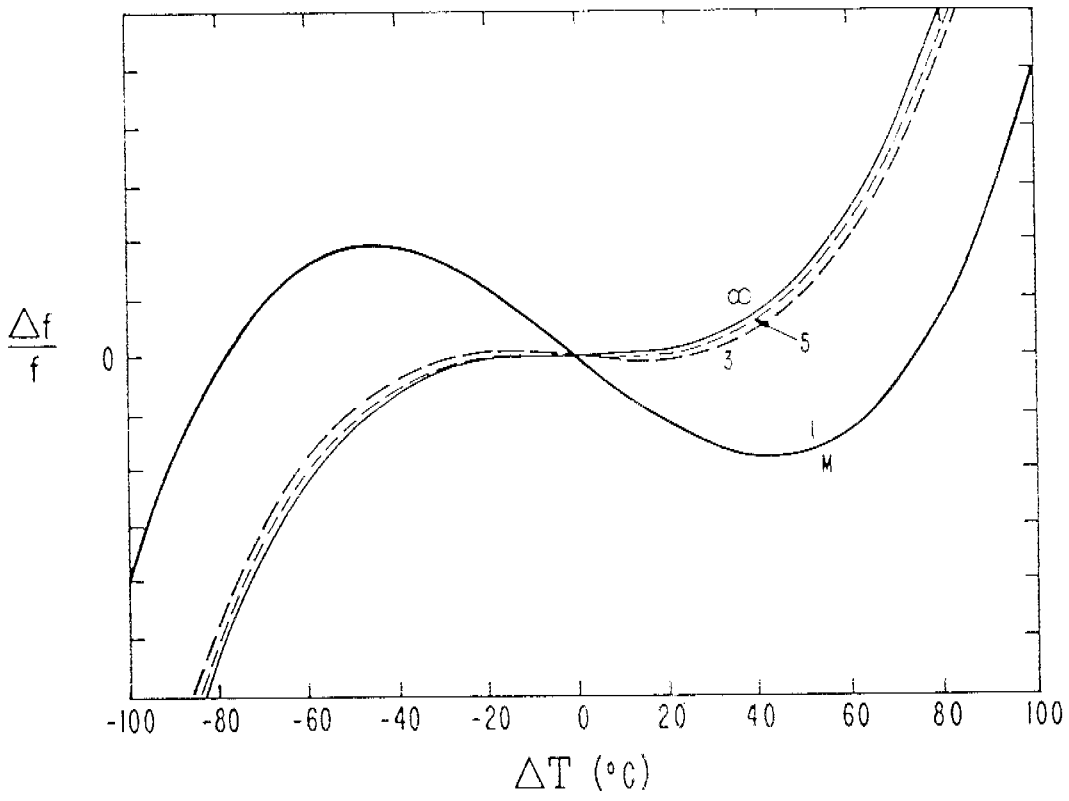

Figure 1. Frequency-temperature-harmonic characteristics of temperature-stable resonators.

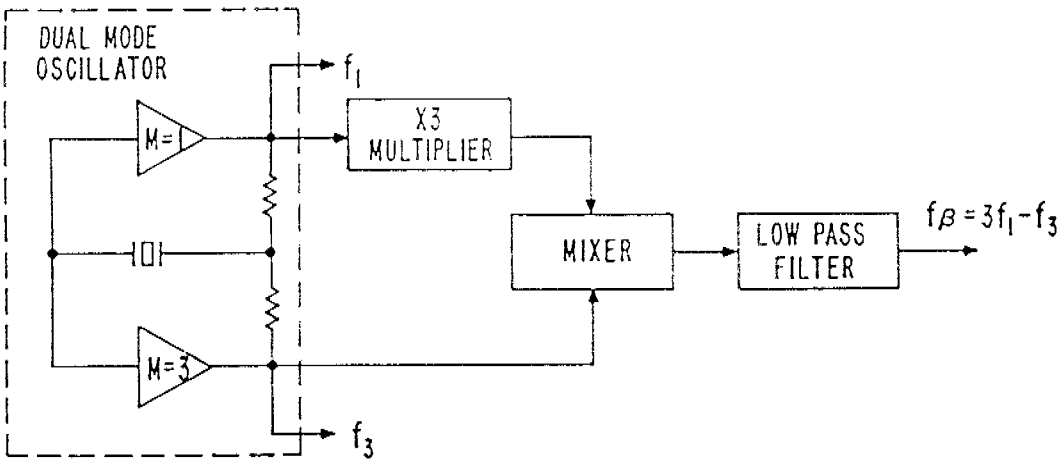

Figure 2. Illustration of the dual c-mode thermometry method.

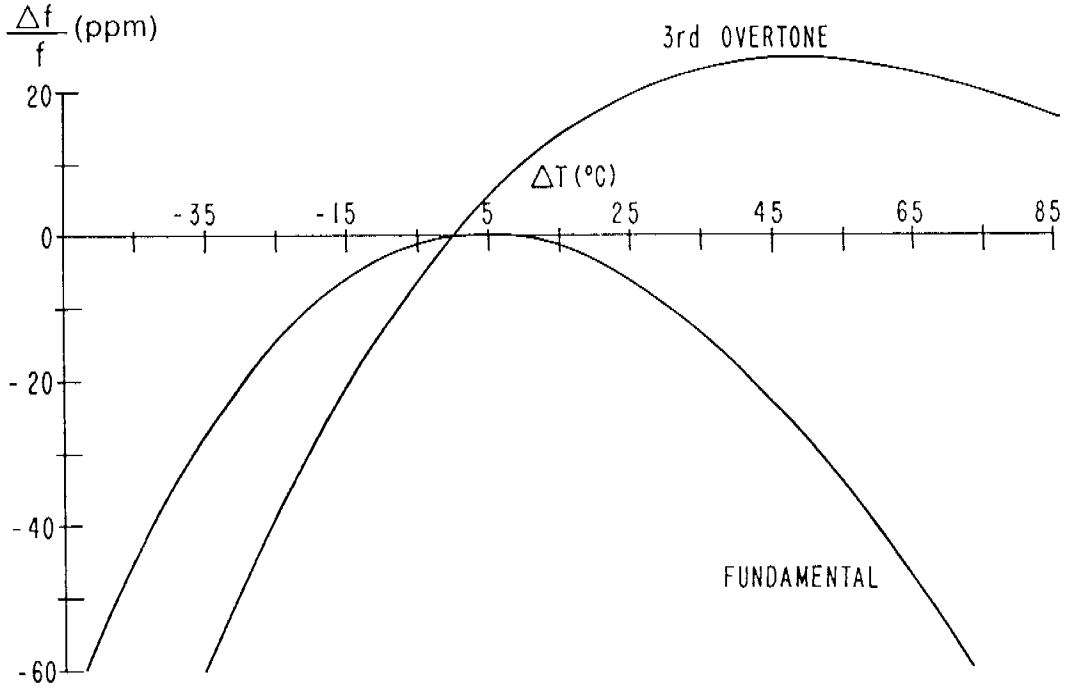

Figure 3. Fundamental and third-overtone characteristics of a typical SC-cut resonator.

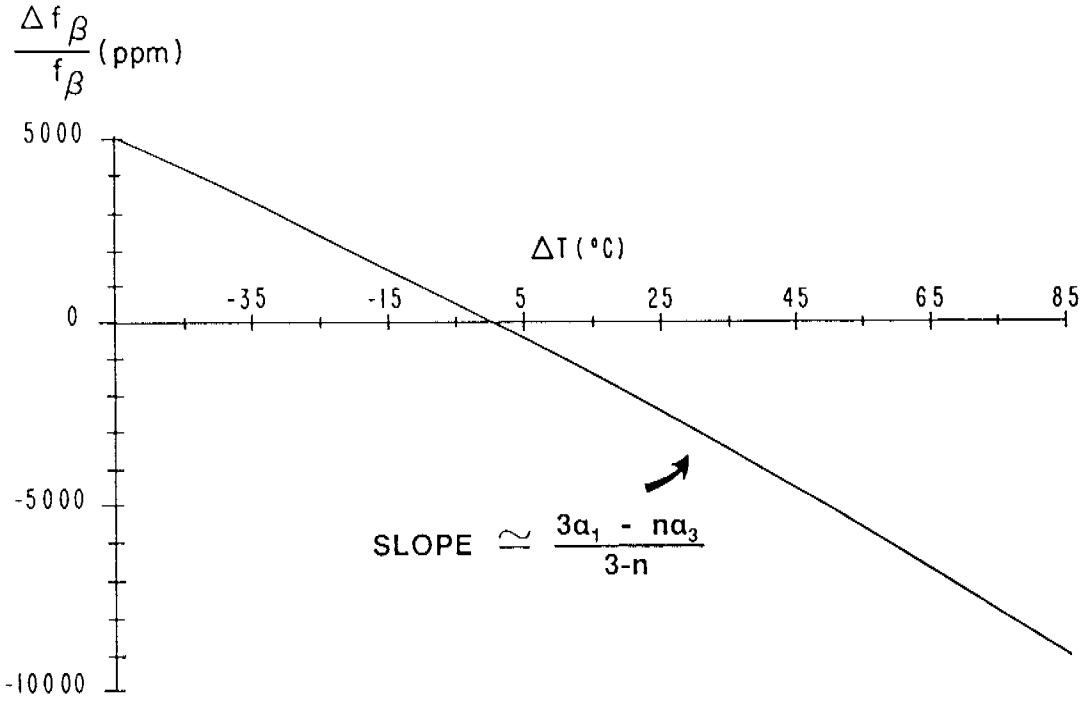

Figure 4. Temperature dependence of the beat frequency.

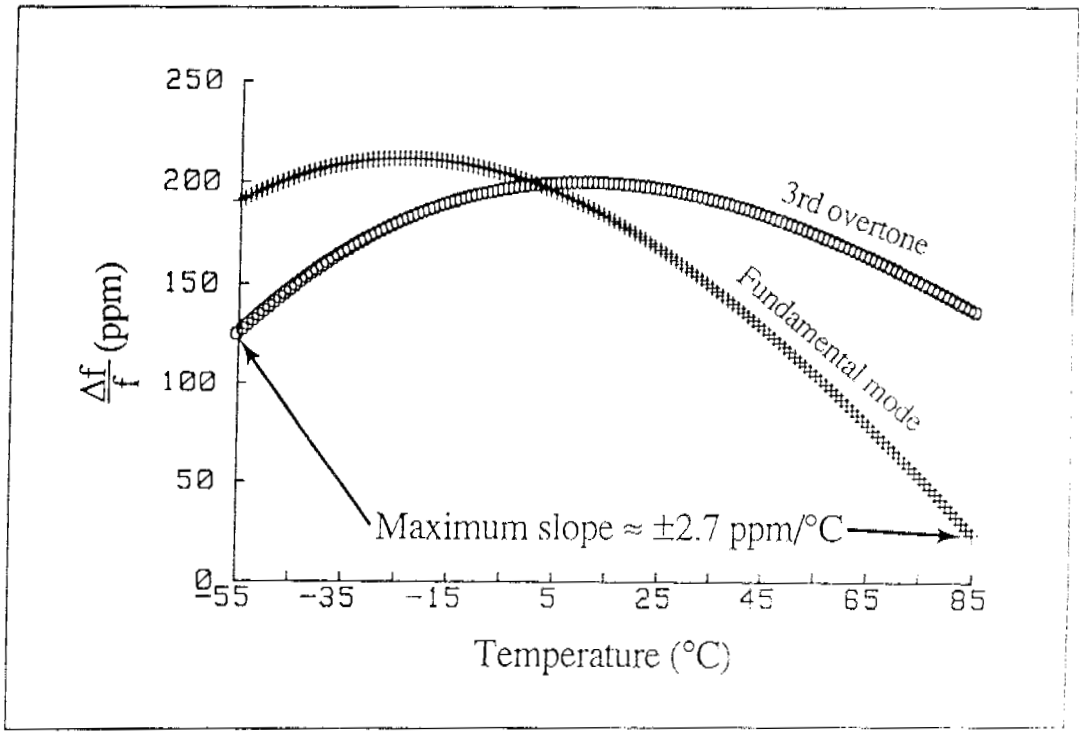

Figure 5. Optimum resonator f vs. T for -55°C to +85°C MCXO.

Table I
MCXO - TCXO Resonator Comparison

Parameter	MCXO	TCXO
Cut, overtone	SC-cut, 3rd	AT-cut, fund.
Angle-of-cut tolerance	Loose	Tight
Blank f and plating tolerance	Loose	Tight
Activity dip incidence	Low	Significant
Hysteresis (-55°C to +85°C)	10^{-9} to 10^{-8}	10^{-7} to 10^{-6}
Aging per year	10^{-8} to 10^{-7}	10^{-7} to 10^{-6}

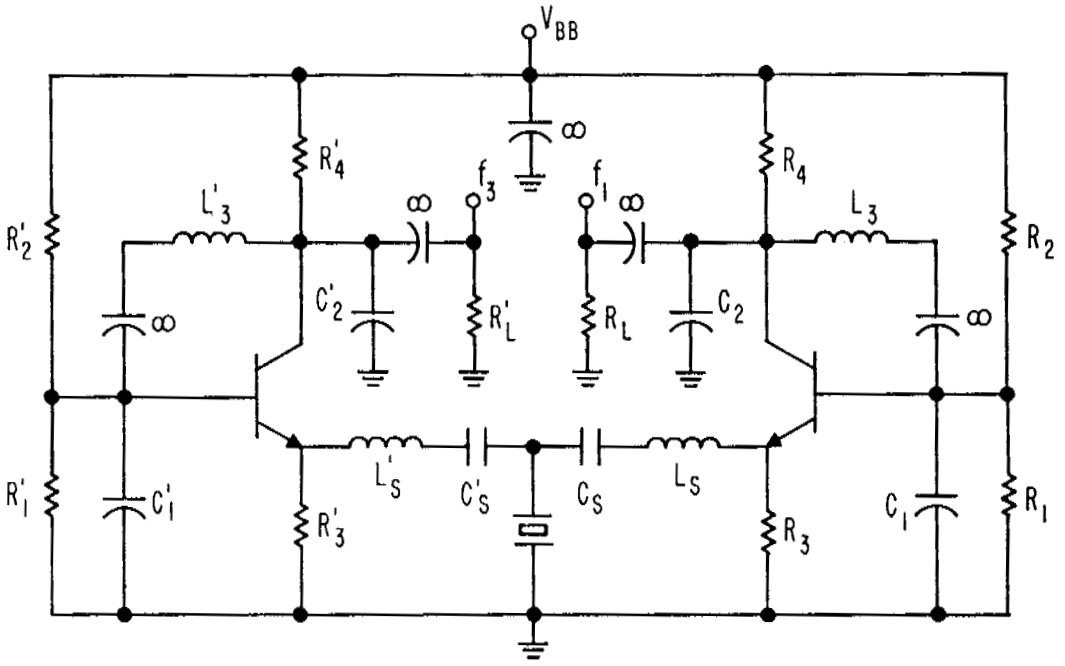

Figure 6. A double-gain-loop DHMXO (emitter degenerative type).

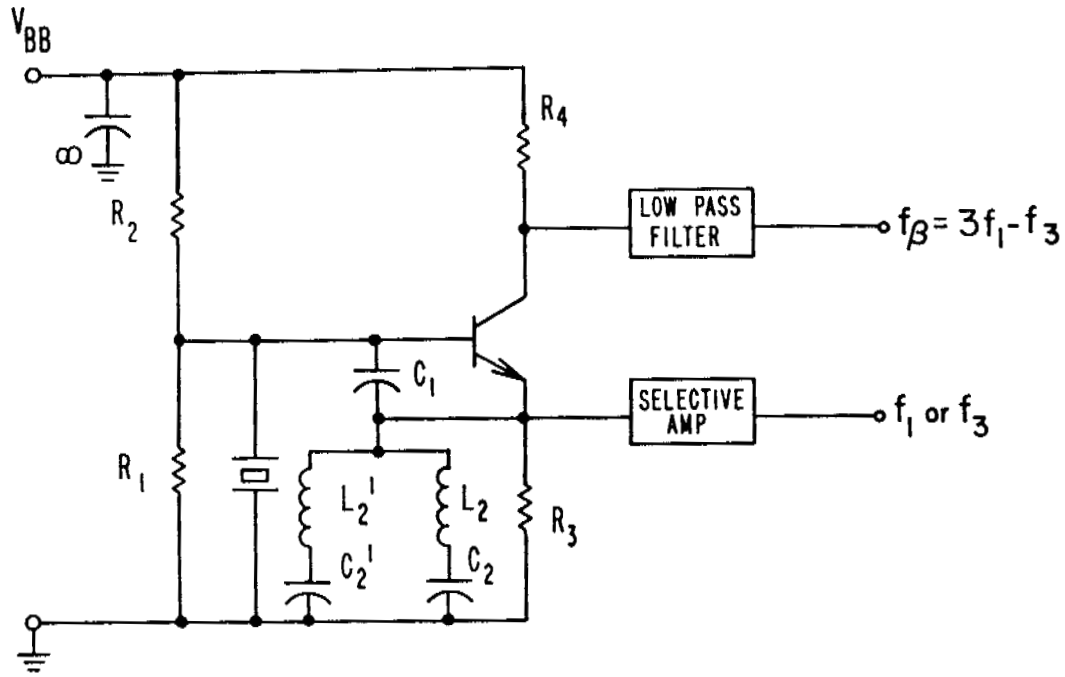

Figure 7. A single-gain-loop DHMXO (colpitts type).

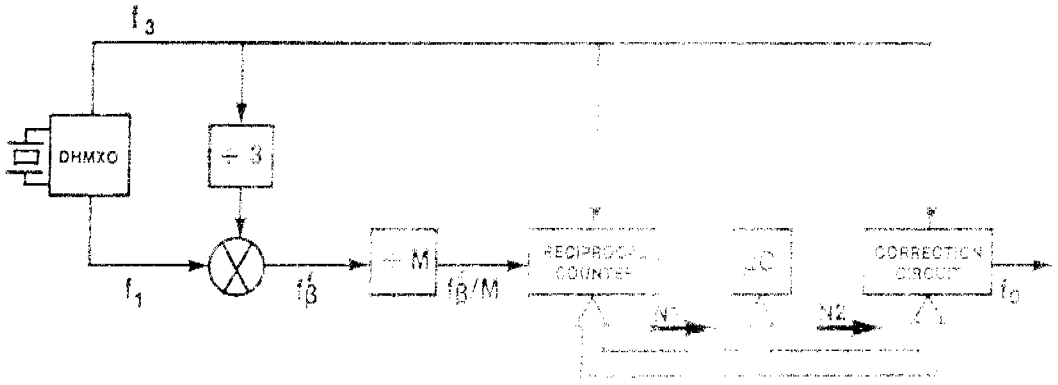

Figure 8. Typical implementation in an MCXO.

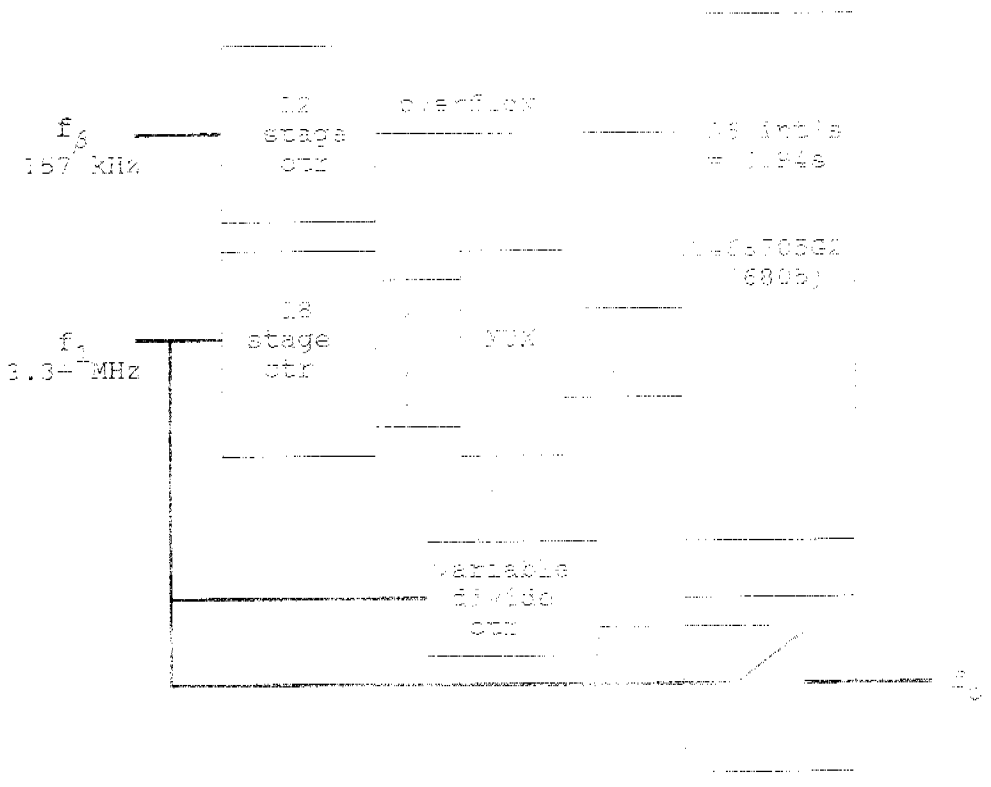

Figure 9. Digital control block diagram (pulse deletion).

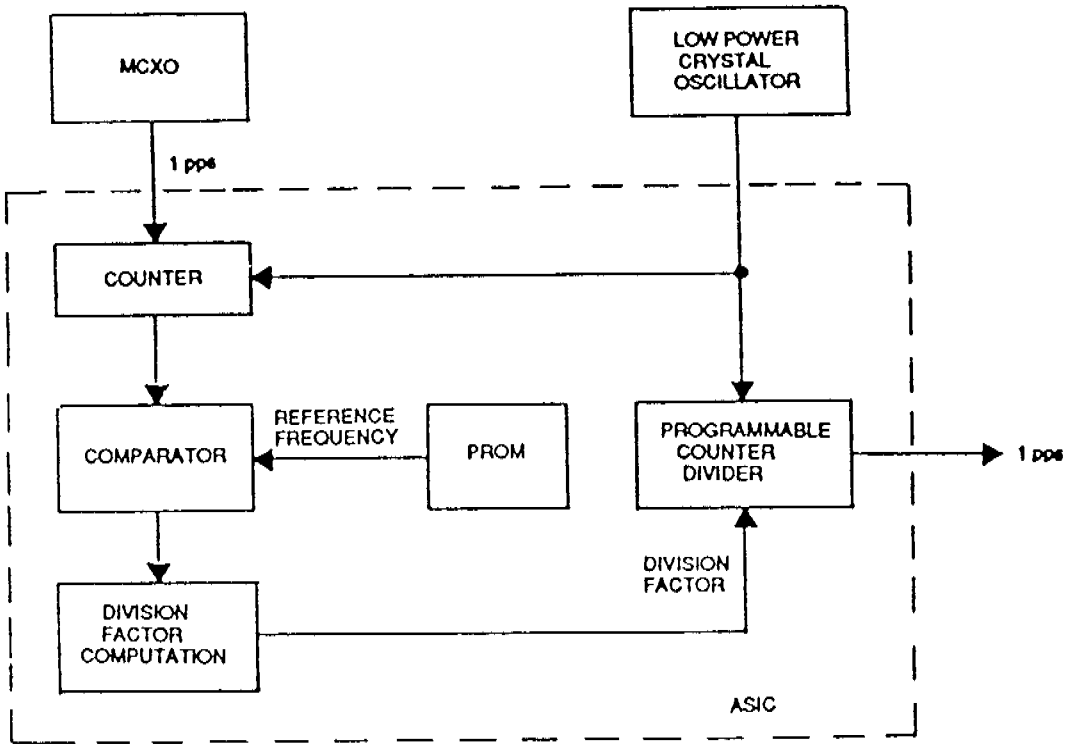

Figure 10. Low-power clock block diagram.

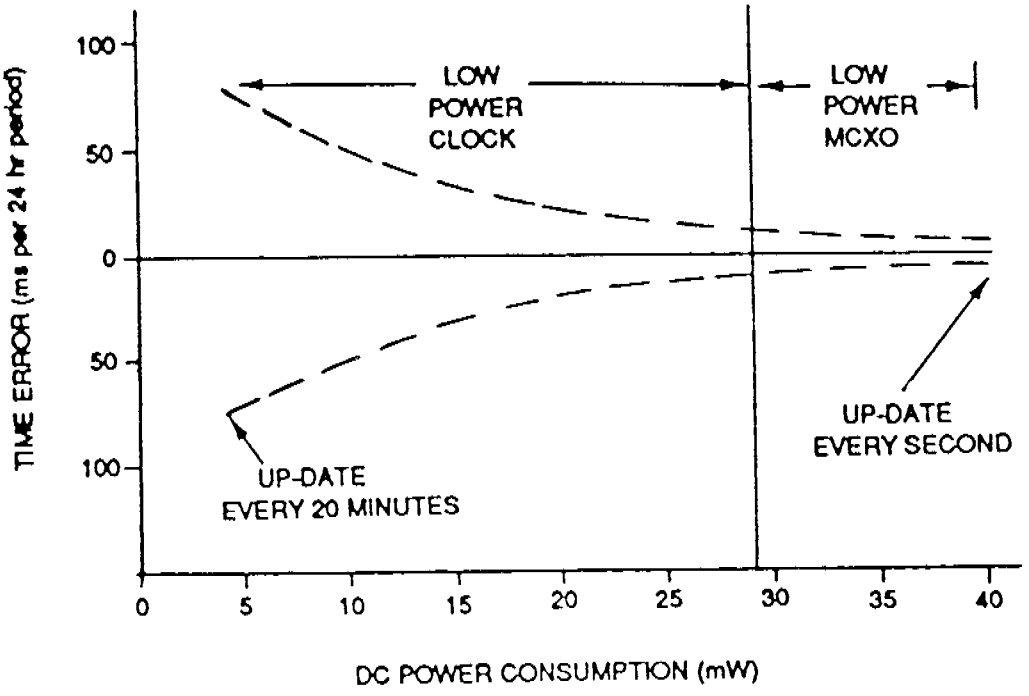

Figure 11. Timing error vs. power consumption.

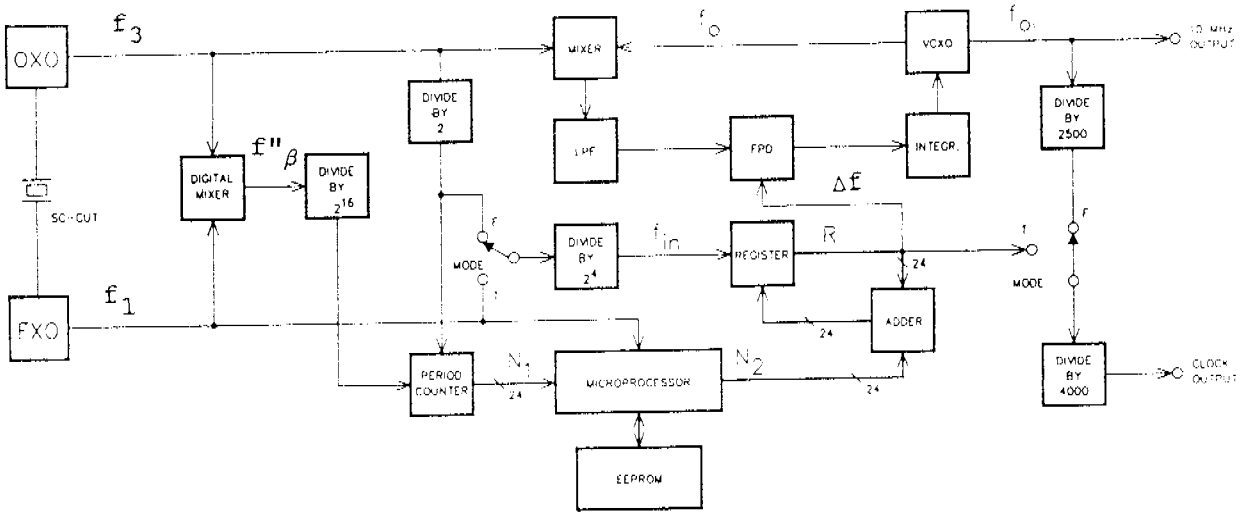

Figure 12. GTS MCXO block diagram.

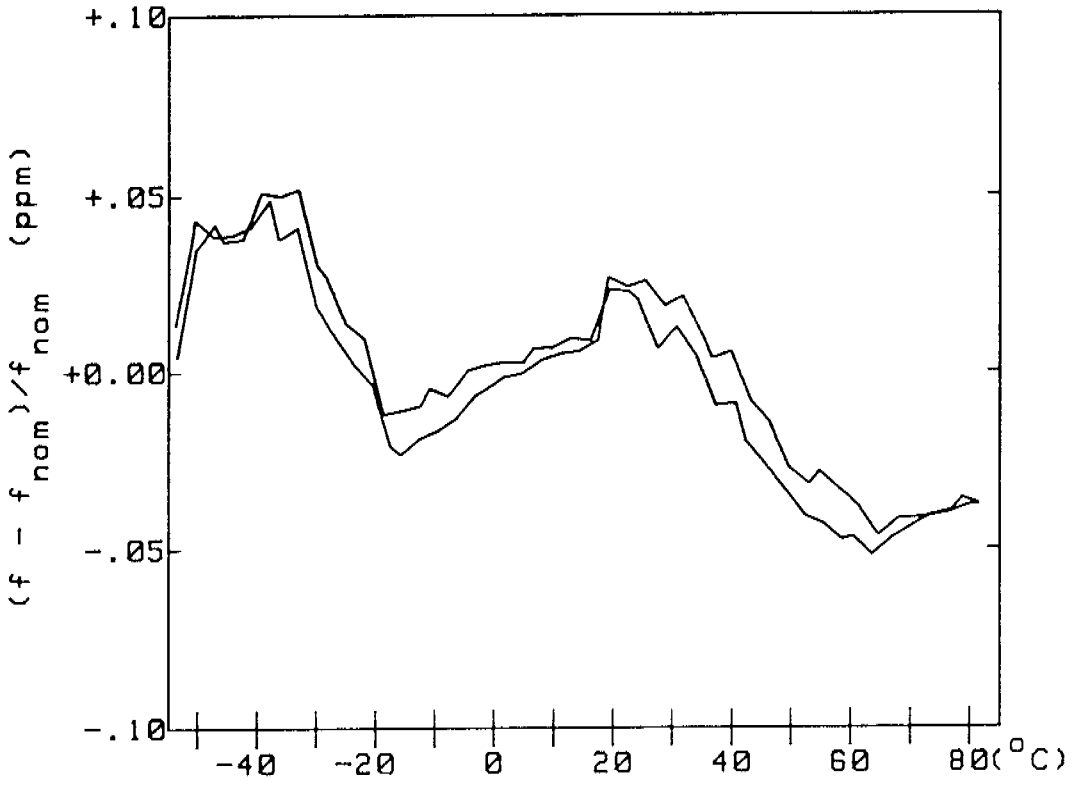

Figure 13. f vs. T characteristic of MCXO 16725.

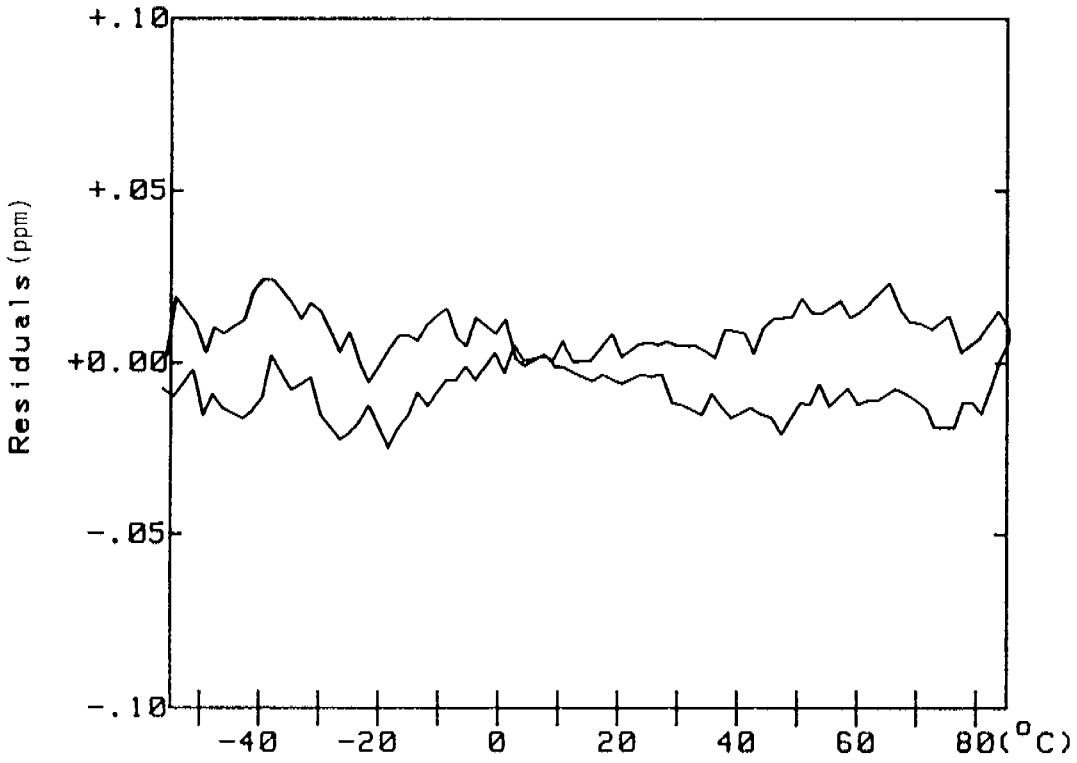

Figure 14. Residuals vs. temperature for the dual-mode oscillator in MCXO 16725.

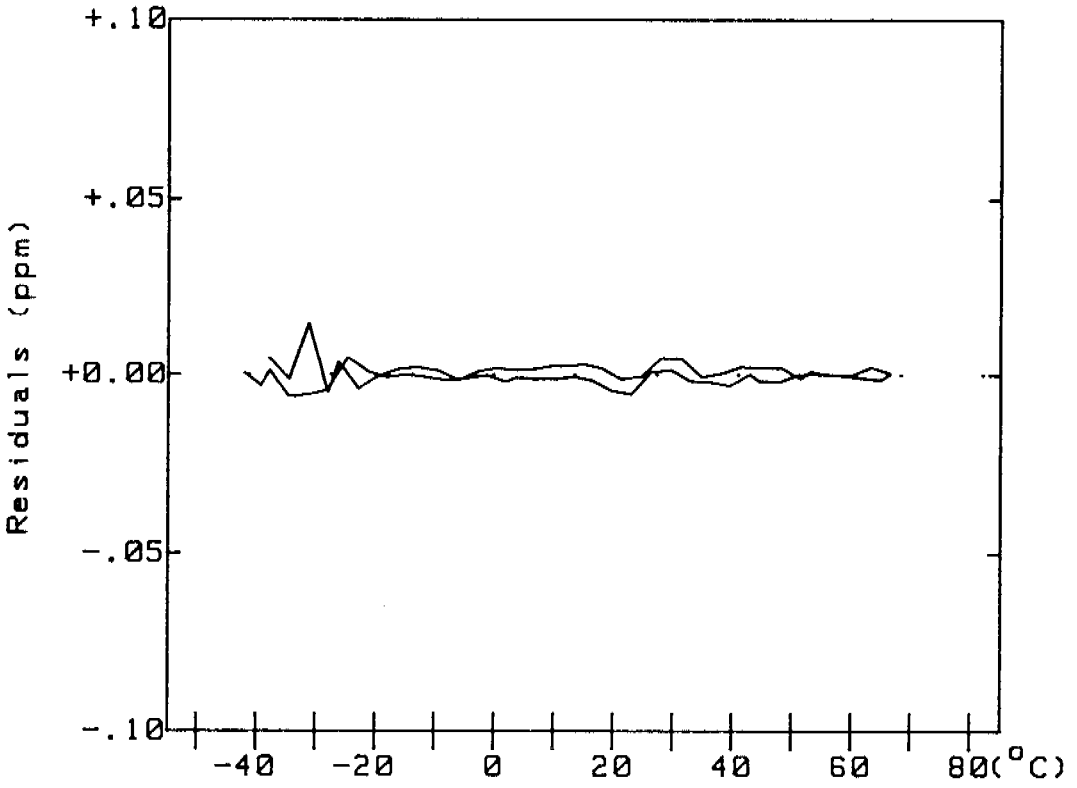

Figure 15. Residuals vs. temperature for the TFR dual-mode oscillator SN AP7.

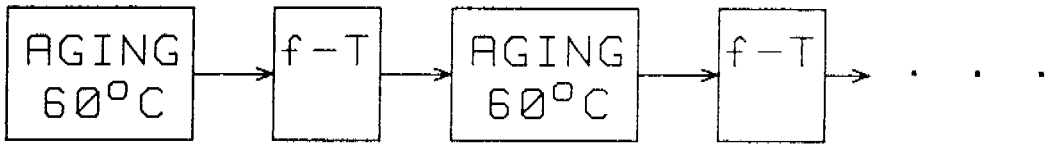

Figure 16. Aging sequence.

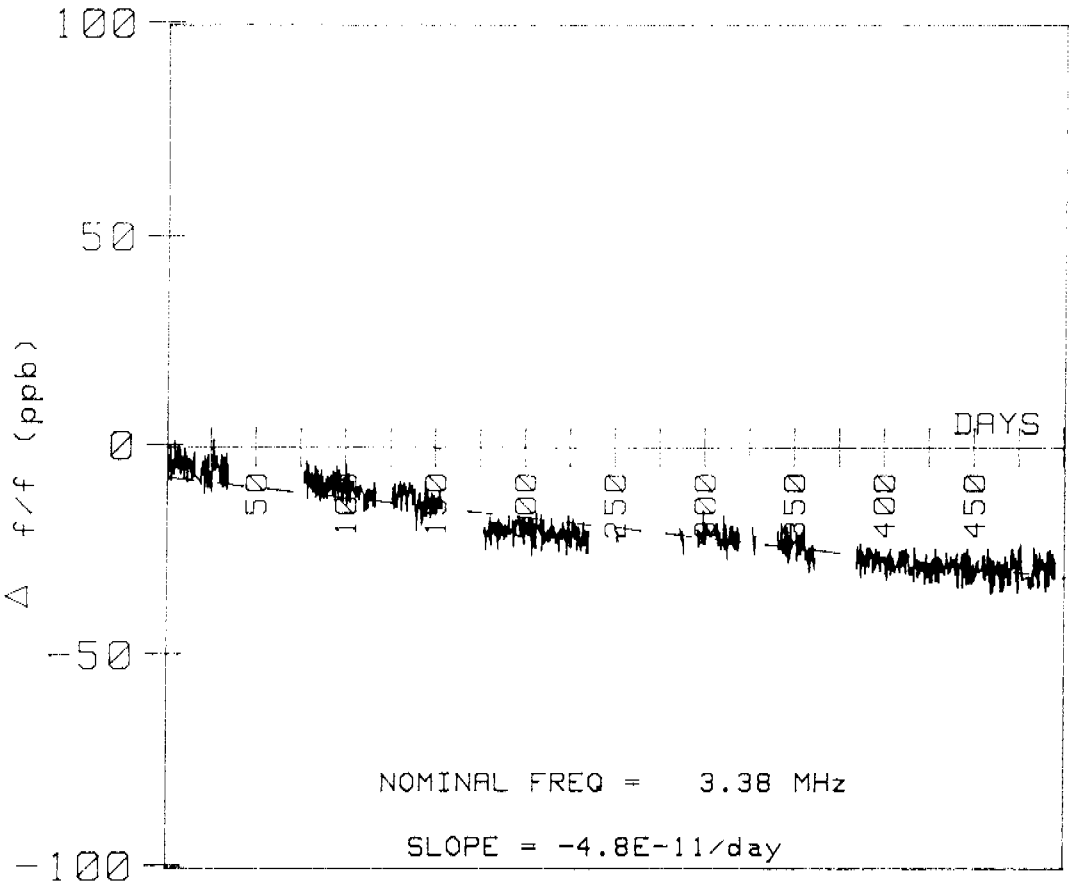

Figure 17. Aging data for MCXO 16725.

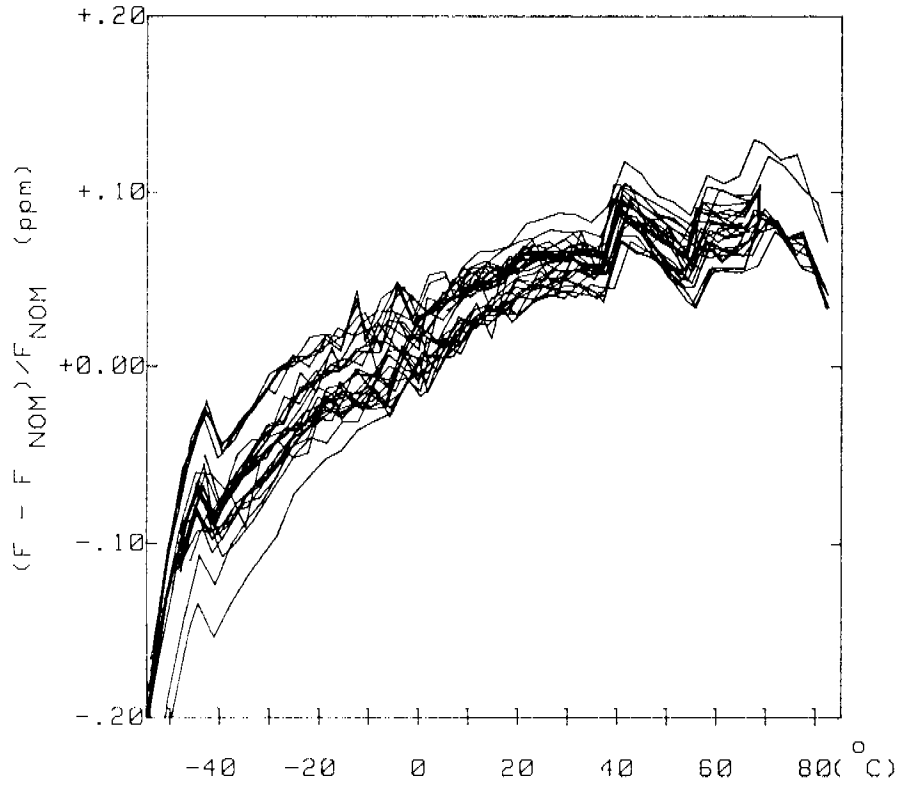

Figure 18. Superposition of 15 f vs. T runs with various parameters for MCXO 16723.

Table II

MCXO PERFORMANCE SUMMARY

Source	Res.	# of f-T	f-T (int) 10^{-8}	f-T (ext) 10^{-8}	Hyst 10^{-8}	Aging 10^{-10}/d
FEI	LFR	21	12.5	3.2	2	2.0
FEI		19	5.5	–	1	0.6
FEI		17	5.0	2.5	1	1.0
FEI		16	11.0	5.7	4.5	0.3
FEI		15	8.5	3.2	1.5	0.7
GTS	TFR	11	–	1.1	0.5	–
GTS		11	–	3.1	2	–
GTS		1	1.0	–	0.8	–

QUESTIONS AND ANSWERS

MR. REINHARDT: Why can't you just use a varactor and just feed back the information and control the frequency directly?

MR. VIG: In the MCXO a varactor wouldn't work because it would affect the two frequencies differently. The fundamental mode would be pulled much more than the third overtone. It would mess up the compensation. The idea is to avoid having to pull the frequency.

UNIDENTIFIED QUESTIONER: Do you have any data on the phase noise of these oscillators and do you see any increase in phase noise due to amplitude variations in the drive?

MR. VIG: These were initially intended for clocks, so phase noise has not been of concern, although we are finding more and more applications where people want both an accurate clock and a source for carrier frequency control. We are looking at various ways of getting both a good clock output and a good frequency output from the MCXO. The obvious way of doing it is to just phase lock a VCXO to it. That, in fact, is being done in both approaches. When you do that, you will get the phase noise of the VCXO, which will be basically signal–to–noise at the far out frequencies and whatever VCXO's do close to the carrier. Since the temperature is not controlled, the close in phase noise will probably be temperature dependent. The is no unique answer as to what the phase noise will be at one Hertz from the carrier.

MR. REINHARDT: If you let the VCXO run at a different frequency than the MCXO, the offset could be controlled by the information from the MCXO and it could all be done in software.

MR. VIG: Exactly. One application that we have right now will use a 10MHz clock and a 60 Mhz frequency reference, using that kind of approach.