Fundamentals Of Crystal Oscillator Design

electronicdesign.com/analog/fundamentals-crystal-oscillator-design

Theron Jones, Maxim Integrated | Sep 07, 2012   September 7, 2012

[Download this article in .PDF format](#)

Appropriately cut quartz crystals can be used as high-quality electromechanical resonators. Their piezoelectric properties (voltage across the crystal deforms it; deforming the crystal generates a voltage) allow them to be the frequency-determining element in electronic circuits. Crystals are widely used in oscillators, timebases, and frequency synthesizers for their high quality factor (QF); excellent frequency stability; tight production tolerances; and relatively low cost.

This article covers the primary design considerations for fundamental-mode oscillators using AT-cut crystals. These include load capacitance; negative resistance; startup time; frequency stability versus temperature; drive-level dependency; crystal aging; frequency error; and spurious modes. This information is based on experience from more than a decade designing ISM-band (industrial, scientific, medical) radios. (Topics relevant in other types of radio systems, such as crystal oscillator

phase noise, aren't limiting factors in ISM radios and aren't covered.)

Table Of Contents

Crystal Modeling Basics

Quartz crystals are modeled electrically as a series LCR branch in parallel with a shunt capacitance _(Fig. 1)_. The series LCR branch, often called the motional arm, models the piezoelectric coupling to the mechanical quartz resonator. The shunt capacitance represents the physical capacitance formed by both the parallel plate capacitance of the electrode metallization and the stray package capacitance.

1. The fundamental resonant mode of a quartz crystal can be modeled as an LCR network shunted by a capacitor.

For crystals operating in the fundamental mode with a 5-MHz to 30-MHz frequency range, typical values of the circuit elements are:

- C1: 2 fF to 20 fF (motional capacitance)
- R1: 10 Ω to 150 Ω (equivalent series resistance, ESR)
- L1: Determined by C1 and the operating frequency (motional inductance)
- C0: 0.5 pF to 5 pF (shunt capacitance)

For a series LCR circuit with no driving voltage, summing the voltages across the elements produces:

$L*dI/dt + I*R + (1/C)*\int I*dt = 0$

By definition, dQ/dt can replace I:

$L*d^2Q/dt^2 + R*dQ/dt + Q/C = 0$

Multiplying both sides by L gives:

$d^2Q/dt^2 + (R/L)*dQ/dt + Q/(L*C) = 0$

which is of the form:

$d^2Q/dt^2 + (\omega_0/QF)*dQ/dt + Q*\omega_0^2 = 0$

This yields the well-known result for LCR circuits: the natural frequency ω_0 is the square root of the inverse of the product of the inductance and capacitance. In the simple mechanical model—mass, spring, dashpot—the forces applied to the crystal (ignoring gravity) accelerate the mass (F = ma) *(Fig. 2)*.

2. The mechanical model of a crystal oscillator is a simple compliance (spring)—inertia (mass)—damping (dashpot) system.

A simple linear model equates two forces: spring force and frictional force with mass multiplied by acceleration (Newton's Second Law). Hooke's Law (F = K*Y) provides the spring force, where K is the spring modulus and Y is displacement from equilibrium. Frictional loss is assumed proportional to the velocity of the dashpot's plunger and the dashpot's friction constant (D). Equating these forces (with no external driving force) gives:

$M*d^2Y/dt^2 + D*dY/dt + K*Y = 0$

Dividing both sides by M gives:

$d^2Y/dt^2 + (D/M)*dY/dt + Y*(K/M) = 0$

which is of the form $d^2Y/dt^2 + (\omega_0/QF)*dY/dt + Y*\omega_0^2 = 0$.

Because the electrical and mechanical models are assumed equivalent, the natural frequency of the mechanical system must equal the natural frequency of the electrical system. This yields:

$\omega_0 = \sqrt{(1/(C*L)} = \sqrt{(K/M)}$

The effective mass of a quartz resonator with electrode metallization on opposing faces of the narrowest dimension is proportional to the product of the electrode area and the electrode spacing (the narrowest dimension or thickness). Then:

$M \sim A*T$

where A is the electrode area and T is the thickness (*Fig. 3*). The spring modulus of the same resonator is proportional to the product of the electrode area and the inverse of the thickness.

---

3. In a cubic quartz resonator, the electrodes are usually positioned across the narrowest dimension.

$K \sim A/T$

Therefore, the natural frequency of the mechanical system is independent of electrode area and proportional to the inverse of the thickness:

$\omega_0 = \sqrt{(K/M)} \sim \sqrt{(A/(T*A*T)} = \sqrt{(1/T^2)} = 1/T$

There are many ways to cut crystal resonators from a piece of quartz. AT-cut crystals are popular for their good temperature-coefficient characteristics and consistency from one sample to another. For AT-cut crystals, the mechanical resonance is a shear mode. In this mode of operation, the center of gravity moves both vertically and horizontally. Thus, the preceding analysis is a one-dimensional approximation, useful for qualitative understanding of the mechanical resonance of an AT-cut crystal.

From a parallel circuit perspective, the overall electrical impedance of the crystal will be inversely proportional to the electrode area, as a larger electrode area is equivalent to multiple smaller electrode area crystals in parallel. Thus, the motional capacitance and the parallel-plate portion of the shunt capacitance will be *directly* proportional to the electrode area, and the series resistance and motional inductance will be *inversely* proportional. The shunt capacitance and motional capacitance have a linear relationship, as they are both proportional to the electrode area for the unpackaged crystal (the "blank"). The relationship would be *strictly* proportional if the shunt capacitance's fringing fields and the package's parasitic shunt capacitance were negligible.

The preceding analysis reveals three design tradeoffs. First, a smaller electrode area can reduce the crystal's cost and the package size. However, reduced area increases series

resistance, which slows startup time and can prevent oscillation. Second, a larger electrode area lowers series resistance, but also increases shunt capacitance. This, in turn, lowers the active circuit negative resistance, which also slows startup time and can prevent oscillation. And third, increased electrode area increases motional capacitance. This causes greater sensitivity to frequency shift ("pulling") from external capacitive loads.

Load Capacitance

Many crystal oscillators operate at the parallel resonance of the crystal and the load capacitance rather than the series resonance of the crystal's LCR model. The load capacitance is defined as the effective capacitance, external to the crystal package, appearing between the crystal's terminals *(Fig. 4)*. Crystal manufacturers therefore specify a load capacitance to set the crystal's intended operating frequency. A different load capacitance will produce the wrong frequency.

4. The load capacitance is the
capacitance "seen" across the crystal's terminals, exclusive of the internal shunt capacitance.

This can be demonstrated by combining the shunt and load capacitances in parallel, then combining this shunt, plus the load capacitance, in series with the motional capacitance, to form the overall effective capacitance.

$$C_{EFF} = C_{MOTIONAL} * (C_{LOAD} + C_{SHUNT})/(C_{LOAD} + C_{SHUNT} + C_{MOTIONAL})$$

The change in effective capacitance is slight because the motional capacitance is generally about three orders of magnitude smaller than the shunt and load capacitances. Therefore, $(C_{LOAD} + C_{SHUNT})/(C_{LOAD} + C_{SHUNT} + C_{MOTIONAL})$ is nearly unity, and the effective capacitance is very near the value of the motional capacitance.

As the load capacitance increases, $(C_{LOAD} + C_{SHUNT})/(C_{LOAD} + C_{SHUNT} + C_{MOTIONAL})$ comes even closer to unity, so the effect of load capacitance on the net capacitance grows even smaller, and frequency pulling is reduced. In the same way, smaller motional capacitances also reduce frequency pulling as $(C_{LOAD} + C_{SHUNT})/(C_{LOAD} + C_{SHUNT} + C_{MOTIONAL})$ comes closer to unity for any given load capacitance.

Negative Resistance

Pierce and Colpitts oscillators are known as "three-point oscillators." The three points A, B, and C are identical for both topologies, except for the ac-ground point *(Fig. 5)*. To determine the impedance presented to the crystal by a transconductor (usually a MOSFET or a bipolar junction transistor, but in some cases a JFET or even a vacuum tube), C2, and C3, we can replace the crystal with a current source that drives current from point A to point C in the Pierce oscillator equivalent circuit *(Fig. 6a)*. From this:

5. The small-signal model of a Colpitts oscillator with only those elements that determine crystal loading can be used to demonstrate the equivalence (other than the ac ground) with the Pierce oscillator as a three-point oscillator. (a). The small-signal model of a Pierce oscillator with only those elements that determine the crystal loading can be used for demonstrating the equivalence (other than the ac ground) with the Colpitts oscillator as a three-point oscillator. (b).

---

6. The small signal-model of a Pierce oscillator can be used for calculating the impedance presented to the crystal by the three-point oscillator so the negative resistance and load capacitance characteristics of the oscillator may be determined (a). The simplified equivalent circuit of a three-point crystal oscillator and crystal includes the negative resistance and load capacitance applied to the crystal used for deriving the negative resistance applied directly to the crystal's motional arm. This is accomplished by taking the shunt capacitance into account in conjunction with the three-point oscillator. (b).

$V_A = -Z3*I$

where:

$Z3 = 1/(j*\omega*C3)$

$V_C = Z2*I - Z2*gM*V_A = Z2*I + Z2*gM*Z3*I = I*(Z2 + gM*Z3*Z2)$

and gM is the small signal change in collector current per change in base to emitter voltage for a bipolar junction transistor ($gM = \Delta I_C/\Delta V_{BE}$), or the small signal change in drain current per change in gate to source voltage for a MOSFET ($gM = \Delta I_D/\Delta V_{GS}$).

where:

$Z2 = 1/(j*\omega*C2)$.

$V_{CA} = V_C - V_A = I*(Z3 + Z2 + gM*Z3*Z2)$

$Z_{IN} = V_{CA}/I = Z3 + Z2 + gM/(C3*C2*(j* \omega)^2) = Z3 + Z2 - gM/(C3*C2*\omega^2)$

Z_{IN} is the impedance presented to the crystal by C2, C3, and the transistor. It's the effective series combination of the capacitors in series with a negative resistance. Choosing an appropriate C2 and C3, independent of the transistor's gain, can set the crystal's load capacitance.

This analysis suggests that an arbitrary negative resistance can be synthesized for a three-point oscillator using the appropriate transfer characteristic and capacitor values. This is possible in the absence of stray capacitance between nodes A and C. In reality, there will always be some stray capacitance between nodes A and C. More importantly, the crystal's shunt capacitance will always reduce the effective negative resistance presented to the LCR motional branch *(Fig. 6b)*.

To assess the effects of the crystal shunt capacitance on a three-point oscillator, let's look again at the equation for the input impedance of a three-point oscillator:

$Z_{IN} = Z3 + Z2 + gM*Z3*Z2$

Placing this impedance (Z_{IN}) in parallel with C_{SHUNT}:

$Z_{APPLIED} = [1/Z_{SHUNT} + 1/(Z3 + Z2 + gM*Z3*Z2)]^{-1}$

$Z_{APPLIED} = [(Z3 + Z2 +Z_{SHUNT} + gM*Z3*Z2)/(Z3*Z_{SHUNT} + Z2*Z_{SHUNT} + gM*Z3*Z2*Z_{SHUNT})]^{-1}$

$Z_{APPLIED} = (Z3*Z_{SHUNT} + Z2*Z_{SHUNT} + gM*Z3*Z2*Z_{SHUNT})/(Z3 + Z2 +Z_{SHUNT} +gM*Z3*Z2)$

Replacing the generic impedances with capacitive impedances, and taking the real part of $Z_{APPLIED}$, the negative impedance presented by the three-point oscillator to the LCR motional branch is:

$Re\{Z_{APPLIED}\} = - (gM*C3*C2) / [\omega^2 *(C3*C2 + C3*C_{SHUNT}+ C2* C_{SHUNT})^2 + (gM*C_{SHUNT})^2]$

Taking the derivative of $Re\{Z_{APPLIED}\}$ with respect to gM and setting the derivative equal to zero yields the transconductance gM(MIN)R at which the minimum (largest magnitude) negative resistance occurs:

$gM(MIN)R = \omega * [(C2*C3) / C_{SHUNT} + C2 + C3]$

Substituting gM(MIN)R yields:

$Re\{Z_{APPLIED}\}|_{MIN} = -1/\{2*\omega*C_{SHUNT}*[1 + C_{SHUNT}*(C2 + C3) / (C2*C3)]\}$

The negative resistance, $Re\{Z_{APPLIED}\}$, is always negative. The absolute value of the negative resistance drops as C_{SHUNT} increases (*Fig. 7*). The maximum achievable absolute value of the negative resistance (at gM(MIN)R) drops as C_{SHUNT} increases. The absolute value of the negative resistance must be larger than the motional resistance of the crystal for oscillation to occur. Generally, the typical or nominal absolute value of the negative resistance should be greater than four times the motional resistance.

---

7. A plot of negative resistance versus capacitance at 10 MHz, with a transconductance of 5 mA/V (5000 µmhos), illustrates the need to use the optimal capacitance for a fixed transconductance to achieve maximum negative resistance magnitude. The capacitance is the series combination of C2 and C3 (a). A plot of negative resistance versus transconductance at 10 MHz, with a capacitance of 10 pF, illustrates the need to use the optimal transconductance for a fixed capacitance to achieve maximum negative resistance magnitude. The capacitance is the series combination of C2 and C3 (20 pF each) (b).

Note C_{SHUNT}'s strong influence on both plots. Even a small increase in C_{SHUNT} decreases the magnitude of the negative resistance in every configuration, especially near the negative resistance's peak magnitude. To be able to use the recommended load capacitance, and retain a higher magnitude of negative resistance, C_{SHUNT} must be small. C2 and C3 are then adjusted to provide the appropriate load.

Consider two cases where the operating frequency is 10 MHz, the crystal C_{SHUNT} is 2 pF, the crystal load capacitance is 8 pF, the IC package and printed-circuit board (PCB) stray capacitances C2 and C3 are 8 pF, and the transconductance is fixed at 1 mA/V (1000 µmhos):

> • Case 1: Use 8-pF ceramic capacitors at the C2 and C3 positions to load the crystal. These capacitors are in parallel with the 8-pF stray capacitances, for total C2 and C3 values of 16 pF. The net crystal load will be 8 pF, as C2 and C3 are in series with respect to the crystal. The negative resistance calculated from the preceding equation for Re{$Z_{APPLIED}$} will be –627 â""¦.
>
> • Case 2: Use a 4-pF ceramic capacitor in parallel with the crystal to save the cost of one surface-mount device (SMD) capacitor and placement. The C2 and C3 stray capacitances of 8 pF each load the crystal with 4 pF. The additional 4 pF of shunt capacitance in parallel sum to a total of 8-pF load capacitance. However, in this case the negative resistance will only be –466 â""¦, due to increasing C_{SHUNT}.

Case 1 is the preferred approach because of its higher absolute value of negative resistance.

Startup Time

The definition for the startup time of a crystal oscillator can vary, depending on the type of system. For a microprocessor system, the startup time is often the time from initial power application to the time a stable clock signal is available. The startup time for a phase-locked loop (PLL) is often the time from initial power application to when a stable reference signal is available, often settled to within an acceptable frequency offset from the final steady state oscillation frequency. The startup time of a crystal oscillator is determined by the noise or transient conditions at turn-on; small-signal envelope expansion due to negative resistance; and large-signal amplitude limiting.

Envelope expansion is a function only of total negative resistance and the motional inductance of the crystal. The simplified equivalent series LCR circuit contains the motional inductance, the sum of the applied negative resistance of the three-point oscillator and the motional resistance of the crystal, and the effective series capacitance of the entire network (dominated by the motional capacitance). The undriven network can be modeled with:

$s*L + R + 1/(s*C) = 0$

Multiplying both sides by s/L gives:

$s^2 + s*(R/L) + 1/(L*C) = 0$

The roots are located at:

$$\tfrac{1}{2} * \{-R/L \pm \sqrt{[(R/L)^2 - 4/(L*C)]}\}$$

Because the R/L term within the square root is swamped by the 1/(L*C) term, the equation can be simplified to

$$-R/(2*L) \pm -j * \sqrt{[1/(L*C)]}$$

Because the net resistance R is negative, the poles are in the right half-plane. The resulting time-domain solution for this differential equation is:

$$V(t) = K * [e^{[(R/2*L)]*t}] * \sin\{2\Pi t\sqrt{[1/(L*C)]} + \Theta\}$$

where K is a constant and Θ is an arbitrary phase, both related to the initial startup conditions. The exponential expansion is valid only for small-signal conditions, as the power available to the circuit is limited.

The time constant for envelope expansion is positive. It is directly proportional to the net negative resistance of the three-point oscillator and the motional resistance, and inversely proportional to the motional inductance. Due to the large motional inductance of crystals and the limited net negative resistance, crystal oscillators have long startup times.

As an example of the envelope expansion time constant, assume a crystal with 5-fF motional capacitance, and an oscillator with 1500-Ω negative resistance, operating at 10 MHz. From the motional capacitance and operating frequency, a motional inductance of 50.7 mH can be calculated using $L = 1/(C*\omega^2)$. This motional inductance yields an envelope expansion time constant of $t = 2*L/|R| = 67.55$ µs.

There is a tradeoff between reduced frequency pulling (due to low motional capacitance) and longer startup times (due to high motional inductance), because high motional inductance is a result of low motional capacitance. (The product of the two is more or less fixed.) This interdependence is partly mitigated by the fact that smaller motional capacitances are associated with smaller shunt capacitances, which yield larger negative resistances and thereby improve startup time.

Startup time is an important consideration in battery-powered applications where the device is often switched on and off. Shorter oscillator startup times reduce wasted energy in low-power radio systems (such as those using the MAX7032 transceiver, and the MAX1472 or MAX7058 transmitter).

Frequency Stability Versus Temperature

All crystals vary in resonant frequency as the temperature changes. The way the frequency varies depends on the angle at which the crystal is cut. The relative frequency shift versus temperature of AT-cut quartz crystals can be represented as a cubic polynomial.

$$\Delta f/f_0 = A_0 + A_1(T - T_0) + A_2(T - T_0)^2 + A_3(T - T_0)^3$$

where the coefficients A_0 through A_3 are functions of the cutting angle.

Frequency stability is important in systems using a crystal as the frequency reference. This is especially true for high-frequency narrow-band applications. Consider a 25 kHz channel in the 863- to 870-MHz European ISM band. A frequency drift of only 5 kHz (less than 6 ppm) could result in communication failure or regulatory noncompliance.

Keeping the drift to substantially less than 6 ppm over the "industrial-spec" temperature range of –40°C to 85°C is essentially impossible, even with zero-tolerance crystals that are perfectly cut and don't age *(Fig. 8)*. In this case, a radio system with an internal temperature sensor and narrow-frequency-step fractional-N synthesizer, such as the MAX7049 transmitter, can be used to compensate for known temperature shifts.

8. This plot shows relative frequency shift versus temperature for AT-cut crystal angles (in minutes).

Aging

A crystal's series-resonance frequency gradually changes. A shift of a few parts per million over several years is common. Most of the change occurs during the first year or two. Aging occurs more rapidly at higher temperatures and oscillation amplitudes. A principal cause of aging is an increase in the crystal's mass. The added mass usually comes from contaminants within the crystal case that land on and stick to the crystal's surface.

Frequency Error

There are four principal sources of frequency error:

- Initial tolerance, or the manufacturer's tolerance at 25°C with the specified load capacitance
- Changes in frequency with temperature change
- "Pulling" from load-capacitance variations
- Aging

Drive-Level Dependency

After a period of inactivity (hours to weeks), a crystal's series resistance can rise well above its maximum specified value, becoming a function of the ac electrical drive level. This effect is known as drive-level dependency (DLD) or "sleepy crystals." If the crystal is mechanically shocked or driven "hard" electrically, the series resistance will often drop back to within the specified limits.

DLD is thought to be caused by additional mechanical losses from contamination within the crystal package. The contamination can be solid particulate or moisture, most often water that condenses or freezes on the crystal. Once vibration removes the contamination, normal series

resistance returns. The contamination might not settle again (or at least not to the same degree), resulting in erratic and unpredictable changes in series resistance.

Irreversible DLD can occur when a particle becomes permanently stuck to the crystal, electrode plating cracks or comes loose, or the crystal is scratched. No crystal is completely free of DLD, but higher quality products exhibit much lower DLD, both in degree of series resistance increase and the percentage of units exhibiting resistance changes.

To reduce problems associated with DLD, operate with a large negative resistance, greater than four times the manufacturer's maximum specified series resistance. This eliminates nearly all DLD issues. You also could purchase from vendors who make high-quality crystals. Or, pay the premium for crystals that have been factory tested for low DLD.

Spurious Modes

Crystals often have undesired mechanical resonances near the fundamental frequency. These "spurious modes" can be modeled as series LCR branches in parallel with the fundamental-frequency branch. Spurious modes usually have higher losses than the desired mode and are thus less likely to oscillate strongly. (Crystal manufacturers usually test for spurious resonances and don't ship units with low losses at those frequencies.)

Oscillators with large negative resistances usually limit or clip throughout most of the oscillation cycle. During limiting, the circuit's effective gain is essentially zero (i.e., a change in input produces little or no change in output). Unless a spurious mode has low losses, or the active circuitry doesn't strongly limit the signal amplitude, there won't be enough gain for spurious modes to oscillate. They will be "choked" by the desired large-signal oscillation. If there is enough gain for these undesired secondary oscillations, these frequencies can wreak havoc with phase-frequency detectors in PLLs and other circuits.

References

1. Vittoz, Eric A., Degrauwe, Marc G. R., and Bitz, Serge, "High-Performance Crystal Oscillator Circuits: Theory and Application," IEEE Journal of Solid-State Circuits, Vol. 23, No. 3, June 1988.
2. Kreyszig, Erwin, Advanced Engineering Mathematics, Fifth Edition, John Wiley and Sons, 1983.
3. Bechmann, Rudolf, "Frequency-Temperature-Angle Characteristics of AT-type Resonators Made of Natural and Synthetic Quartz," Proceedings of the IRE, November 1956.