

Crystal Oscillator Design with CMOS Tutorial

Bob Hay May 2008

Introduction

This tutorial discusses the design of an oscillator comprised of a crystal resonator, a few passive components, and a CMOS gain block. It begins with a discussion of oscillator concepts and Q followed by a review of the general model of the crystal resonator, including discussion and simulation of both the resonant and anti-resonant operating modes of operation. This is followed by a discussion and simulations of the gain block as well as simulations of a system comprising the gain block and the resonator in the anti-resonant mode in a Pierce configuration.

What is a frequency resonator?

One of the first frequency resonators was a pendulum. A stable pendulum is one that has large mass and very good bearings which generate minimal loss and noise. It would ideally operate in a vacuum to minimize loss (as does a good crystal). One of many definitions of the quality factor Q is

$$Q = 2\pi * \frac{\text{Energy Stored}}{\text{Energy dissipated per cycle}} \quad (1)$$

Clearly in the case of our pendulum the value of Q is increased with more mass (more energy stored) and lower loss. Even with a large value of Q it is obviously necessary to supply some energy periodically to maintain a reasonably constant amplitude of oscillation. Since this additional energy is in the form of an occasional push (probably once each cycle) and the pushing mechanism typically never provides the push at exactly the ideal moment in time, we can achieve the maximum stability with a very weak pusher and a correspondingly high Q . We might also note that the time required to get a system started from standstill will increase as the Q increases due to the increased mass and the weaker pusher that injects energy at a lower rate.

This pendulum example has many parallels to the crystal oscillator and illustrates the importance of Q and its relation to stability.

Why use a crystal oscillator?

In many applications the performance of a circuit can fail to meet its objectives if the reference frequency is not sufficiently stable. In digital circuits, problems arise when variation in spacing between clock transitions (AKA clock jitter) becomes great enough that some clock transitions occur prior to the arrival of accurate data and setup time requirements are not met. In communication circuits, problems occur when phase jitter (AKA phase noise) exacerbates adjacent channel interference and causes errors in phase modulated/demodulated baseband signals. In both cases low-rate noise (AKA frequency

drift, generally caused by aging or temperature variation) can cause all the problems that would be expected when the frequency does not have the desired value.

The general solution to frequency drift, time-domain jitter, and frequency-domain noise is to use a more stable oscillator (or reference clock). For many applications the best choice of frequency controlling element is a crystal resonator. These resonators are generally available at reasonable cost, are made with a wide variety of frequency values, and can achieve very good frequency stability due to their high available Q and relatively low temperature sensitivity.

The crystal resonator

A crystal resonator is typically made of a thin quartz crystal disk with gold electrodes plated on its opposite sides as shown in Figure 1. The typical equivalent circuit is shown in Figure 2 where the motional elements R_s , L_s , and C_s are the result of the electromechanical parameters of the piezoelectric quartz blank and C_p is the electrical capacitance between the gold electrode pair combined with the capacitance of the base pins. The significant benefit of crystal resonators derives from the extremely high value of $\frac{X_{L_s}}{R_s}$, a value that would not be achievable using an actual inductor.




Figure 1 Typical quartz crystal resonator




Figure 2 Equivalent circuit

When operating in the series resonant (or “resonant”) mode the values of the resonant frequency and the Q can be determined by the equations (2) and (3).

$$\omega_s = \frac{1}{\sqrt{L_s C_s}} \quad (2)$$

$$Q = \frac{1}{R_S} \sqrt{\frac{L_S}{C_S}} \quad (3)$$

Typical values for R_S vary with frequency and the specific crystal but tend to follow the trends shown in Figure 3.




Figure 3 Typical R_S vs. fundamental frequency

For the purposes of this paper we begin with the simulation of a high quality 10MHz crystal. We shall assume it has a typical Q of 10^6 a typical R_S as indicated in Figure 3 of 65Ω , and a typical value for C_P of 3pF. Using these values and rearranging equations (2) and (3) we can generate the values for C_S and L_S shown in the simulation file:

[*1_CrystalModel_10MHz_0Z load.asc*](#)

where we drive the crystal with an ideal zero-impedance voltage source and load it with an ideal zero-impedance load through which we observe the current.

With these values of source and load impedance the capacitor C_P plays no role and we observe only the effect of the series resonant motional components.

The resulting plot shows the load current which indicates a gain due to R_S of

$20 \log_{10} \left(\frac{1}{65} \right) = -36.25dB$ and a phase shift of 0° at resonance. Also plotted is the group delay τ_G and from the relationship

$$\tau_G = \frac{d\phi}{d\omega} \quad (4)$$

and the relationship:

$$Q = \frac{\omega_0}{2} \frac{d\phi}{d\omega} \quad (5)$$

where ω_0 is the resonant frequency and where noting the group delay of 32mS from the simulation we can calculate

$$Q = \pi * 10^7 * 32 * 10^{-3} \approx 10^6 \quad (6)$$

which is the value of Q used to generate the values for L_S and C_S in the simulation. It is convenient to be able to use group delay to determine the loaded Q of a resonator. The same value of Q can also be derived by determining the 3dB bandwidth ω_{3dB} of the amplitude response and using the equation

$$Q = \frac{\omega_0}{\omega_{3dB}} \quad (7)$$

Creating a crystal simulation subcircuit

In order to simplify the schematics we can create a crystal simulation subcircuit. This can be done either by incorporating the four components shown in Figure 2 in to a subcircuit or, as suggested in the LTSpice manual, by using the features of the capacitor model shown in Figure 4 to include all of the components.




Figure 4 LTSpice capacitor model

It is also convenient to have the crystal subcircuit compute the values of the motional inductor and capacitor from a desired value of Q , f_s , and R_s as shown in [2_xtal_param.asc](#).

An example of how this circuit is used by passing in the appropriate parameters is shown in the file [3_Xtal Symbol and params.asc](#).

We show the implementation of this symbolic circuit in

[4_CrystalModel_10MHz_0Z_load_symbol_param.asc](#)

where the results should be identical to those previously found in

[1_CrystalModel_10MHz_0Z_load.asc](#)

Creating an ideal series-resonant crystal oscillator

As with the pendulum example, it is necessary to add energy at the resonant frequency to sustain oscillation in any resonator. This added energy is generally provided to a crystal resonator by adding a gain stage shown in the open loop configuration of Figure 5.

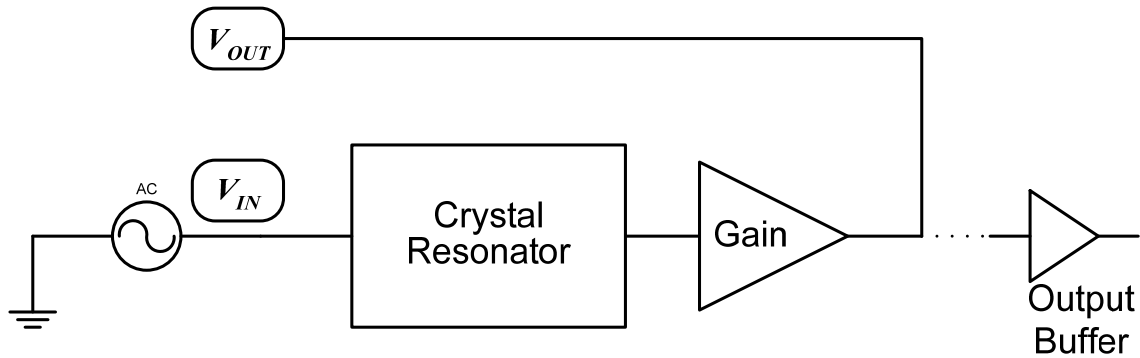


Figure 5 Open loop response

The output buffer shown is generally added with the objective of isolating the oscillator from variations in load impedance that could modulate the frequency of the oscillator through a process generally referred to as “frequency pulling.”

We define the terms

$$G = \left| \frac{V_{OUT}}{V_{IN}} \right| \quad (8)$$

and

$$\phi = \angle V_{OUT} - \angle V_{IN} \quad (9)$$

According to the Barkhausen criteria for oscillation, when $G > 1$ and $\phi = 2N\pi$ where N is an integer, if we were to remove the AC stimulus and connect the V_{OUT} node to the V_{IN} node, thus closing the loop, the system should sustain oscillation since the new V_{IN} is similar to but greater than the V_{IN} provided by the AC stimulus.

We can see from the simulation file

[5_CrystalModel_10MHz_0Z_ICVS_symbol.asc](#)

using ideal components that we can meet these criteria if we make the gain of the B1 stage sufficiently large to overcome the loss of the R_S element of the resonator.

We can then close the loop as shown in Figure 6 where the gain block has gain sufficient to exceed the Barkhausen criteria.




Figure 6 Closed loop oscillator

This simple system using an ideal zero input and output impedance gain block is simulated in

[6_CrystalModel_Oscillatorl_10MHz_0Z_ICVS.asc](#)

There are several important items to note in this simulation.

- Because the ideal gain components have constant gain at essentially infinite frequency, it was necessary to run this simulation with no parallel capacitance C_p in the crystal resonator in order to achieve stability.
- In this simulation the resonator Q was degraded to 10^5 . As mentioned in the pendulum example, startup time (and thus simulation time) increases with increasing Q , so reducing the Q reduces the simulation time.
- It is necessary to add a startup circuit to initiate oscillation. In a real oscillator there is a small amount of noise generated by the active gain circuit. This noise includes sufficient power at the resonant frequency to get the oscillator started. The ideal components used in these simulations contain no noise sources so the additional startup source is necessary.
- The ideal components contain no level-based saturation properties. Extending the simulation time of this circuit will result in exponential growth in the amplitude of the signals.

Operating in the anti-resonant (parallel resonant) mode

The circuits necessary for creating a series resonant oscillator are generally not very practical. Semiconductor devices naturally have relatively high input and output impedances and methods to reduce those impedances using feedback are limited at high frequencies due to gain-bandwidth constraints. Use of L-C based transformers (as seen in the Butler oscillator, for example) can improve this situation but at the cost of including an inductor in the circuit.

One way to improve the match between the natural capabilities of semiconductor devices and the properties of a crystal resonator is to operate the resonator in the parallel resonant mode. In this mode the ideal source and load impedances for maximum Q are infinite, a value which is also impossible to achieve. However, particularly for high frequency designs, the anti-resonant mode is a more commonly used configuration. An additional benefit of the anti-resonant mode operation is that the device biasing tends to be a little easier to manage because of the phase inversion that takes place.

Recall the crystal equivalent circuit of Figure 2. At the series resonant frequency f_s the impedance of the circuit is largely resistive, with a small amount of parallel capacitance. At frequencies slightly above f_s the impedance will be highly inductive. We can add a capacitor as shown in Figure 7 that will appear in parallel with the C_p of the crystal and the combination of the two capacitors will, at some frequency $f_p > f_s$ create a parallel resonance. The stimulus shown in Figure 7 is a high impedance voltage source.




Figure 7 Anti-resonant mode

This circuit is simulated in

[7_Crystal_inParallelResonance.asc](#)

where the stimulus is replaced by a current source, the inductor is replaced by the crystal model, and the anti-resonant frequency f_p is shown to be slightly higher than f_s .

We can slightly modify the presentation of this circuit in a way that creates no change in function as shown in Figure 8 where we rearranged the parallel capacitor into two series capacitors and connected their common node to an arbitrary ground.




Figure 8 Anti-resonant mode with split capacitors

We also show the current in these capacitors indicated by I_1 and I_2 . According to KCL, these two currents as shown must be equal in magnitude and have a phase relationship of 180° . In order for this to be the case, it is also necessary that $V(\text{Pin1})$ and $V(\text{Pin2})$ be of equal magnitude and opposite phase.

This circuit is simulated in

[*8_CrystalSymbolInParallelResonance2.asc*](#)

where as before the stimulus is replaced by a current source and the actual crystal model is used. The simulation shows as expected the 180° relationship between the two pin voltages and the equal magnitude of the voltages.

Implementation with a real gain stage.

We now look at approaches to implementing these circuits with real gain stages. In Figure 9 we show a typical implementation where the crystal is operating in the anti-resonant mode and the gain stage provides the appropriate phase inversion. Note the presence of the two series-connected capacitors that appear in parallel with the crystal. Also note that the crystal/capacitor node pair are connected to the relatively high impedance nodes (collector and base) of the active device. Of course, the NPN BJT could be replaced with alternative devices including a PNP BJT or MOS or JFET devices with appropriate bias configurations.




Figure 9 Generic gain stage

There are three standard topologies for creating anti-resonant oscillators shown in Figure 10, where the only difference is the placement of the ground node.

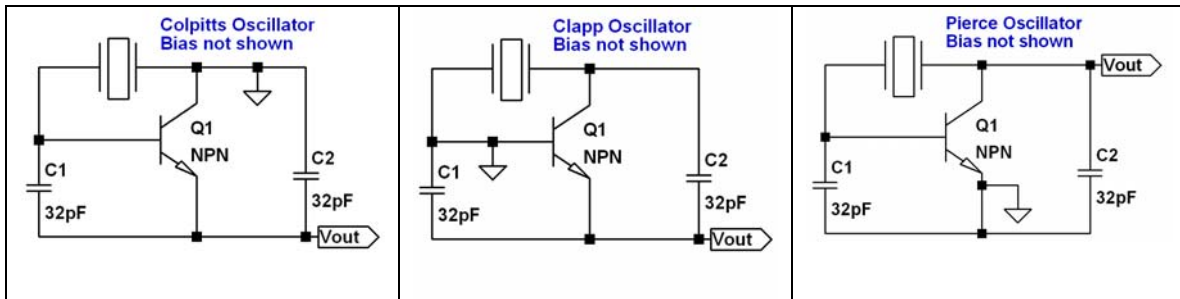


Figure 10 Alternative topologies for crystal oscillator gain stages

Since our objective is to create an oscillator incorporating a CMOS device, the obvious selection is the Pierce topology, since a CMOS inverter most resembles a common emitter (or source) active device.

In order to create the required inversion for the gain stage we will need an odd number of inverters in the active device. The choice of number of stages is based on several factors including

- The required gain
- The input loading
- The bias current

For a high Q crystal the gain requirements are not extremely high so we might have sufficient gain with a single stage but answering that question will require more exploration. Input resistive load of the gain stage can cause degradation of loaded Q but since the input load for a CMOS device is largely capacitive at moderate frequencies that will also not be a primary consideration.

Bias current is a major consideration in many applications. The input of the active gain block will typically be biased at the switching point and have very little AC component and thus a relatively high bias current. This stage should therefore have minimal width if bias current is an important consideration. Conversely, if noise is an important consideration, it may be desirable to increase the width of the input devices.

For our initial approach to this simulation we choose to use a 3-stage CMOS device with minimum-width first stage devices followed by two other stages, each with devices having a width of 3 times the width of the preceding devices. This provides for nearly minimum propagation delay and operation of the wider devices further from the switching point. We can

bias the input through a low-pass filter connected to the output of the inverter. The filter raises the AC input impedance since it isolates the inverted output from the bias, thus increasing the input resistance.

The basic inverter gain block is as shown in Figure 11 where reasonably small widths are shown in the parameter statement. Modifying this statement will cause all devices to scale appropriately.

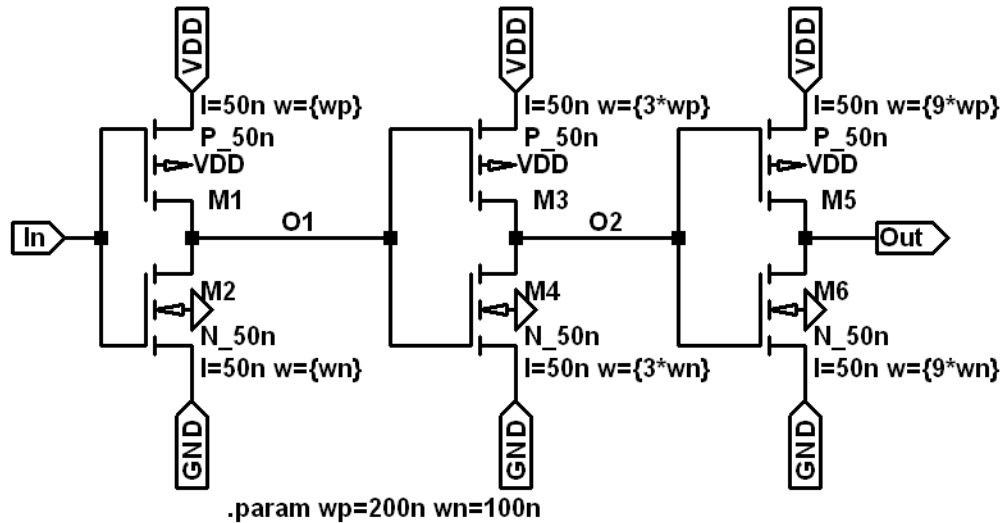


Figure 11 Basic inverter gain block

The gain block simulation with appropriate biasing is available in the file

[09a_Inv3.asc](#)

which shows approximately 9° of excess phase shift through the inverter and a gain of about 21dB per stage.

It is also useful to characterize the output impedance of the inverter, which we do with the simulation in

[09b_Inv3.asc](#)

where we see that the output impedance at 10MHz is approximately $3.76K\Omega$ and from the angle it appears to be nearly all resistive.

This inverter design is incorporated with a crystal and associated components, and the AC simulation is shown in

[10_CrystalModel_10MHz_Network_Inverter3_Qparam3_xtal.asc](#)

Note that the series resonant frequency f_s is shifted downward slightly to compensate for the fact that we are operating the crystal in the inductive mode. Also note that we use the output impedance of the inverter as determined from the previous simulation.

Here we can determine from the group delay plot that our Q_L , the loaded Q , has deteriorated to approximately 600k based on Eq. (5) combined with the observation that the group delay τ_G from this simulation is approximately 19ms. This degradation is due almost entirely to the low output impedance of the inverter. We could increase the Q_L by inserting a series resistor between the inverter and the crystal but the downside would be less power in the crystal and therefore a poorer SNR at the input of the inverter. Another option might be to redesign the output stage with either smaller width and/or greater length to increase the output impedance while attempting to retain sufficient gain.

If we are going to explore the possibility of a smaller output stage, we might begin with exploring a single stage inverter.

We can see from

[*11a_Inv1.asc*](#)

that the simulated gain is about 20dB, the excess phase shift is less than 2° , and from

[*11b_Inv1.asc*](#)

that the output impedance is approximately 77K Ω . Unfortunately, when we include the crystal resonator in the AC simulation in

[*12_CrystalModel_10MHz_Network_Inverter1_Qparam3_xtal.asc*](#)

we find that the gain is insufficient to sustain oscillation.

However, the higher output impedance remains attractive for a higher Q_L so our next design iteration is to return to the 3-stage design for higher gain but incorporate an output device with smaller channel width and/or greater channel length. The example chosen uses the 50nm process but the output device pair has width of 200/100 nm and length of 100nm. The simulation at

[*13a_Inv3s.asc*](#)

indicates that the gain is about 67dB, the excess phase shift has increased to about 12° , and from

[*13b_Inv3s.asc*](#)

we see that the output impedance has risen to about 129K Ω . We can check the AC simulation of this gain block with our crystal in the circuit in

[*14_CrystalModel_10MHz_Network_Inverter3S_Qparam3_xtal.asc*](#)

where we find an open loop gain of about 34dB and a group delay τ_G of about 30ms and from Eq. (5) we find this corresponds to a Q_L of about 940K. This is a significant improvement in the Q_L compared to the previous 3-stage gain block. Note that we must use the group delay at the point where the phase crosses the zero line and not at its maximum point. This is a good example of how excess phase shift in the gain block causes Q_L deterioration.

The final test is a transient simulation of the oscillator which is done in

15_CrystalModel_10MHz_Network_Inverter3S_Qparam4_transient_xtal.asc

This simulation requires several hours to run to completion. In order to reduce the simulation time to that level the Q of the resonator has been decreased to 100K and the startup circuit delivers a huge 200 amps peak-peak to the resonator (can only be done in a simulator!) for 20 cycles. The purpose of the simulation is first to determine whether there is sufficient gain to sustain the amplitude of the oscillation at some level and if so to determine what that level is. Due to non-linearity of the gain block it is normal for the level to be somewhat less than the maximum possible excursion of the unloaded amplifier. The change in amplitude takes place very slowly so the simulation time is fairly long. In this simulation the amplitude at both Pin1 and Pin2 appears to stabilize with excursions between 426mV and 540mV for an amplitude of 114mV P-P. Another parameter of interest from the transient simulation is the current consumption of the gain block. In this case the circuit is very efficient, with an average current consumption of 14.5 μ A at 1V for a power consumption of about 14.5 μ W. The average power dissipated in the crystal is 675nW. The schematic for the gain block is shown in Figure 12.

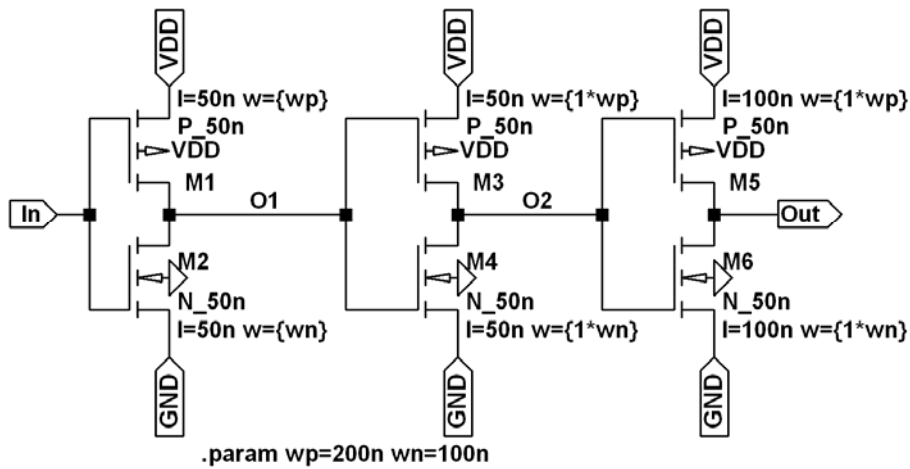


Figure 12 Gain block for high Q performance

Conclusion.

LTSpice appears to be a useful tool for simulating crystal oscillator performance. One observation is that the AC simulations seem to suggest greater open loop gain than the transient simulations suggest, even considering the effects of nonlinearities of the gain block.

It is also apparent that the design of a crystal oscillator system including the gain block requires a rather iterative process to optimize the loaded Q, the crystal power, and the system noise.

For applications where noise is of concern, it would likely be preferable to use a CMOS process that could support higher voltages. This would enable more power in the crystal resonator which, while causing some increase in aging rate, would likely improve the SNR at the input to the gain block. Exploring gain blocks using the 1 micron process would be a good next step.