1. Introduction

This paper begins with a review of conventional oscillator technology, with particular attention to temperature sensitivity. Next, the bridge stabilized oscillators of Meacham and Sulzer are discussed. This will establish why a new bridge stabilized circuit is needed, what it needs to do, and how it might work. Then the new bridge network and ancillary circuits are described. Within this context, the ALC techniques are developed. Some manufacturability issues and experimental results are presented. The oven and other peripheral hardware used with the oscillator are described in [1].

2. Existing technology

2.1. Conventional circuits

Virtually all non-bridge-stabilized crystal oscillators are derived from the free-running Colpitts oscillator. In the Pierce oscillator [2], a parallel resonant crystal, replaces the inductor. In the bridged-T oscillator (also known as grounded-base or Butler[1]), a series resonant crystal is inserted in series between the emitter and the rest of the tank circuit (fig. 1). In either case, the oscillator consists of an amplifier coupled to a feedback network that is essentially a crystal filter. The crystal filter's phase vs. frequency slope and the amplifier's phase vs. temperature slope in combination determine the oscillator tempco with respect to the amplifier. The temperature coefficients of the passive components make additional contributions.

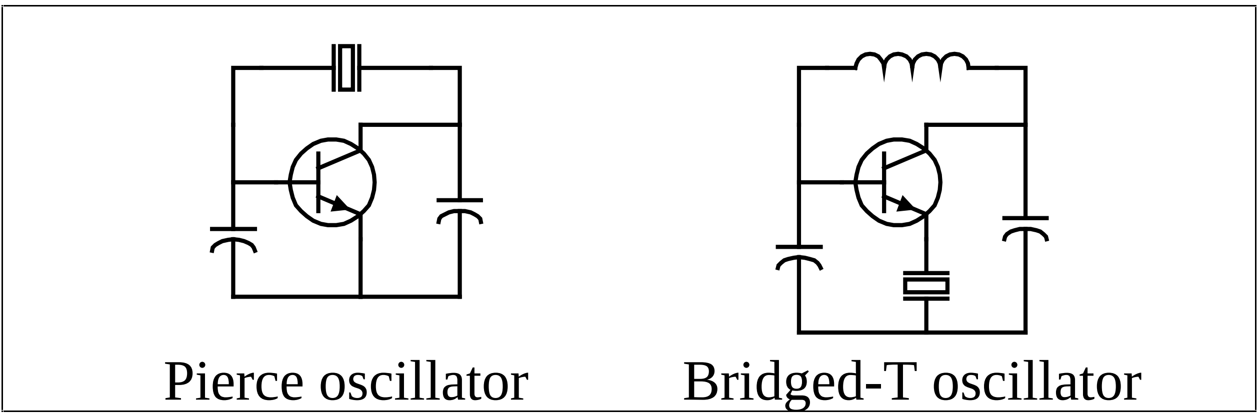

Figure 1. Basic oscillator circuits.

To minimize the oscillator tempco, low tempco passives can be used, and the phase vs. frequency slope of the crystal filter network can be maximized by maximizing loaded Q. There are limits to what can be done in these areas. ±30 ppm/°C ("NP0") capacitors are available, but inductor tempco generally runs at least a few hundred ppm/°C positive. The loaded Q of the crystal network can only approach the unloaded Q of the crystal, which is in turn limited by the intrinsic QF product of quartz. If an SC cut crystal is used, a additional problem is created by the need for a

[1] Butler's "earthed grid" oscillator was actually based on the Hartley, not the Colpitts oscillator [3,5] .

high Q mode suppression tank. This tank must have a loaded Q on the order of 10 because the frequencies of mode B and mode C are only about 10% apart, and the B mode has somewhat lower resistance. Because of the leverage created by the high Q, the sensitivity of frequency to the value of the mode suppressor components is stepped up an order of magnitude compared to a design intended for AT cut crystals.

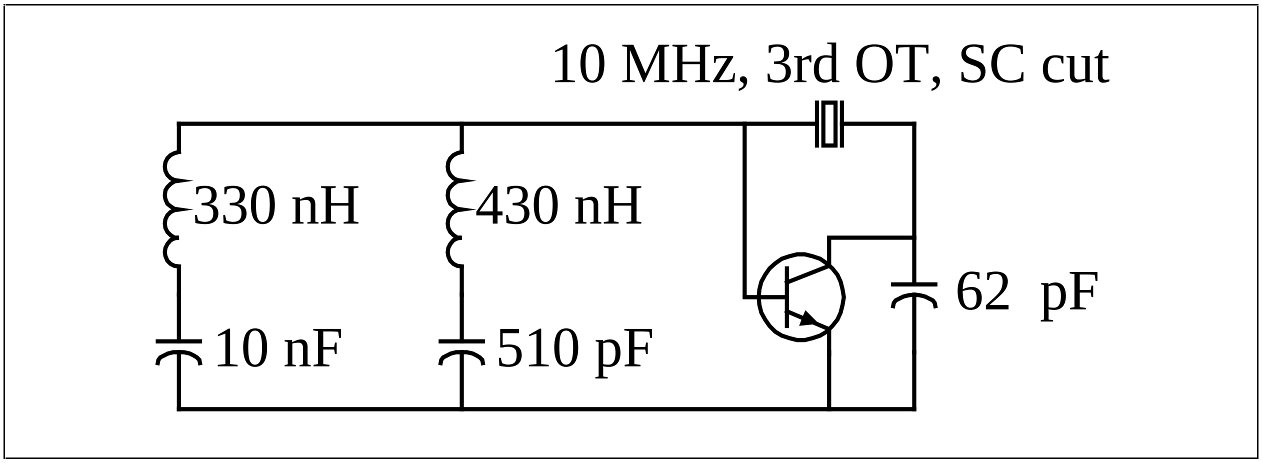

Figure 2. Pierce oscillator with mode suppressor.

Fig 2 shows a simplified schematic of a typical 10 MHz third overtone SC cut crystal oscillator [6]. The active device is the commonly used 2N5179. It is derived from the Pierce oscillator of fig. 1, with the base-emitter capacitor replaced by the double LC mode suppressor circuit shown. The tempco contribution of the passive components can be calculated using network analysis based on their specified performance. The 2N5179 is not specified for phase and gain vs. temperature, so these parameters were measured experimentally with the circuit of fig. 3.

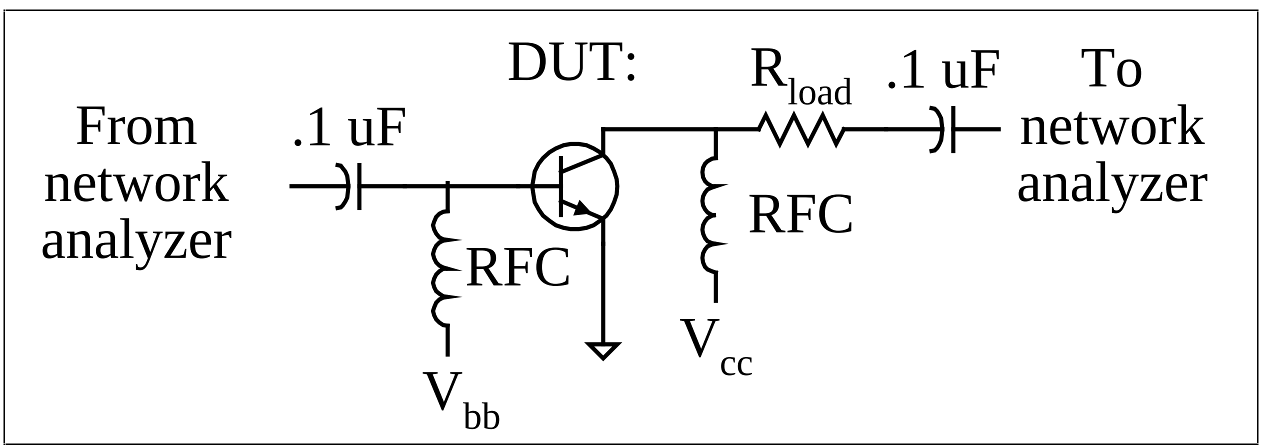

Figure 3. Transistor tempco test fixture.

The effect of phase on the frequency of oscillation is obvious, but the effect of gain is indirect. As the transistor temperature changes, the gain (for a given collector current) changes. This gain error must be corrected by the ALC circuit, which typically operates by varying the bias of the transistor. The phase coefficient of bias current then results in an AM to PM conversion process.

Analysis of the circuit of fig. 2 shows at least six tempco effects on the order of 10^{-10} / ºC each:

1. Transistor phase shift.
2. Transistor gain shift.
3. 62 pF "NP0" cap (±30 ppm / ºC).
4. Inductance tempco (+250 ppm / ºC).
5. Inductor Q tempco ($\Delta Q \rightarrow \Delta L_p$).
6. ALC tempco, assuming 0.02 dB / ºC.

None of these are particularly predictable, they tend to interact, and measurements on real oscillators show these numbers to be optimistic.

Experimental measurements of three complete oscillators of different designs were made for additional confirmation. The three designs were a conventional Colpitts oscillator, a bridged-T oscillator, and a two-transistor Butler oscillator.[2] The Butler oscillator used a transistor array IC instead of 2N5179's and used a saturated amplifier in place of a conventional ALC scheme. All three oscillators showed a tempco with respect to the active device(s) on the order of 10^{-9}/°C. From these results, it was concluded that a breakthrough in tempco reduction of any oscillator of this general class was highly unlikely.

2.2. The Meacham bridge oscillator

The earliest known attempt to eliminate active device tempco in crystal oscillators was the bridge-stabilized circuit described by Meacham in 1938 [4, 7] . In the Meacham bridge oscillator, the crystal replaces one of the resistors in a Wheatstone bridge (fig. 4). The bridge is connected so as to provide negative feedback to the active device at all frequencies except a narrow band around the crystal's series resonant frequency, where the feedback becomes positive. Note that the bridge must be slightly out of balance at the frequency of oscillation; if it were balanced there would be no feedback to sustain oscillation. The idea is that the frequency of oscillation will be determined by where the bridge balances (or more properly is maximally unbalanced in the positive direction), with minimal impact by the phase shift through the active device.

Meacham showed that indeed the circuit becomes less sensitive to perturbations outside the bridge as the gain of the active device increases, which causes the bridge to operate closer to being balanced. He described this circuit as a sort of Q-multiplier, which makes the Q of the crystal appear to increase by a factor equal to the amount of excess gain in the active device. The increase in effective Q makes the crystal harder to pull off frequency.

[2] Not the oscillator described in [5], rather Butler's previous "two-valve" design.

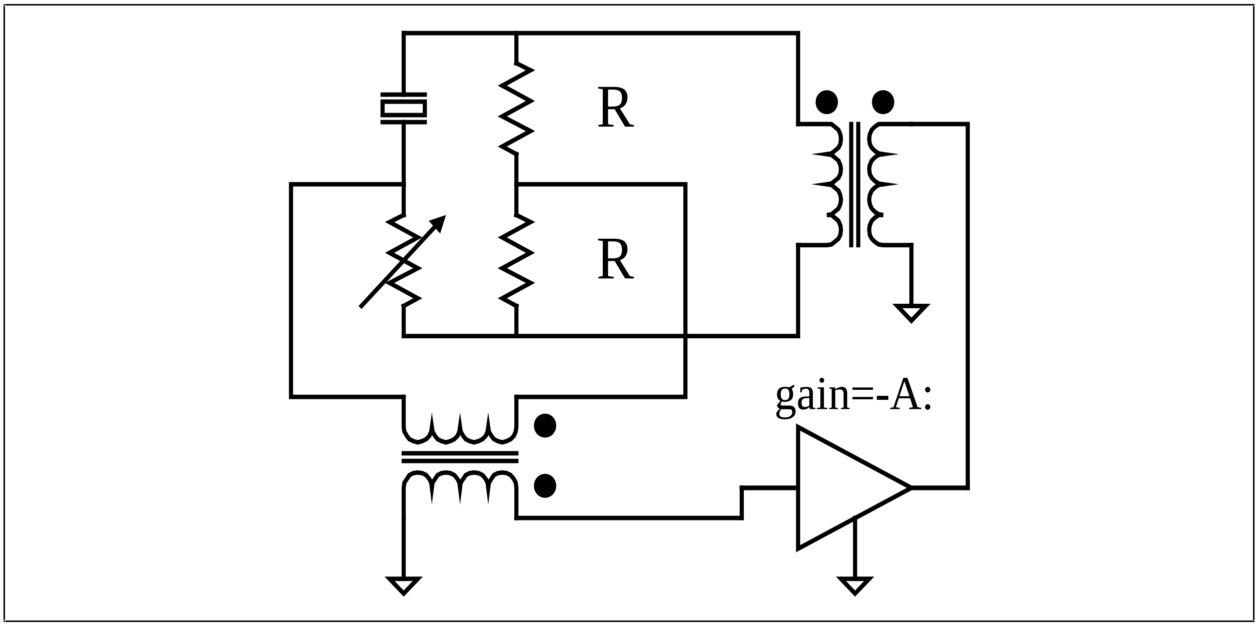

Figure 4. Meacham bridge-stabilized oscillator.

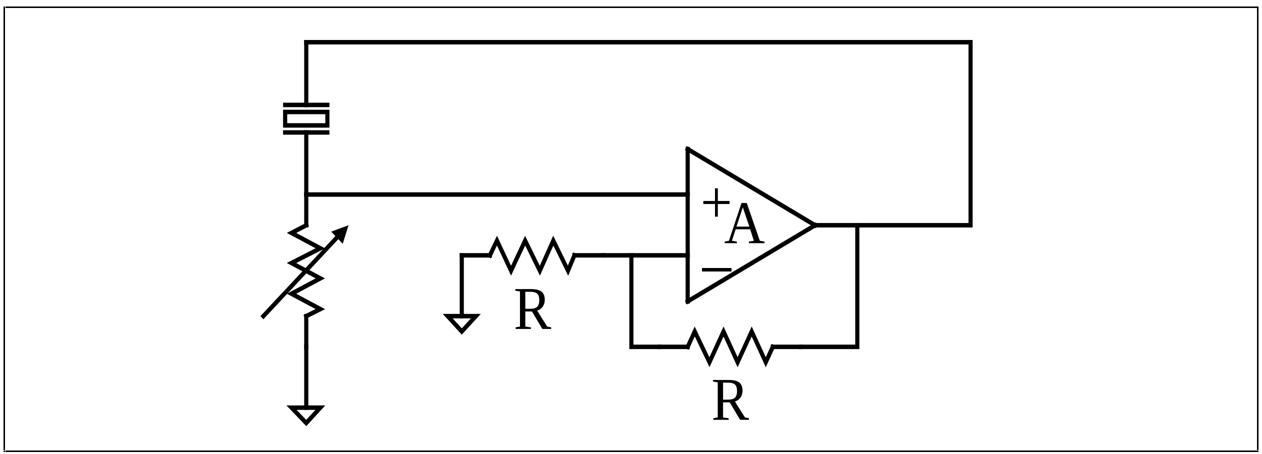

Figure 5. Meacham oscillator utilizing op amp.

Fig. 5 shows another way of looking at Meacham's circuit [8]. Here it can be seen that what is really going on is that negative feedback has been applied to the active device as if it were an op amp. And this is the key to understanding the limitations of the Meacham bridge oscillator. For any frequency not in the vicinity of the crystal resonant frequency, the crystal can be considered an open circuit, and the conditions for stability of the active device are the same as for an op amp with a closed loop gain of -1. This requires that the active device be equivalent to an op amp that is unity-gain stable. It is extremely difficult to realize a unity-gain op amp using conventional dominant pole compensation with a gain bandwidth product of more than about 300 MHz. This limits the active device gain at 10 MHz to about 30 dB. Since 6 dB of this gain is necessary to overcome the voltage drop through the crystal, only 24 dB of negative feedback is being applied.

This is not a sufficient amount of feedback to remove the dependence on the phase shift of the active device.

One way to overcome this limitation is to change the compensation from a dominant low pass pole to a dominant band pass complex pole pair. This could be implemented, for example, by shunting the input with a parallel LC tank. However, the tempco of the tank components creates a new source of drift. This dilemma is a general problem that would apparently occur with any scheme involving the use of negative feedback to stabilize the active device. The arguments presented here are similar to the reasons why active filters are never used at 10 MHz, at least if a high degree of precision and stability is required of them.

An additional problem with the Meacham bridge oscillator is that the mode suppressor required for SC cut crystals would have to be inside the bridge, hence it would have just as high a tempco as it would in a non-bridge circuit.

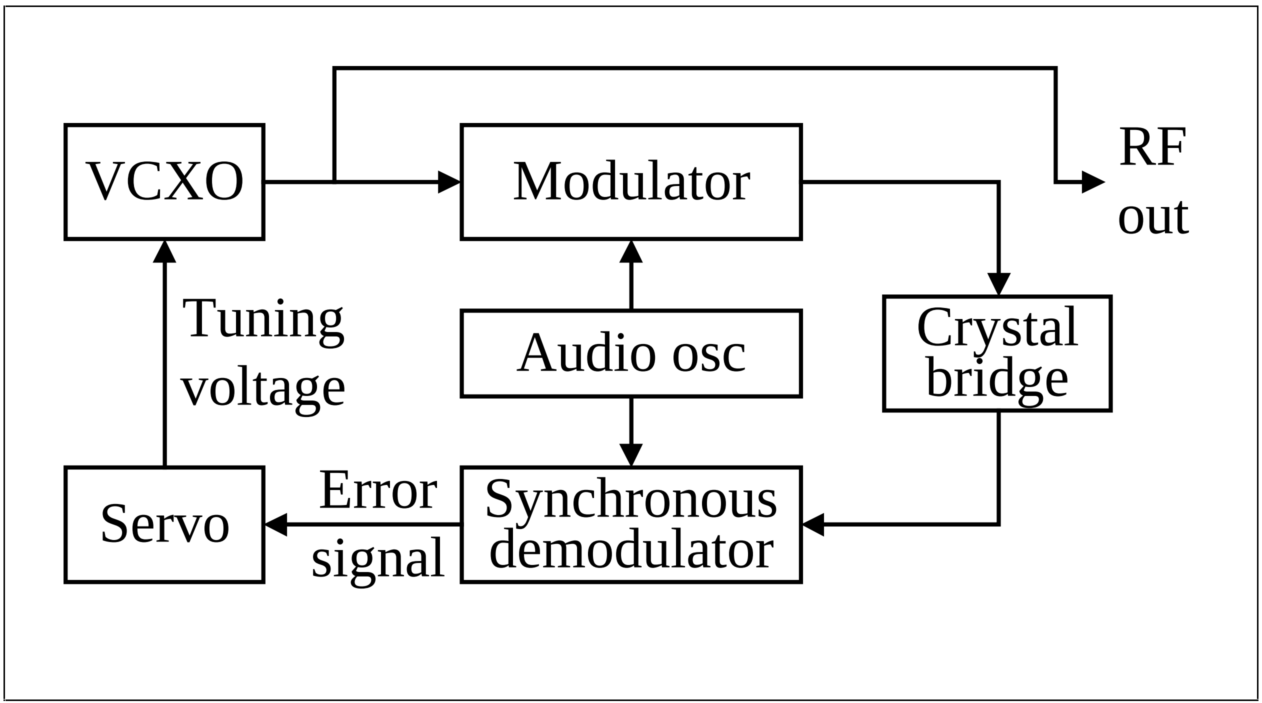

Figure 6. Bridge servo'ed oscillator.

2.3. Bridge servo'ed oscillators

Crystal oscillators stabilized by an auxiliary crystal in a balanced bridge were constructed in the early 1950's [9]. A servo loop tuned the oscillator to keep the bridge balanced, which would force the frequency of oscillator to coincide with the unperturbed resonant frequency of the crystal (fig. 6). The bridge was the same modified Wheatstone bridge used in the Meacham oscillator, and the servo loop was a conventional lock-in amplifier operating via synchronously detected audio modulation, as commonly used in atomic standards. The advantage of this scheme is that the bridge operates with arbitrarily small deviation from balance, since there is no limit to the servo loop gain. The obvious disadvantage was the requirement for two crystals.

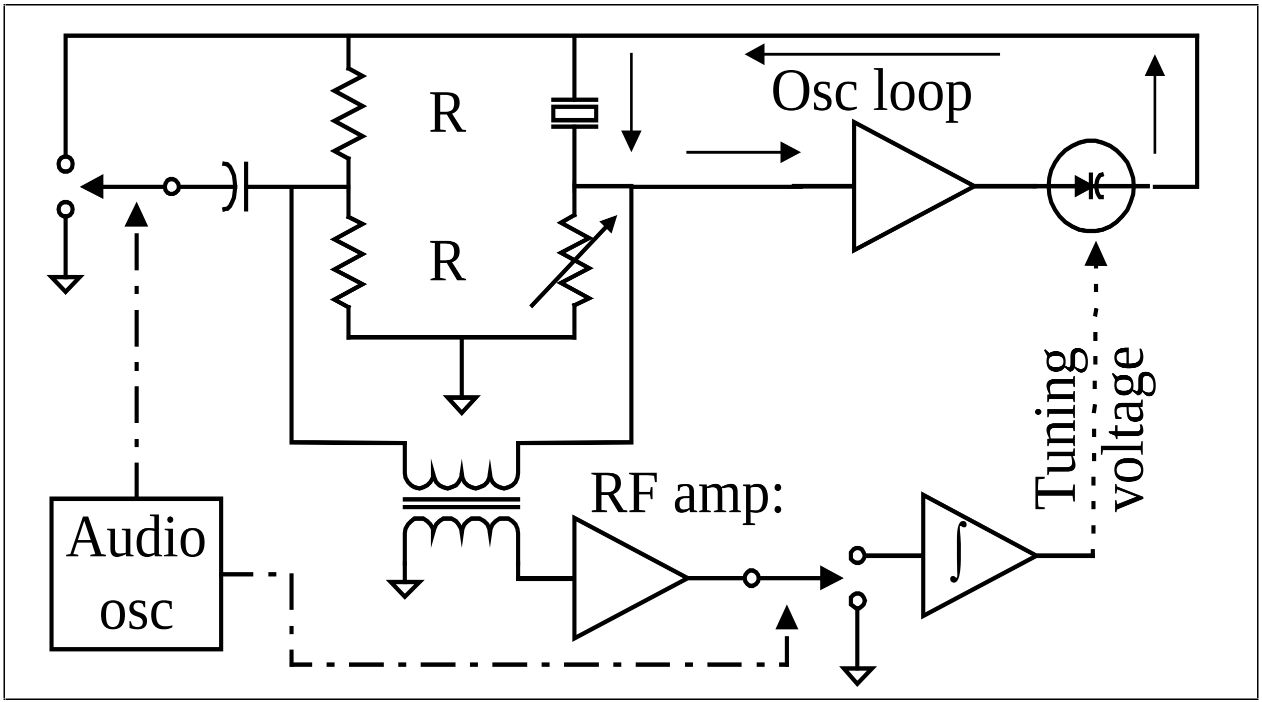

Figure 7. Sulzer bridge-balancing oscillator.

2.4. Sulzer's "bridge-balancing" oscillator

In 1955, Sulzer [10] described an oscillator and bridge using a single crystal that servo'ed the frequency so as to maintain bridge balance (fig. 7). The basic bridge is again the modified Wheatstone bridge. Instead of driving the bridge with a separate VCXO, the crystal in the bridge is made part of an oscillator circuit. The output of the amplifier drives the top of the bridge as with the VCXO in the figure 6. The input of this amplifier is driven by feedback from the bridge.

The proper output port of the bridge cannot be used because it has a null, not a peak, at resonance. Instead, the amplifier input is connected in shunt with one of the bridge resistors. A lock-in amplifier is used to servo the bridge into balance, but the audio modulation is cleverly added by modulating the impedance of the bridge, instead of by adding undesirable sidebands to the oscillator signal. This was done by an electromechanical chopper relay.The basic problem with this design is that the integrity of the bridge is violated due to the shunt loading of the arms by the oscillator active device and the modulator. The active device must present a high impedance load to the bridge, which compromises phase noise. In a modern design, the mechanical chopper would have to be replaced by some sort of electronic equivalent, such as FET switches, which would add their own capacitive loading across the bridge. Finally, there would inevitably be crosstalk between the bridge modulator and oscillator sections that would impress audio sidebands on the output signal at some level. Although Sulzer was able to demonstrate greatly reduced sensitivity of the frequency to such perturbations as power supply voltage, this design clearly left a lot to be desired.

3. The new oscillator

3.1. The ideal bridge-stabilized oscillator

The previously described oscillators each have unique advantages and disadvantages. What is needed is an ideal bridge stabilized oscillator that combines all the advantages with none of the disadvantages by meeting the following requirements:

1. One crystal is used for both the bridge and the oscillator.
2. The bridge operates at balance and controls the frequency of oscillation.
3. The integrity of the bridge is maintained.
4. No audio modulation is added.
5. Low phase noise is facilitated.
6. Overtone and mode control can be effected outside the bridge and without compromising stability.

3.2. A true balanced-bridge controlled oscillator

The approach chosen to meet these requirements is shown in fig. 8. The crystal is removed from a VCXO and embedded in a two port bridge network. Meanwhile, port 1 of the network is connected to the VCXO in place of the crystal. A null detector operates an AFC (automatic frequency control) loop. The key is to devise a two-port bridge network that meets two requirements: (1) The transfer function from one port to the other has a null output when the crystal is excited at resonance (as in the modified Wheatstone bridge) and (2) port 1, the input, has a driving point impedance such that it emulates the impedance of a freestanding crystal (i.e. the bridge is "transparent").

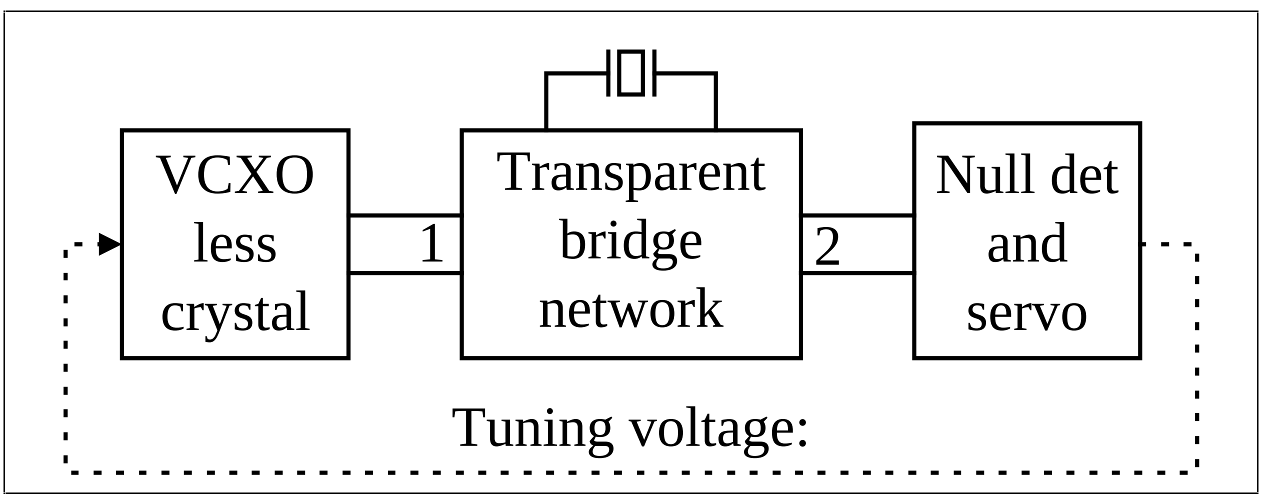

Figure 8. Balanced bridge controlled oscillator.

Transparency, as used here, means that if the oscillator were connected to port 1 of the network, it would operate just as if it were connected directly to a crystal. In other words, port 1 emulates a virtual crystal. The modified Wheatstone bridge originated by Meacham is not transparent because of the shunting effect of the resistors on the non-crystal side of the bridge

3.3. Development of the bridge network

Although the bridge network does not require any particular oscillator design for its proper operation, some initial assumptions will be made about the oscillator to facilitate the analysis. The oscillator will be assumed to be of the series resonant type, and will be assumed to have some sort of ALC system that maintains constant (virtual) crystal current. Thus, when port 1 of the bridge network is substituted for the crystal, the oscillator will then maintain constant current into port 1. Also, the holder capacitance (C_0) of the crystal will be assumed to be negligible at this time. These assumptions will be revisited after the basic bridge network is determined. An additional assumption is that the output port can be considered to be unloaded because the input admittance of the null detector is negligible.

The network will be analyzed with z parameters. Since $i_2 = 0$, the driving point impedance at port 1 is the same as z_{11} by definition.[3] The desired driving point impedance of the bridge network is of the same form as a crystal (without C_0), hence $z_{11} = [(Q_v/\omega_0)s^2+s+Q_v\omega_0](R_v/s)$, representing a virtual crystal[4] with motional resistance $R_m = R_v$, quality factor $Q_m = Q_v$, and resonant frequency of ω_0. The desired form of v_2 vs frequency is a null at resonance. Since i_1 is constant (due to oscillator ALC), v_2/i_1 should also be of this form. And because $i_2 = 0$, $z_{21} = v_2/i_1$. A null at resonance can be implemented simply by a complex pair of zeros on the jω axis at $\pm j\omega_0$ (complemented by poles at 0 and ∞ for realizability.) Hence, $z_{21} = (s^2+\omega_0^2)(K/s)$, where K is a gain constant to be chosen later. The bridge will not utilize non-reciprocal elements, therefore $z_{12} = z_{21}$. The remaining z-parameter, z_{22}, is not constrained, so it is advantageous to set $z_{22} = z_{11}$ so that the network is of the more easily realized symmetrical class.

Network theory predicts that if a symmetrical two-port is realizable in any form, it is realizable as a balanced lattice network. Furthermore, a balanced lattice network is actually a form of bridge circuit. Hence if a lattice solution can be found, a bridge solution has been found. For a balanced lattice network with known z-parameters, the impedances of the straight through arms, z_a, and the crossover arms, z_b are given by standard formulas [11] as follows:

$$\begin{aligned} z_a = z_{11}+z_{21} &= [(Q_v/\omega_0)s^2+s+Q_v\omega_o](R_v/s) + (s^2+\omega_0^2)(K/s) \\ &= [(Q_v/\omega_0+K/R_v)s^2+s+(Q_v/\omega_o+K/R_v)\omega_0^2](R_v/s) \end{aligned}$$

$$\begin{aligned} z_b = z_{11}-z_{21} &= [(Q_v/\omega_0)s^2+s+Q_v\omega_o](R_v/s) - (s^2+\omega_0^2)(K/s) \\ &= [(Q_v/\omega_0-K/R_v)s^2+s+(Q_v/\omega_o-K/R_v)\omega_0^2](R_v/s) \end{aligned}$$

---

[3] The z_{11} parameter is defined as the driving point impedance of port 1 with port 2 open.

[4] Denoted by the v subscript.

By setting arbitrary gain constant $K = R_v Q_v/\omega_0$, the crossover arms are reduced to $z_b = R_v$. The straight through arms become $z_a = [(2Q_v/\omega_0)s^2+s+2Q_v\omega_0](R_v/s)$, which, like z_{11}, represents the impedance of a crystal resonant at ω_o with $R_m = R_v$, but with a doubling of $Q_m = 2Q_v$, which also implies a doubling of L_m and a halving of C_m. The resulting lattice is shown in fig. 9. The network will appear to the oscillator as a crystal with half the q of the actual crystals. This minor deviation from the original transparency requirement is tolerable and is not an obstacle to the ultimate objective.

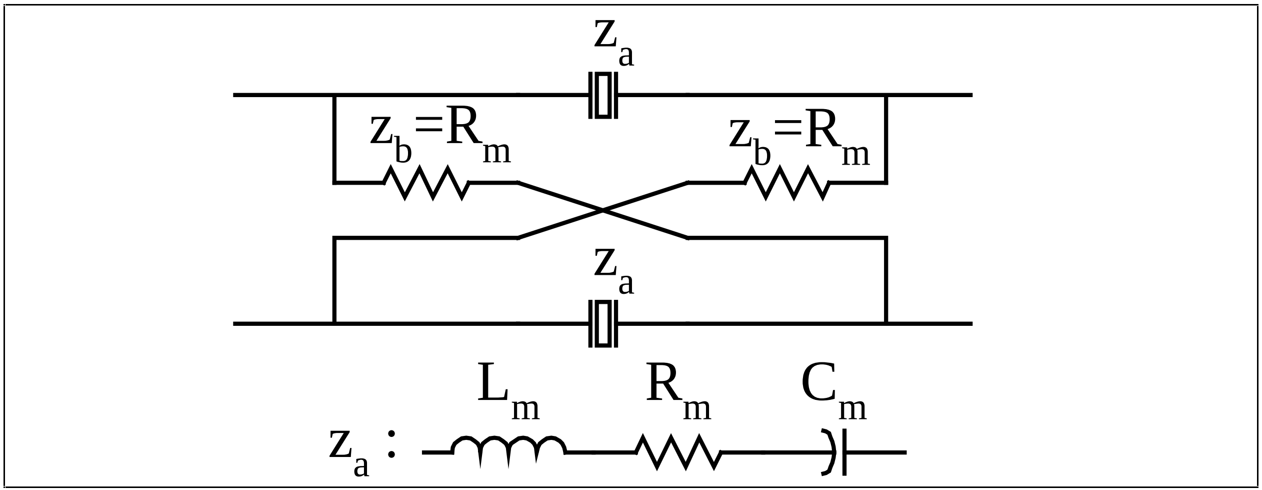

Figure 9. Theoretical transparent crystal bridge.

Having chosen $K = R_v Q_v/\omega_0$, it follows that $z_{21} = (s^2+\omega_0^2)(R_v Q_v/\omega_0 s)$ which results in a sensitivity on a fractional frequency basis at the null detector of $|z_{21}| = 2R_v Q_v \omega_o(\Delta\omega_0/\omega_0) = (1/C_m)(\Delta\omega_0/\omega_0)$, where C_m is the motional capacitance of the crystal.

A matched pair of crystals (or perhaps a monolithic dual resonator) would be required to build this lattice. This problem of component duplication comes up frequently in crystal filter design [12] and can be solved in this case by using the standard filter technique of converting the balanced lattice to a half lattice consisting of the crystal, a resistor and a transformer (fig. 10)[5]. The half lattice also has the advantage of having ports that can be used in either balanced or unbalanced mode. The resistor, matched to R_m, will be referred to as the image resistor.

The conversion to a half lattice has the side effect of transforming the impedance level at port 1 down by a factor of 4. The impedance level at port 2 is unchanged if the turns ratio of the transformer is 1:1, although any ratio can be used, subject to constraints caused by implementation details. Because of the symmetry of the network, the ports are interchangeable allowing port 2 to be connected to the oscillator with port 1 connected to the null detector. The apparent impedance of the crystal can then be transformed up or down to suit the oscillator. If desired, the unbalanc-

[5] Patented [15].

ing effects of the crystal holder capacitance, C_0, can be compensated out by adding an equal shunt capacitance across the image resistor.

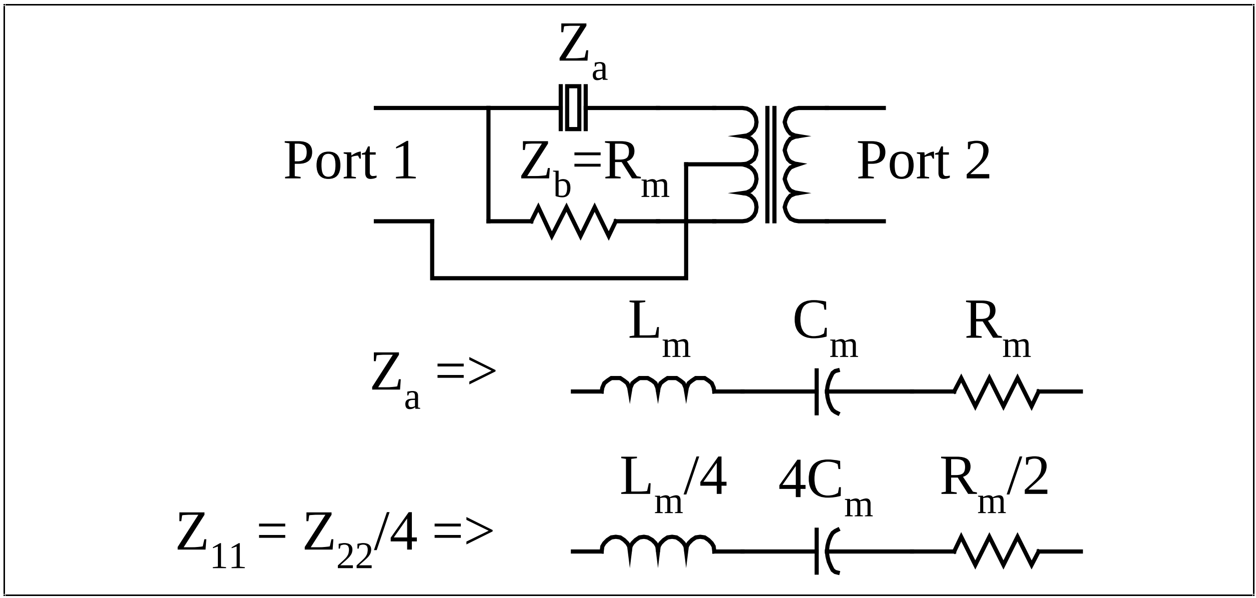

Figure 10. Practical transparent crystal bridge.

3.4. Development of a null detector

The first requirement for the null detector is that it have a high input-impedance front end to avoid drawing significant current from port 2. In practice, it is sufficient to have an input impedance about an order of magnitude higher than the magnetizing inductance/impedance of the tranformer, which will be a few thousand ohms at most. If the null detector simply consisted of a buffer amplifier driving a scalar power detector, it could detect the presence or absence of a null, but not the sign of the frequency error. That would require using audio modulation to sweep across the null as in previous schemes since sign information is essential to any servo loop.

To obtain sign information without resorting to auxiliary modulation, a vector type detector operating at the crystal frequency must be used. A vector detector requires a phase reference, and port 1 of the bridge is the obvious choice. A vector voltmeter connected to ports 1 and 2 could be made to work after a fashion as an null detector, but has the problem that the phase slope reverses outside the 3 dB point of the crystal, which would cause convergence problems. Also, the voltage at port 1 varies widely in amplitude. On the other hand, the current at port 1 is constant, and its relationship to the voltage at port 2 is given by z_{21}, a function having the ideal form to steer a servo loop. If it were easy to compute the quotient z_{21}, the problem would be solved. Unfortunately, division is difficult to implement using analog computer techniques at RF.

However, in the time domain, the following observations can be made. Assume that the ALC fixes i_1 at $I_X \cos \omega t$, where I_X is the magnitude of the virtual crystal current (twice the actual crystal current). Then $v_2 = I_X z_{21} = I_X(1/C_m)(\Delta\omega_0/\omega_0)(-\sin \omega t)$. Now if $-I_x \sin \omega t$ (i.e. i_1 retarded by

90°) is multiplied by v_2, the result is $(I_X^2/C_m)(\Delta\omega_0/\omega_0)(\sin^2 \omega t) = (I_X^2/C_m)(\Delta\omega_0/\omega_o)(1-\cos 2\omega_0 t)/2$, which has the DC component $(I_X^2/2C_m)(\Delta\omega_0/\omega_0)$. The multiplication is implemented by a Gilbert multiplier. Although the preceding mathematics imply that both ports of the multiplier operates in the linear region, this is not necessary and the LO port is actually operated in the saturated mode. Under these conditions, the multiplier operates in manner more similar to a synchronous demodulator than a phase detector.

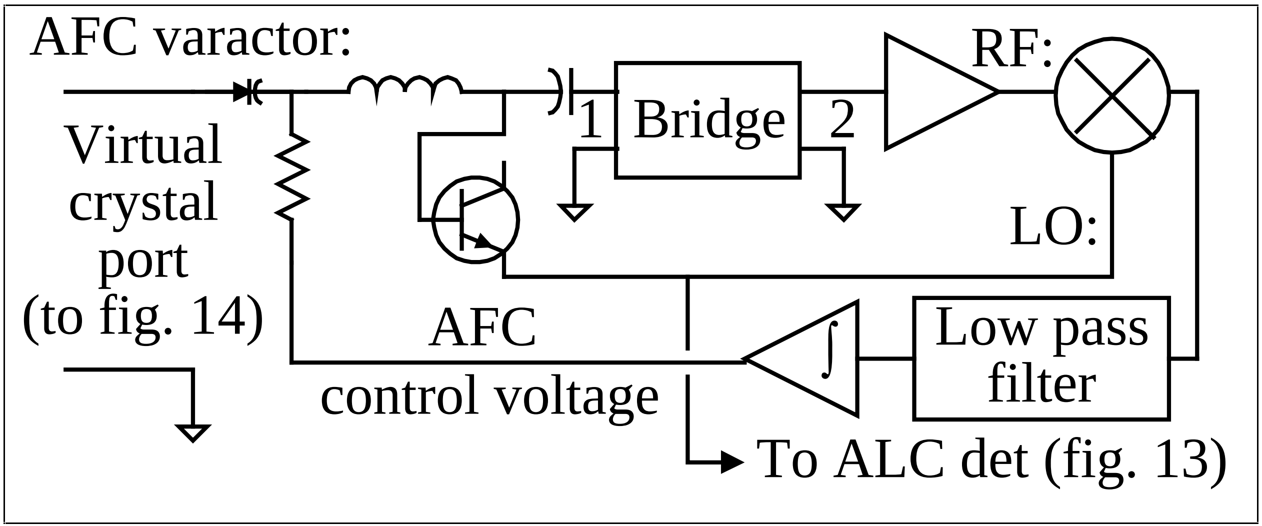

Figure 11. AFC system block diagram.

Although the LO port is saturated, it is still necessary to keep the drive level under control because of limited dynamic range within the Gilbert multiplier. The LO port is driven by a phase-retarded voltage derived from i_1. This provides a constant amplitude LO signal, since i_1 is stabilized by the oscillator ALC. The RF port of the multiplier is driven by v_2, with about 30 dB of gain added by the AFC RF amplifier. The DC output of the multiplier is then extracted by a low pass filter (fig. 11). A capacitor is inserted in series with port 1 of the bridge to convert i_1 to a voltage with an inherent 90° phase lag. Ideally, the voltage directly across the capacitor would be utilized. However, it is much easier to use the voltage from the capacitor to ground. If $X_c >> z_{11} = R_m/2$, the phase error contributed by v_1 will be negligible. An emitter follower buffers this voltage to drive the LO port of the multiplier. The emitter follower also drives the ALC detector of fig. 13 that is described below. The reactance added by the port 1 series capacitor is cancelled out by an inductor of equal and opposite reactance.

3.5. AFC loop design

The DC output from the synchronous demodulator drives a lead-lag type integrator (fig. 11) that controls the bias of the AFC varactor, which is in series with port 1. The operation of the servo loop pulls the frequency of oscillation up or down as necessary to keep the bridge balanced. This varactor has no effect on the frequency at which the bridge balances, but it does modify the resonant frequency of the virtual crystal and hence pulls the frequency of any oscillator connected to the bridge. It is convenient to increase the value of the inductor in series with port 1such that the combined reactance of the AFC varactor, inductor, and capacitor is zero when the AFC voltage is at mid-range. A loop time constant of about 1 second is used in the AFC loop which results in no measurable noise modulation of the oscillator by the AFC loop. In the event there is an abrupt perturbation to the oscillator (such as a step in the power supply voltage), a recovery time of a few seconds is necessary for the AFC loop to return the frequency to its original value.

3.6. Automatic level control (ALC)

In a precision crystal oscillator, the crystal drive current must be very well stabilized, because the crystal exhibits significant amplitude to frequency (AM/FM) conversion due to nonlinearities inherent in the quartz. Fig. 12 shows the AM/FM curve of a typical SC-cut crystal, as measured in the oscillator described in this paper by adjusting the ALC set point. The design goal was to hold the drive current stable to 0.0002 dB/°C, for a tempco contribution of about 10^{-12}/°C.

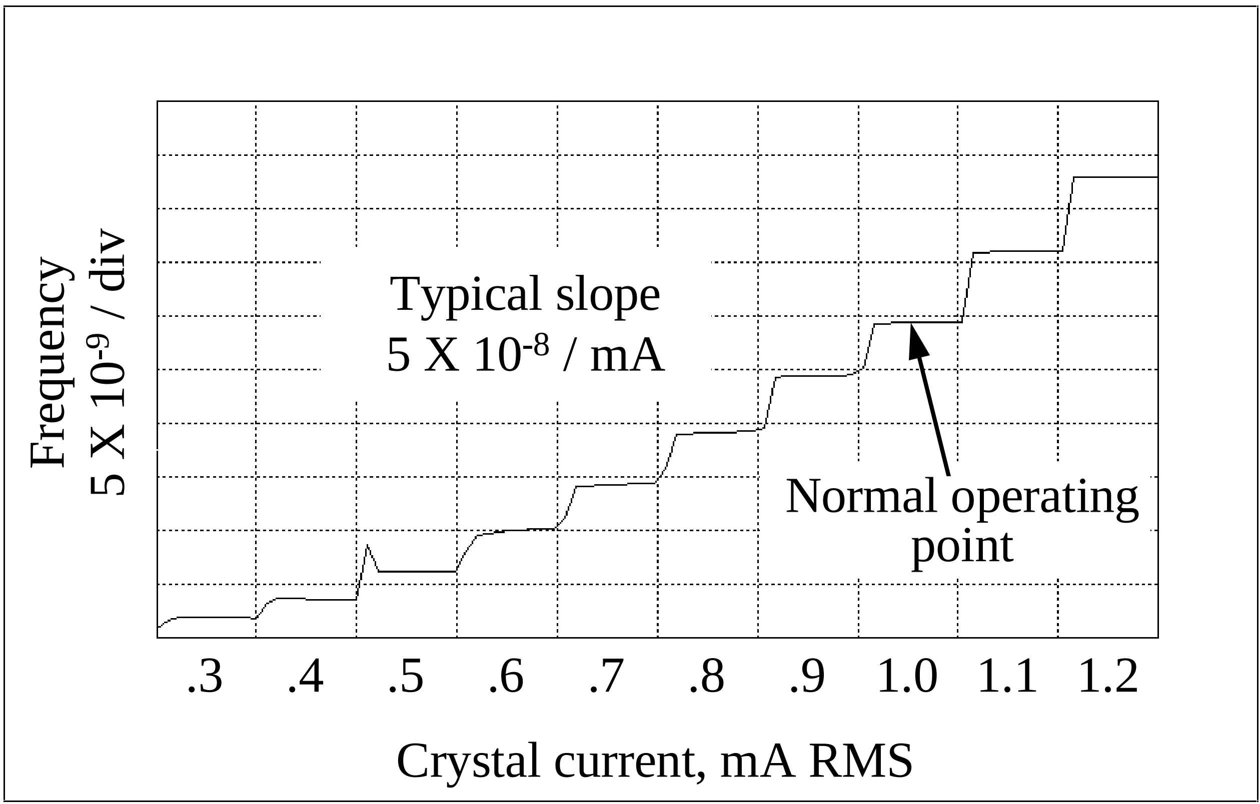

Figure 12. Drive level induced frequency shift.

A diode peak detector was chosen as the basis for the ALC detector because it is certainly the simplest method, and is probably as temperature stable as any method. Fig. 13 shows a conventional temperature compensated diode peak detector. A matched pair of thermally coupled diodes is used, one acting as an RF peak detector, the other as a DC reference subtracted from the rectified output. Ideally, the ≅ -2mV/°C tempcos are equal and cancel each other out. The compensation can never be perfect, even for perfectly matched diodes at identical temperatures, because the diodes operate under the different conditions of pulsed current vs steady state current. The temperature induced offset error in the detector system mainly consists of this residual tempco mismatch. In practice, this mismatch, for reasonable DC bias currents, is about ½%, corresponding to 10 μV/°C. For a given rectified DC output current, the residual tempco is not very sensitive to input voltage amplitude. Linearity and frequency response, traditionally important issues with diode detectors are irrelevant here since the amplitude and frequency are constant.

Thus it can be concluded that the error in the detector is dominated by offset voltage tempco. Since the absolute error doesn't increase significantly with amplitude, the amount of error relative to the voltage being measured is inversely proportional to that voltage. I. E., a 10 μV error in measuring a .1V signal represents 0.001 dB, while a 10 μV error measuring a 1V signal represents 0.0001 dB. Therefore, if a temperature stable way of stepping up the voltage to be measured can be devised, the effective ALC stability can be improved.

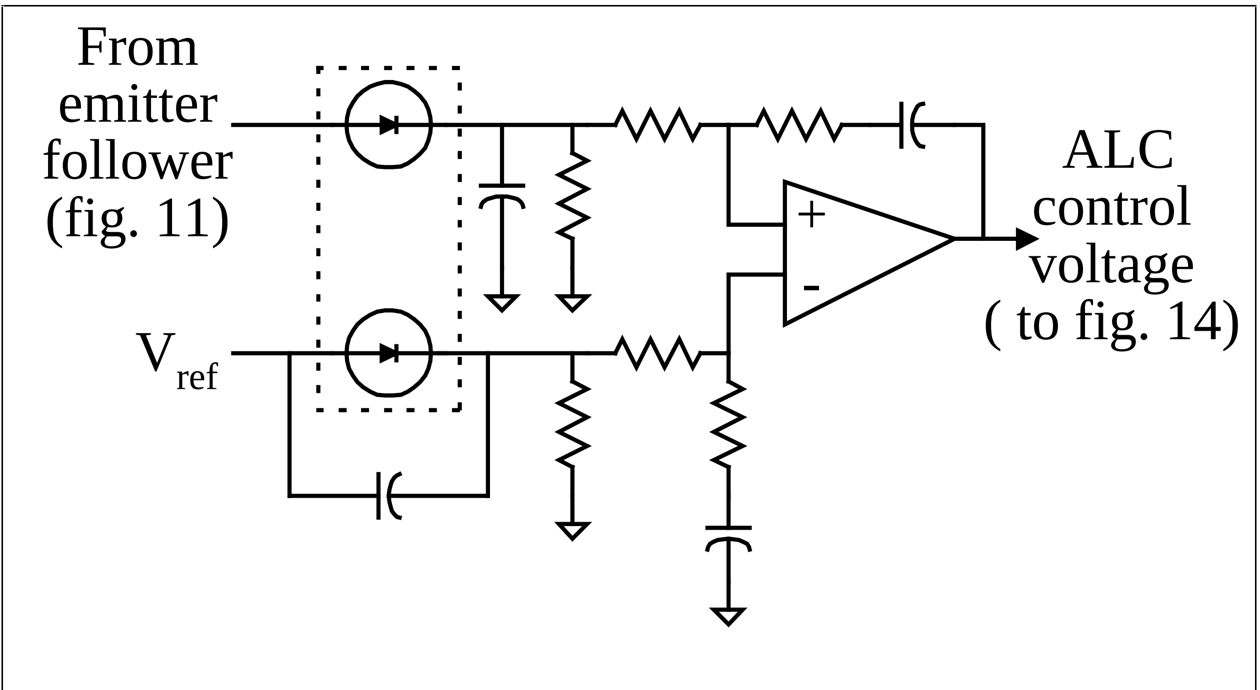

Figure 13. ALC detector.

The potential for such a method already exists in the circuit of fig. 11. By making the capacitor in series with port 1 very small, the voltage at the emitter follower can be made arbitrarily large. Of course, the series inductor has to be increased to compensate for the capacitor change. The temperature stability of the step up ratio can be guaranteed by using an NP0 capacitor. The emitter follower has a negligibly low tempco of gain because it is lightly loaded. The emitter follower in fig. 11 drives the diode detector of fig. 13, which is used to control the gain of the oscillator shown in fig. 14 as discussed below. The ALC integrator has standard lead-lag compensation and a time constant of about .1 second.

The use of a bridge to control the oscillator is essential as an enabling technology for the stepped up voltage ALC technique.[6] If it were attempted in an ordinary oscillator, the high temperature sensitivity of the compensating inductor in series with the step up capacitor would greatly degrade the stability of the oscillator, making the method unusable.

3.7. The oscillator proper

At this stage, the oscillator itself is almost an afterthought. The bridge should be compatible with most series resonant crystal oscillators. A bridged-T oscillator was chosen for this design (figures 1 and 14), because it is the simplest possible oscillator circuit for overtone SC cut crystals. A single pi-network tank provides feedback, overtone selection, and undesired mode suppression. By making the Q of the tank at least 10, it suppresses undesired mode B oscillations at 10.95 MHz, making an explicit mode suppressor superfluous.

Fig. 14 shows a simplified schematic of the oscillator proper. The circuit is based on the simple bridged-T oscillator of fig. 1, but has two modifications. The virtual crystal represented by the bridge (fig. 11) replaces the actual crystal, and the base now gets its feedback through the capacitive voltage divider rather than directly from C_b. The dotted lines in fig. 14 show the original topology of fig. 1. The capacitive voltage divider is utilized by the ALC system, as described below.

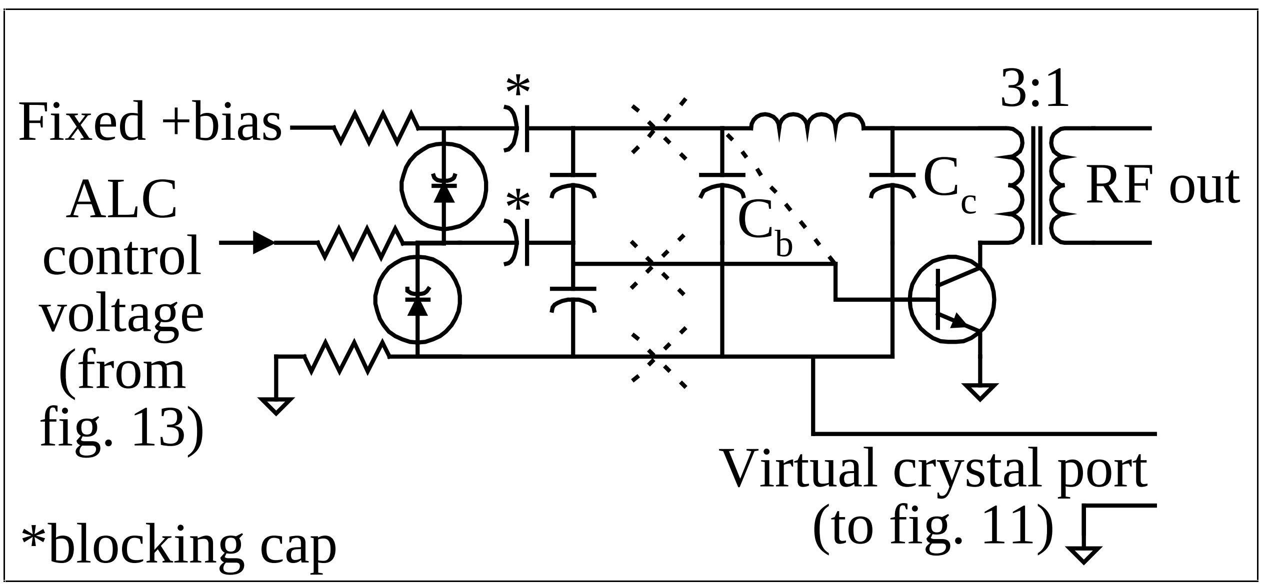

Figure 14. Bridged-T oscillator with gain control.

[6] Patented [15].

3.8. Oscillator gain control

The oscillator must have a gain control input to connect to the ALC integrator. Typically, this is done by varying the oscillator collector current. While clearly the simplest method, this requires that the oscillator run in a starved bias condition that increases distortion, and it is difficult to get a lot of dynamic range while keeping collector current within reasonable limits. Another problem is that collector current tends to get out of control at turn on, and can cause oscillator starting problems. In a bridge-controlled oscillator, extra dynamic range in the ALC is needed to maintain oscillation during AFC acquisition.

To address all of these ALC issues, the variable capacitive voltage divider in fig. 14 is used. A portion of the base capacitance has been replaced with a capacitive voltage divider made from two varactors. The base operates from the divided-down feedback voltage, hence the division ratio controls the loop gain. The division ratio is controlled by the ALC control voltage, which causes the bias voltages of the two diodes to move in opposite directions as the control voltage varies. Thus at low ALC voltage, the upper capacitor has high voltage, hence low capacitance and high reactance. The lower capacitor has low voltage, hence high capacitance and low reactance. This causes the base voltage to be a small fraction of the voltage across C_b. The situation is reversed at high ALC voltage resulting in most of the C_b voltage being available at the base.

If varactors with a $\gamma = 1$ characteristic are used, the equivalent capacitance of the two diodes in series is invariant as the ALC voltage is varied. However, varactors having sufficiently high capacitance for this application are only available with a $\gamma = 2$ characteristic. Shunt capacitors (fig. 14) can be added to the varactors to "linearize" them fairly effectively in the sense that the series equivalent capacitance is constant to within a few percent over the ALC range. This together with the desensitizing effect of C_b results in a design where the total base capacitance is nearly independent of ALC voltage. Thus, the ALC loop does not significantly perturb the AFC loop; i.e. inducing a change in ALC voltage by adding extra loading to the oscillator has a negligible effect on AFC voltage. The reverse is not true; when the AFC control voltage is at one of the rails (maximum pulling), the ALC control voltage rises substantially to counteract the inevitable gain reduction that accompanies frequency pulling.

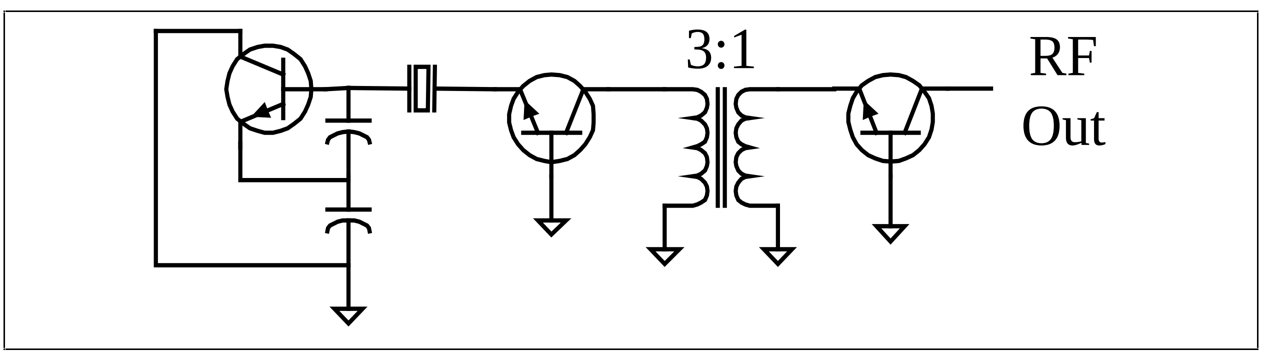

Figure 15. Burgoon low phase noise output circuit.

3.9. Extracting the output

Fig. 15 shows a conventional crystal oscillator with an output scheme having the lowest possible phase noise floor [6,13]. The right hand terminal of the crystal would normally go to ground in a Pierce configuration. Instead it is returned to ground through the very low input impedance of the grounded base amplifier, thus only minimally disrupting normal oscillator operation. Many kHz from the frequency of oscillation, where the phase noise floor is reached, the crystal impedance is so high as to essentially constitute an open circuit. Open circuiting the emitter eliminates base recombination current noise, leaving only collector shot noise current. The output of the first grounded base amplifier is then stepped up so that its noise dominates the noise of any subsequent buffer amplifiers.

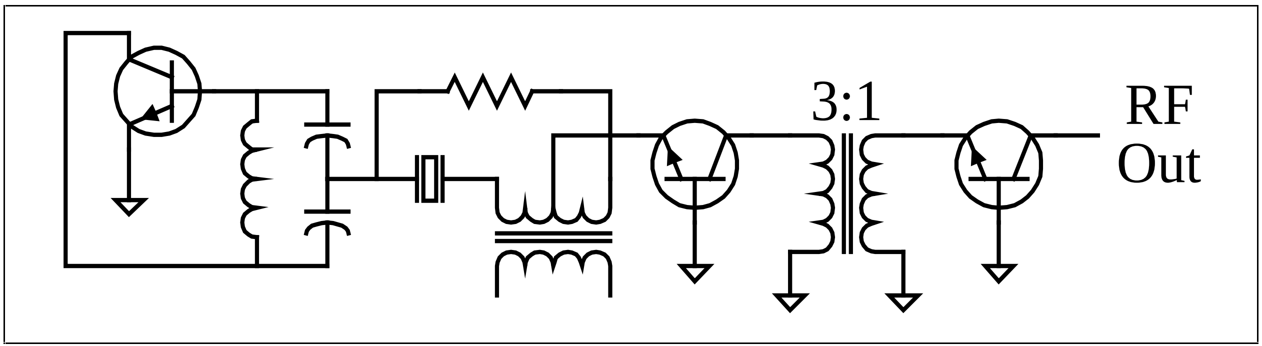

Figure 16. Output adaptation to bridge oscillator.

Fig 16 shows a way to adapt this technique to the bridge-controlled oscillator. In this case, the transformer center tap, which normally goes to ground, now returns to ground through the input impedance of the grounded base transistor. This topology is certainly usable, but it does have the disadvantage that the buffer amplifier must be colocated with the oscillator. Also, there is some risk that the extra impedance in series with the center tap could be detrimental to the accuracy of the bridge.

The method actually used, however, was to place the load in series with the collector of the oscillator transistor, using a current transformer to make the load appear to be floating (fig. 14). This can then be used to a drive subsequent grounded base amplifier. It can be shown that this configuration, like the ones of figures 15 and 16, also has optimally low phase noise due to lack of base recombination current noise. There are only two connections from the oscillator in fig. 14 to RF ground: The emitter and the virtual crystal port. Since this port emulates a crystal, it, too, is effectively an open circuit at phase noise floor frequencies. Since nothing else in the circuit has any RF path to ground either, the emitter, which is grounded, has no RF path back to the base or collector. Therefore, for noise purposes, the emitter is open circuited thus eliminating base recombination current. It should be noted that the bridge circuit is an enabling technology for this technique; otherwise the transformer would degrade the tempco.

The output impedance at the secondary is fairly high because of the high output impedance of the transistor. If it is used to drive a grounded base amplifier, that amplifier will effectively be driven by a current source and therefore have no base recombination noise. The current transformer quadruples the secondary current so that the oscillator shot noise dominates over any subsequent grounded base buffer amplifier shot noise.

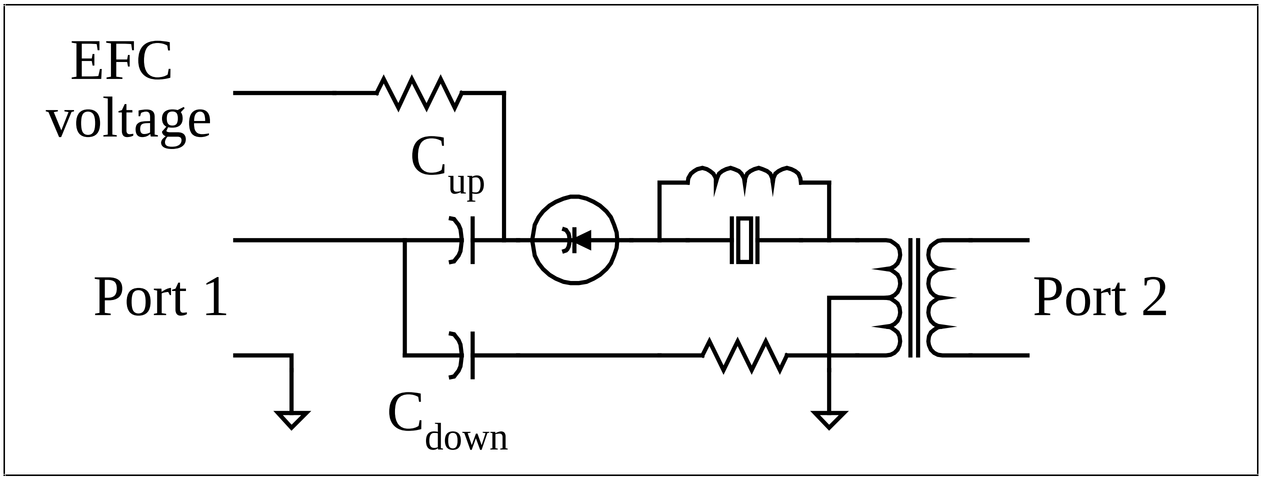

Figure 17. Adding EFC to bridge.

3.10. Electronic frequency control (EFC) capability

For many applications it is desirable to have EFC capability so that the oscillator can be used in VCXO mode. A method of doing this is shown in fig. 17. Note that the EFC varactor located inside the bridge should not be confused with the AFC varactor located outside the bridge. The EFC varactor changes the frequency at which the bridge balances. The AFC varactor is used to force the oscillator to oscillate at this bridge balance frequency. The crystal resonant frequency can be pulled above series resonance in a straightforward manner by adding series capacitance with the EFC varactor and/or C_{up}, a "pull up" capacitor that can be used to remove crystal calibration frequency error. The bridge enables an additional unique feature. A "pull down" capacitor, C_{down}, in series with the image resistor allows tuning below series resonance. The operation of the bridge makes the pull down capacitor behave as if it adds negative load capacitance to the crystal; hence it has the unusual property of decreasing the frequency as its value decreases. This technique was previously applied to a Meacham bridge oscillator [13].

In practice, if the crystal has been fabricated to be series resonant at the exact design frequency, the pull down capacitance is chosen to be equal to the varactor capacitance at mid-range, and the pull up capacitor is increased to an effectively infinite value. Then at mid- range, the varactor and pull down capacitor balance out and the bridge balances at series resonance. It can then be tuned equal amounts above and below resonance as the EFC voltage is varied. If the crystal has a frequency calibration error, it can be corrected by changing the values of C_{up} and C_{down} appropriately. C_{up} acts as a conventional trimmer and pulls the frequency *higher* as it gets *smaller* while C_{down}, unlike conventional crystal trimmer capacitors, pulls the frequency *lower* as it gets *smaller*.

3.11. EFC varactor tempco

Varactors have a typical capacitance tempco of 500 ppm/°C. If the maximum varactor pulling is 10 Hz, this will result in a frequency tempco of 5×10^{-10}/°C. This is much worse than the rest of the oscillator, so it is essential to plan on colocating the EFC varactor with the crystal to obtain high thermal gain.

3.12. Bridge balance issues

The problem of balancing the bridge involves two degrees of freedom: the crystal must be at series resonance to remove any reactive imbalance, and the image resistor must match the ESR of the crystal to remove any resistive imbalance. In practice, the AFC loop has no difficulty servo'ing out the reactive imbalance by controlling the frequency. The resistive imbalance must be dealt with on an open loop basis. The ESR of crystals varies somewhat from unit to unit and is also has a moderate dependence on temperature and even a slight dependence on drive level. As a practical matter, then, the image resistor needs to be adjustable so that the bridges in oscillators can be resistively balanced on an individual basis. Setting the image resistor adjustment such that the signal at the RF input of the multiplier reaches a null makes this adjustment. This null must at least be good enough to avoid exceeding the dynamic range of the multiplier. It becomes more critical as the gain of the RF amplifier increases, which puts a practical limit on gain.

The bridge balance is affected somewhat by the value of the EFC voltage, because the ESR of the varactor is a function of bias voltage. However, this effect can be reduced to an acceptably low amount (equivalent to an image resistor error of about 1 ohm) by proper choice of varactor, component values and topology in the bridge.

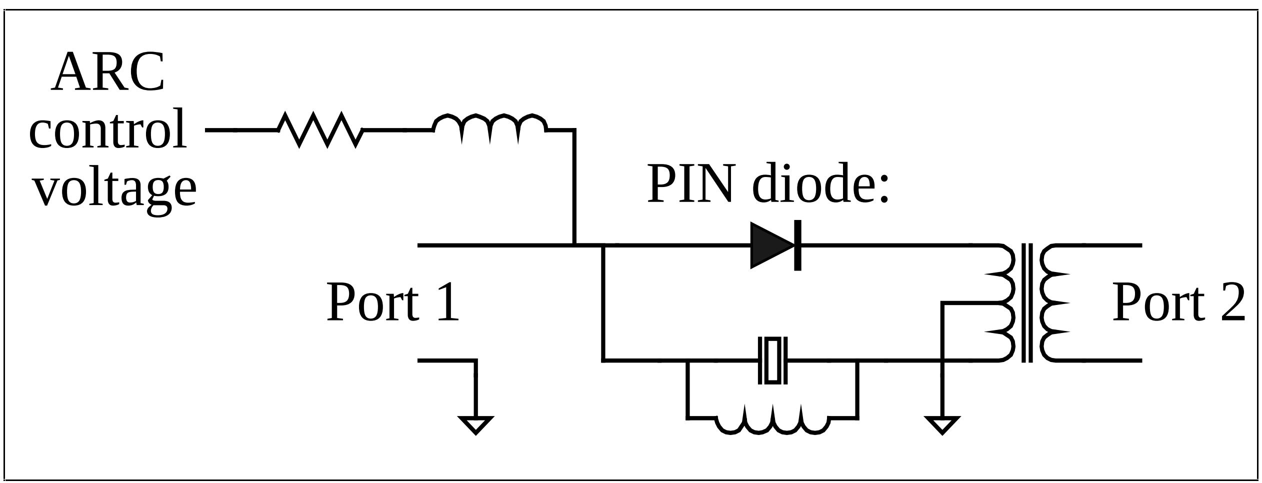

Figure 18. Adding ARC to bridge.

3.13. Servo control of image resistance

A second servo loop operating in quadrature with the AFC loop can be added to automatically achieve resistive balance. Fig. 18 shows how a PIN diode can be substituted for the image resistor in the bridge. Fig. 19 shows the quadrature servo loop that controls the RF resistance of the PIN diode. This automatic resistance control (ARC) loop has the same error detection architecture as the AFC loop, except that the error signal input to the synchronous detector is shifted 90°. The error signal after integration is used to control the bias current of a PIN diode used in place of the image resistor. The extra cost of the ARC loop may be economically justified by the labor saved by not having a manual adjustment. It also reduces temperature induced frequency error as explained later in the section on error sources. Experimental results of an ARC loop are given later in this paper.

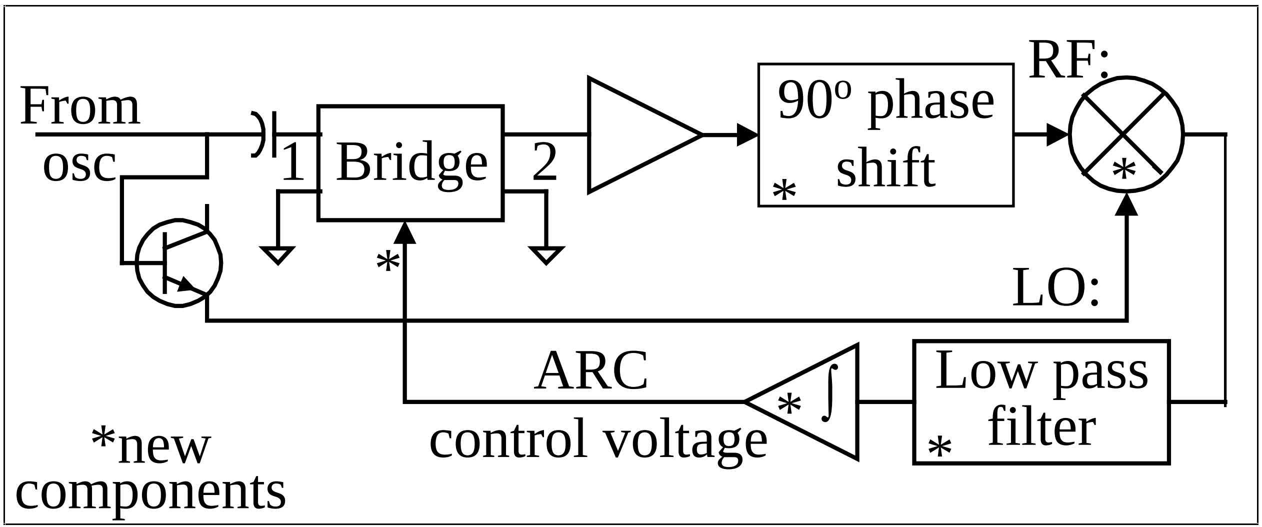

Figure 19. Automatic resistance control (ARC) loop.

3.14. Principal error sources

The main cause of frequency drift in the bridge-controlled oscillator is the combination of image resistor error and AFC phase error. Image resistor error, discussed above, results in an error signal at the output of the bridge that cannot be servo'ed out by the AFC loop. If the multiplier inputs are in exact quadrature with respect to the i_1 and v_2 of the bridge, the AFC loop will respond only to reactive bridge imbalance and will have zero sensitivity to resistive bridge imbalance caused by image resistor error.

The problem comes about when both image resistor error and phase error are present. With respect to the AFC loop behavior, a phase error of ϕ at the multiplier effectively adds a phase shift of ϕ to the impedance of the bridge components. Thus, an image resistor error (denoted ΔR) will be converted to a partially reactive error. For example, if $\Delta R = 1$, and the $\phi = 6°$, there will appear to be an additional reactance of $1 \sin 6° = 0.1$ ohms in series with the crystal. The AFC loop will "correct" the apparent reactive imbalance by shifting the frequency to create an actual reactive imbalance of equal and opposite value, resulting in a frequency error. In the example, the frequency error will be equal to whatever pulling effect 0.1 ohms of reactance would normally have on the crystal. The frequency error (Hz) is given by $f_{error} = (\Delta r)(\phi)/(720L_m)$ where Δr is the resistive imbalance (ohms), ϕ is the phase error (°) and L_m is the crystal motional inductance (H). For the crystal used in this oscillator, $L_m = .7H$, resulting in an error slope of 0.01 Hz/ohm for a 5° phase error.

An alternate description of this problem can be made along the lines of the Q multiplication discussion in [7] as follows: for the crystal used here, $R_m = 40$ ohms, hence $Q_u = 1.1$ million at 10 MHz, which is equivalent to a 3 dB/±45° bandwidth of 9 Hz. This implies a frequency shift of .5 Hz for a 5° phase shift in Z_m. This is 40 times as much frequency shift as would be caused by 5° phase error at the synchronous demodulator with $\Delta R = 1$ ohm. This is equivalent to a Q multiplication of 40:1 where the ratio $R_m/\Delta R = 40/1 = 40$. A critical difference here is that there is no fundamental limitation on the amount of Q multiplication that is possible.

The effective phase error can easily be measured on an oscillator by making a slight adjustment of the image resistance pot. The frequency before vs after adjustment is measured with a counter and compared to the pot resistance before vs after as indicated on an ohmmeter. When the phase error is zero, the frequency is unaffected (to a first order) by the pot setting.

Since the object is to build an oscillator, not a network analyzer, stability is the primary objective, as opposed to absolute accuracy. To this end, it is relevant to take the temperature derivative of the frequency error equation introduced above:

$$
\begin{aligned}
df_{error}/dT &= d/dT\,[(\Delta R)(\phi)/720L_m] \\
&= [d\Delta R/dT]\,[\phi]/720L_m + [\Delta R][d\phi/dT]/720L_m
\end{aligned}
$$

Resistive balance in a bridge should be relatively insensitive to temperature assuming quality resistors are used. The implication of this with respect to the first term above is that it is probably not worth the effort to accurately adjust the phase error to exactly zero. On the other hand, the phase shift at 10 MHz in the ALC RF amplifier is probably always going to be somewhat temperature sensitive. The implication of this with respect to the second term is that it is worth some effort to minimize the image resistor error, especially since it has to be adjusted anyway.

3.15. Effect of ARC loop on errors

If the previously mentioned ARC loop is implemented, then it is possible to reduce the resistive imbalance to negligible proportions. With that factor effectively at zero, the phase shift in the AFC loop no longer is a factor. Under these conditions, the dominant error outside the bridge becomes the AM/FM conversion resulting from ALC drift. This is difficult to reduce below about 5×10^{-12}/°C. Even if a perfect peak detector were to be realized, there is some residual drift due to temperature dependent harmonic distortion, because harmonics affect the peak detector differently than the crystal.

The preceding discussion assumes does not take into account any possible tempco of the PIN diode itself. Certainly the RF resistance of the PIN diode has a non-zero tempco. Unlike the case of the EFC varactor, which is not inside a feedback loop, the PIN diode is controlled by the ARC loop. Hence in principle, it shouldn't matter if the resistance varies with temperature because the servo loop will readjust the diode current to keep the bridge balanced. It might matter if the variation in current, as a side effect, changes the parasitic reactance of the diode. For that matter, this reactance may have a tempco in its own right. Further discussion of this appears in the section on experimental results of the ARC loop.

4. Experimental results

4.1. Temperature

Fig. 20 shows the frequency vs. temperature of a typical oscillator (*circuit only*), with a tempco of about 2×10^{-11}/°C. Some oscillators have been better than 10^{-11}/°C. In this measurement, the crystal and oscillator were ovenized in separate ovens, and the set point of the oscillator oven only was incremented. No EFC varactor was installed during these temperature tests.

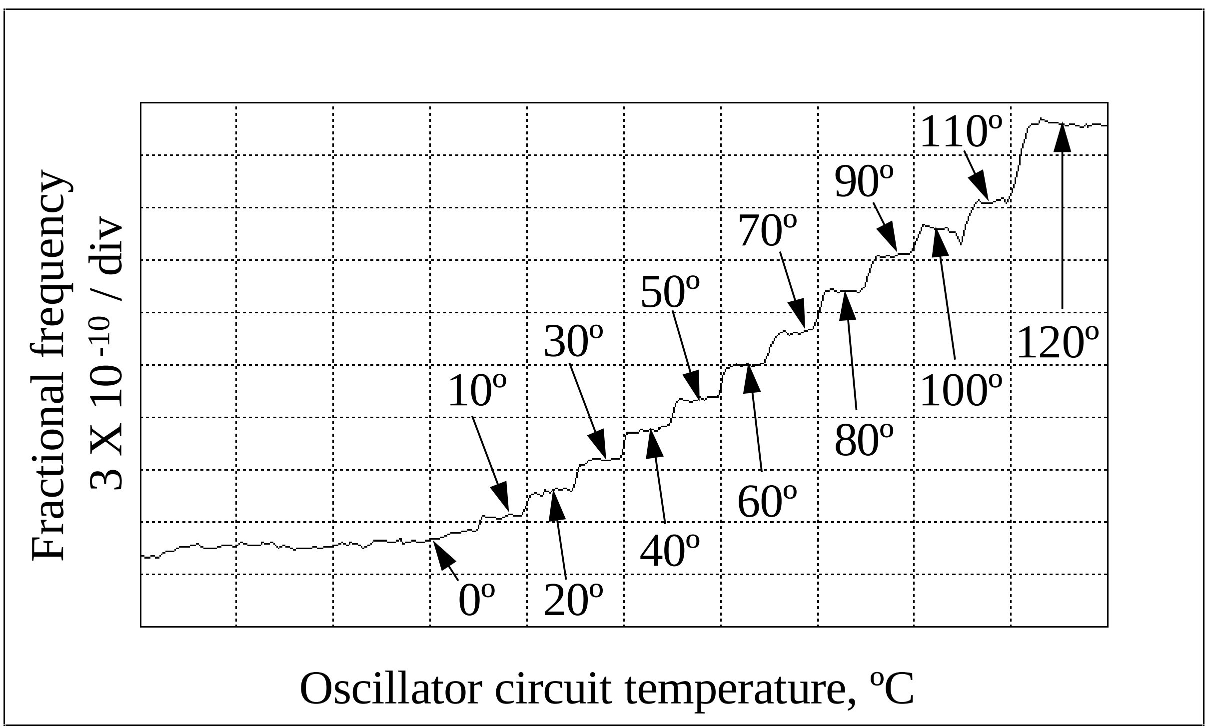

Figure 20. Oscillator circuit tempco (less crystal).

4.2. Humidity

Fig. 21 shows the effect of humidity. It is important to note that this test was performed on an **unsealed** oscillator. A frequency baseline was established at an ambient temperature of 70° C and a relative humidity of a few percent or less. Then sufficient moisture was added to raise the relative humidity to 25% for an hour and a half, at which point the moisture was removed. Although 25% R.H. doesn't sound very high, at 70° C it results in a dewpoint of 32° C (87° F), meaning that the absolute humidity is equivalent to 100% R.H. at 32° C.

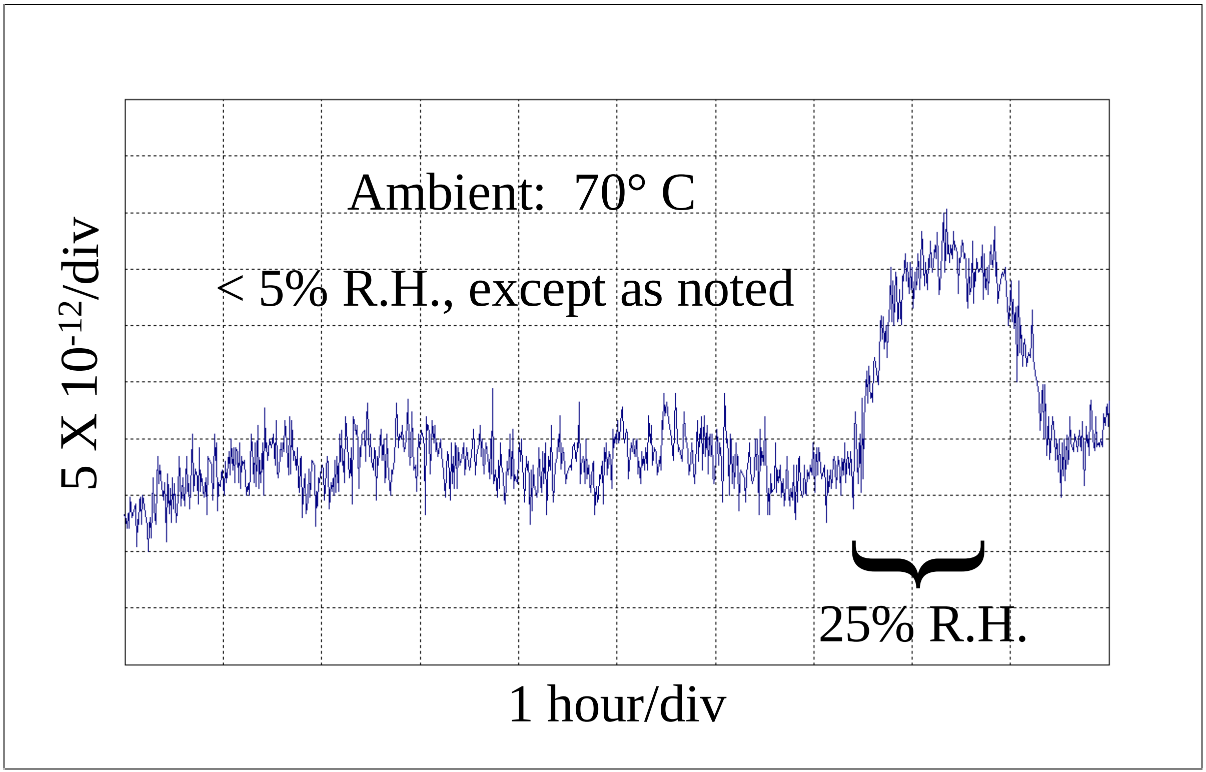

Figure 21. Effect of humidity on frequency.

Experimental results consistently indicate that for humidity effects (as opposed to condensation effects), absolute humidity, not relative humidity, is the critical parameter. Of course, for an ovenized oscillator, condensation is not an issue.

The reason for conducting the tests at such an elevated temperature is that it permits fast changes in humidity to be made easily and completely avoids the danger of condensation in the test chamber and icing of the evaporator coils in the cooling system. The results show a reversible frequency excursion of about 2 parts in 10^{11} due to humidity. It is believed that most of this effect is in the bridge section. In any event, no special precautions had to be observed for selecting components outside the bridge for low humidity sensitivity, whereas inside the bridge, some components had to be replaced with ones having lower sensitivity to humidity, although it was not necessary to use hermetically sealed components anywhere. (In fig 21, a small amount of linear aging has been removed for clarity.)

4.3. Power supply and load

Fig. 22 shows the unmeasurably high immunity to changes in output loading, and power supply voltage. These tests were performed directly on the oscillator circuit with no intervening buffer amplifier and no voltage regulators (except for the ALC voltage reference).

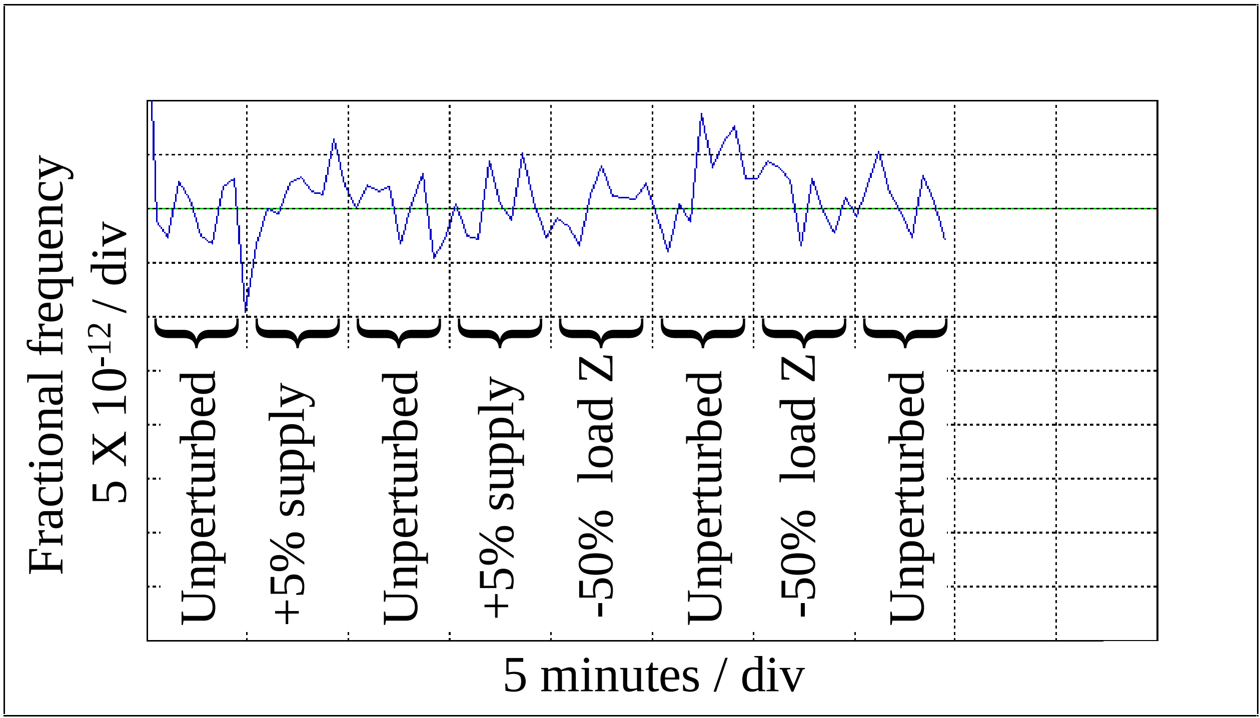

Figure 22. Power supply and load impedance effects.

4.4. Repeatability

To check the repeatability of the oscillator design, the same crystal was installed in 25 different oscillator boards. The frequencies of oscillation all fell within 5×10^{-8} without individual adjustment.

4.5. ARC loop performance

Prior to constructing an ARC loop, some PIN diodes were characterized. As mentioned before, tempco of resistance wasn't an issue. However, reactance was of interest. Diode capacitance can be measured at low currents where the diode resistance is at least several hundred ohms. A value of about .5 pF was fairly consistent as diode resistance was varied and seemed to have no tendency to vary as the current was increased. It also had no measurable tempco, although the resolution of the measurement was not great. A typical crystal has a resistance of 40 ohms, hence the diode will operate at 40 ohms as well. A model of 40 ohms in parallel with .5 pF at 10 MHz is equivalent to a series model of 40 ohms and .3 μF. The frequency pulling effect of that value of capacitance on a typical crystal is less than 10^{-9}. If it had a tempco of 1000 ppm/°C, it would contribute less than 10^{-12}/°C. The conclusion after the previous analysis seemed to be that the PIN diode would be an excellent candidate for the ARC loop and not pose a threat to temperature stability.

In practice with an actual ARC loop, it was found that the tempco of the diode was in the neighborhood of 10^{-10}/°C. It is not known what causes this effect, but it may be nonlinearity in the diode. The diode had a carrier lifetime of 1.5 μs, which is adequate for 10 MHz but not high enough to eliminate all distortion. The practical implication of this is that it is necessary to highly ovenize the PIN diode as the EFC varactor diode and crystal must be. This is not much of a burden since the PIN diode is in the bridge interconnected closely with the crystal and is considerably smaller than the crystal.

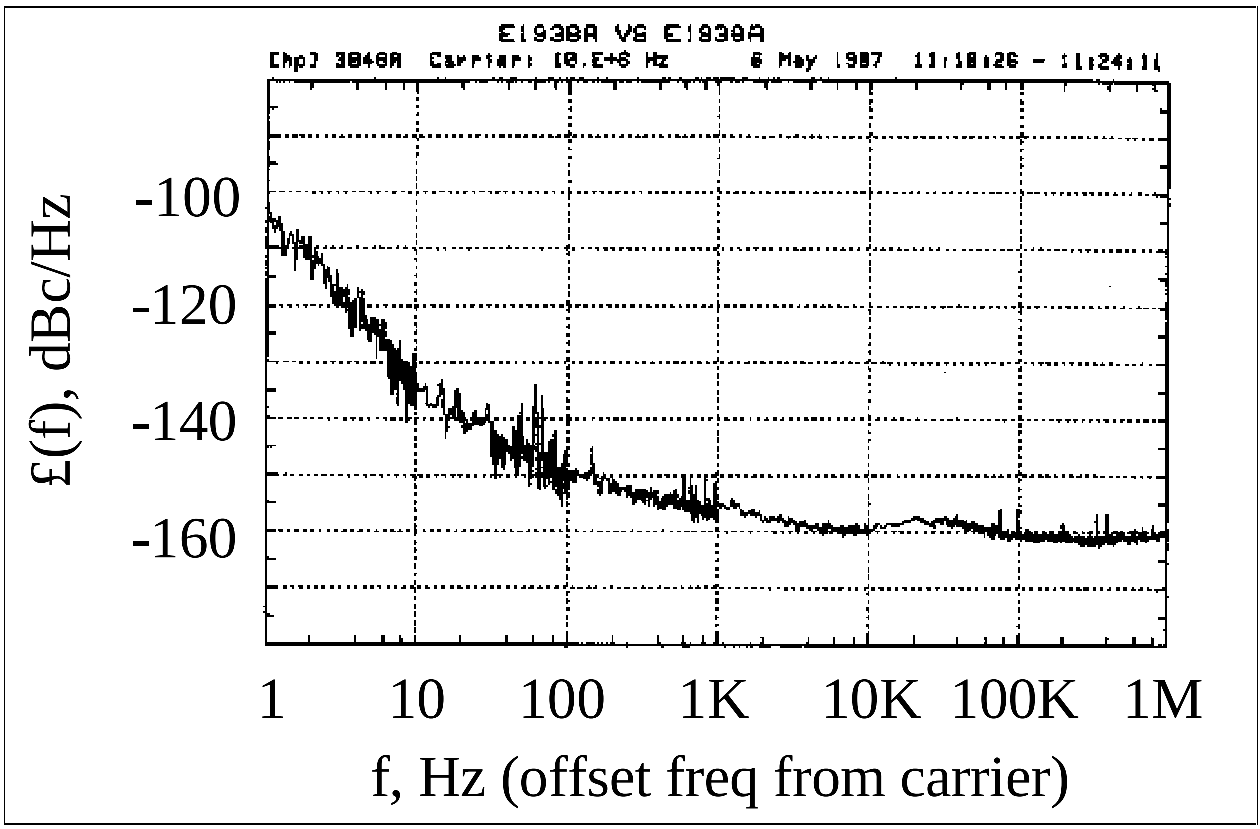

Figure 23. Phase noise.

4.6. Phase noise

Fig. 23 shows the measured phase noise. The close in phase noise is comparable to a conventional oscillator using a similar crystal. The phase noise floor of at -160 dBc/Hz (which includes a buffer amplifier) is reasonable for this class of oscillator, but theoretical noise calculations indicate it is capable of being improved about 10 dB by optimizing the output transformer.

4.7. Short term stability

Fig. 24 shows the short term stability. Again, this is comparable to a conventional oscillator using a similar crystal. The rise around 10 to 100 seconds is believed to be due to the measurement system. Hence there is experimental verification that the AFC loop is not adding significant noise modulation to the oscillator.

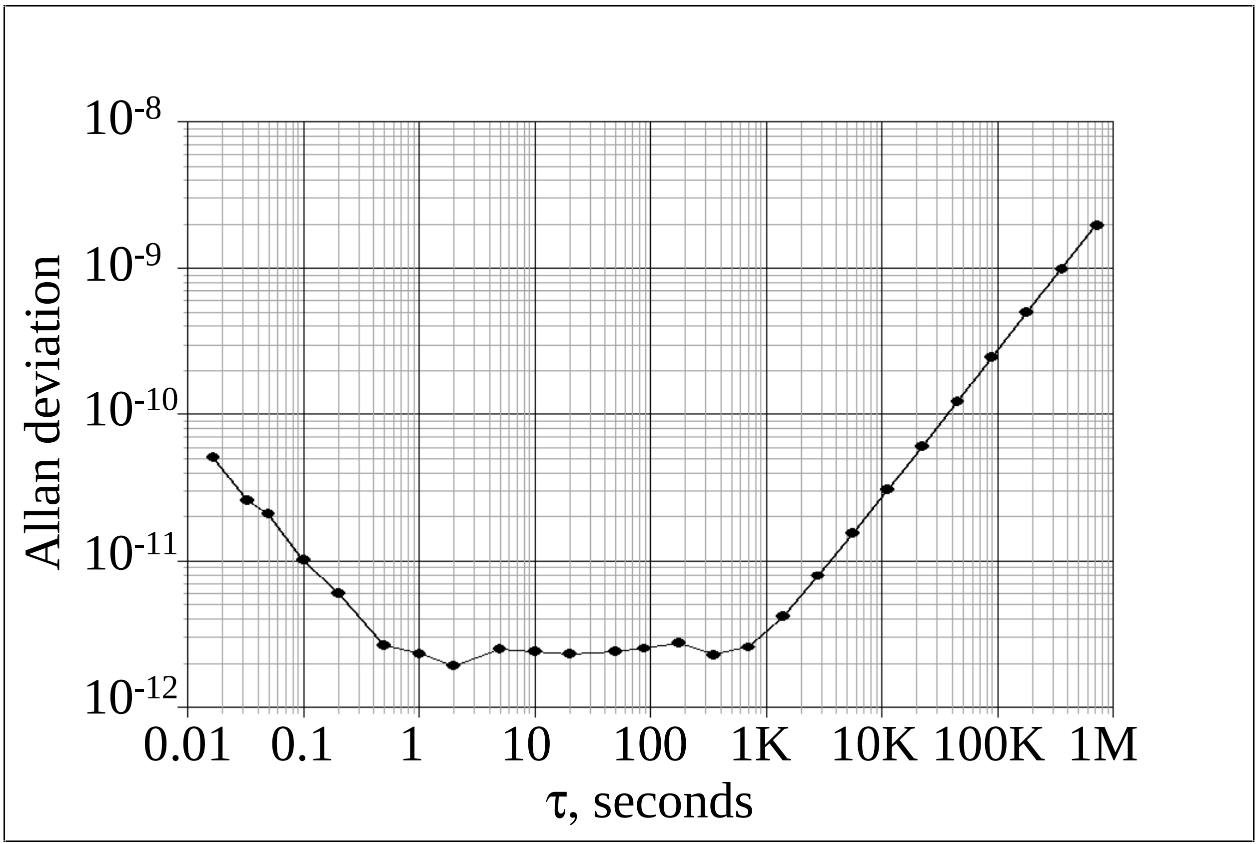

Figure 24. Short term stability.

5. Conclusions

A practical technique for controlling the frequency of a crystal oscillator with a bridge has been described. The high performance of this technique has been demonstrated. Critical design issues have been outlined along with methods of resolving them.

5.1. A comment on complexity

The oscillator described here may seem to have the disadvantage of being more complex than a conventional oscillator. Additional circuitry consisting of the bridge, RF amplifier, Gilbert multiplier and AFC integrator has been added. However, no ovenized buffer amplifiers or voltage regulators are necessary. The AFC integrator can be half of a dual op amp, the other half being the ALC integrator. Even a conventional oscillator needs an ALC (or a very good limiter) because of the non-linearity of quartz. Therefore, the increase in complexity is not severe. In the implementation used here, the oscillator occupied about 16 cm^2 (2.5 in^2) of space on its PC board, not counting the crystal.

6. Acknowledgments

Len Cutler suggested some tests for possible sources of error. Robin Giffard supplied a dual transistor Butler oscillator for testing. Jack Kusters assisted with characterization. Lee Cosart and Matt Semersky wrote test software. Eric Ingman did the humidity characterization and suggested many improvements to enhance humidity performance. Charles Adams was the project manager. This work was done while the author was with the HP Santa Clara Division.

7. References:

[1] R. Karlquist, et al, "A Low Profile High Performance Crystal Oscillator for Timekeeping Applications," Proceedings of the 1997 IEEE International Frequency Control Symposium, pp. 873-884, 28-30 May 1997.

[2] W. Edson, Vacuum Tube Oscillators, New York, John Wiley & Sons, 1953, pp. 212-218.

[3] Ibid, pp. 205-208.

[4] Ibid, pp. 142-151.

[5] F. Butler, "Series-Resonant Crystal Oscillators," Wireless Engineer, V. 23, pp. 157-160, June, 1946.

[6] J. Burgoon and R. L. Wilson, "SC-Cut Oscillator Offers Improved Performance," Hewlett-Packard Journal, V. 32, pp. 20-29, March 1981.

[7] L. Meacham, "Bridge-Stabilized Oscillator", Proc. IRE, V. 26, pp. 1278-1294, Oct. 1938.

[8] L. Ferris, "Transimpedance Oscillator Having High Gain Amplifier," U.S. Patent #4,358,742, Nov. 9, 1982.

[9] T. Pendleton, A System for Precision Frequency Control of a 100 kc. Oscillator by Means of a Quartz-Crystal Resonator, M.S. Thesis, Univ. of Maryland, May 1953.

[10] P. Sulzer, "High Stability Bridge Balancing Oscillator," Proc. IRE, V. 43, pp. 701-707, June 1955.

[11] M. Van Valkenburg, Introduction to Modern Network Synthesis, New York, John Wiley & Sons, 1953, pp. 337-338.

[12] A. Zverev, Handbook of Filter Synthesis, New York, John Wiley & Sons, 1967, pp. 423-425.

[13] J. Burgoon, "Crystal Oscillator Having Low Noise Signal Extraction Circuit," U.S. Patent # 4,283,691, Nov. 5, 1981.

[14] A. Benjaminson, "A Crystal Oscillator With Bidirectional Frequency Control and Feedback ALC," Proceedings of the 40th Annual Symposium on Frequency Control, pp. 344-349, May 28-30, 1986.

[15] R. Karlquist, "Bridge-Stabilized Oscillator Circuit and Method," U.S. Patent # 5,708,394, Jan. 13, 1998.