Ron ZL1TW's Crystal Oscillator

Gary ZL1AN, `g.bold@auckland.ac.nz`

Background

The single-tube QRP transmitter constructed by Ron ZL1TW (now a silent key) in my January/February *Morseman* column is shown in figure 1.

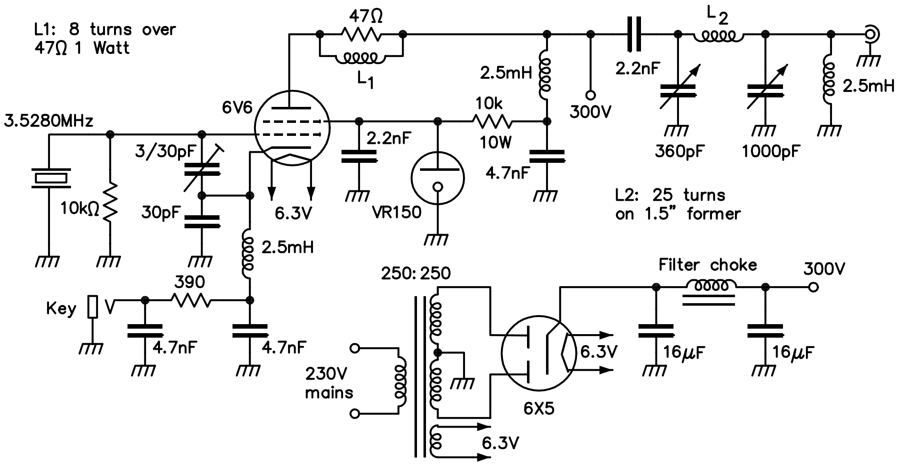

Figure 1: The QRP transmitter of Ron, ZL1TW.

I received this message from an overseas reader:

"In the *Morseman* column of January/February 2008, Ron, ZL1TW, describes his simple crystal controlled transmitter with one valve 6V6. Ron complains that the output is only 1.5 W.

"To me the reason is clear. The cathode of the 6V6 is connected to common via a capacitor of 30pF. At 3.5MHz this capacitor has a reactance of 1514 ohm. Such a high reactance causes heavy negative feedback, decreasing the amplification of the 6V6 to a small value. It would be better to delete the capacitor of 30pF, the trimmer of 3/30pF and the 2.5mH choke. If the circuit won't oscillate a small capacitor between anode and control grid of the 6V6 will help."

This indicates a misunderstanding of how the circuit works. In fact it seems that many others are perplexed by it too. This is my attempt to explain it.

Development of the Circuit

Ron's oscillator is similar to many others that have been published. Two other examples are shown in figure 2. You'll see that both use the same combination of two capacitors from grid to ground, and the apparent "bypass" capacitors are also small values. However, these are not bypass capacitors at all, but parts of the feedback network.

This type of crystal oscillator is a derivative of the *Colpitts Oscillator*, invented by Edwin Henry Colpitts (January 19, 1872 - 1949), who was research chief for Western Electric in the early 1900s.

His colleague, Ralph Hartley, had invented an inductive coupled oscillator which Colpitts modified in 1915. Colpitts first reported it in a paper with Edward Craft, in 1919. He patented it in 1920.

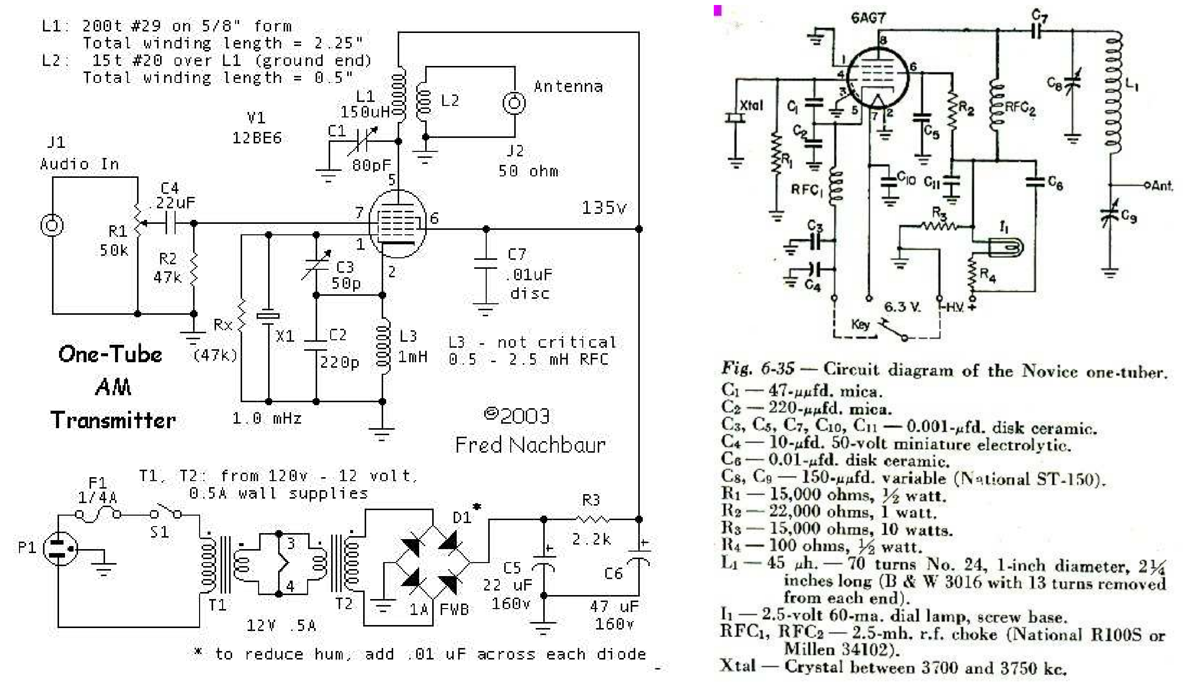

Figure 2: Left: Oscillator published on the web by Fred Nachbaur. Right: One-tube transmitter in 1953 ARRL Handbook

The Standard Colpitts Oscillator.

Figure 3, below, shows the topology of a Colpitts oscillator as now typically drawn, here implemented with an n channel FET. The FET has essentially infinite infinite input (gate-source) resistance, like the pentodes in the circuits above, and acts in the same way. Substituting it doesn't affect the analysis. (I used this diagram because I already include it in my text.)[1] The load, in the drain (the pi or inductively coupled antenna in the plate circuit of the pentodes) is omitted because it doesn't essentially modify the principle of operation.

Most texts and websites showing this circuit omit any analysis. But when drawn like this it is not at all clear what the feedback path is, or how the oscillation conditions can be derived. It's best analysed by changing the common point of the circuit to be the source of the transistor, and drawing the resulting small-signal equivalent circuit. This doesn't change the operation of the circuit.

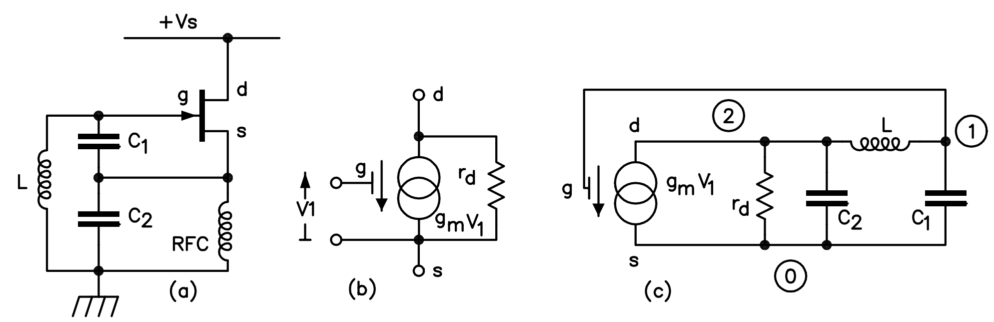

Figure 3: (a) Simplified Colpitts Oscillator. (b) Small-signal equivalent circuit of the transistor. (c) Transformed circuit.

[1]See *"Theoretical and computer analysis of systems and networks"*, Bold and Tan, Chapter 4, page 67, published by the Auckland University Physics Department.

In this simplified circuit with no drain (plate) load, the output is taken from across the RFC. This is adequate to illustrate the oscillation principle.

The small-signal equivalent circuit of the transistor, expressed as a current source, is shown in figure 3 (b), where r_d is the dynamic drain resistance, and g_m is the transconductance. A pentode has the same equivalent circuit.

The transformed circuit is shown in figure 3 (c). The connection from the source to the capacitor's midpoints in (a) is now the common terminal, the RFC is assumed to have infinite impedance and is not shown, and the *dc* power supply (V_s) is a short circuit at the oscillation frequency and is connected to ground.

This version makes a lot more sense, since it shows the transistor energising a *pi* network which can be considered the feedback element, with its output connected back to the gate (input) of the transistor. Since the transistor gives a phase-shift of 180^o, the *pi* section must *also* give a phase-shift of 180^o.

It's worth noting that in early textbooks and papers, both Colpitts and Hartley Oscillators were *always* drawn this way. Figure 4 shows typical diagrams from a text of the 1960's.

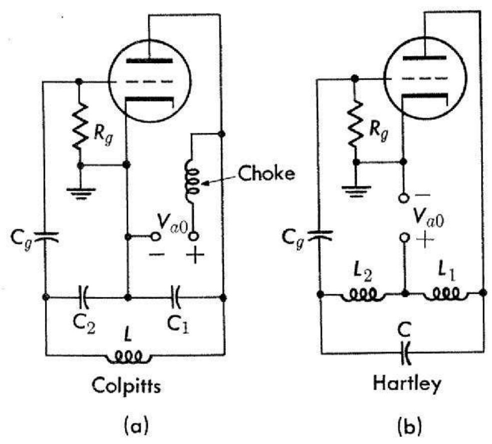

Figure 4: Left: A Colpitts oscillator. Right: A Hartley oscillator.

Again, the only difference is in the common (ground) point. But again, it is clear that the Colpitts feedback network is a low-pass *pi* (drawn upside down), and that of the Hartley is a high-pass *pi*, making circuit operation easier to understand.

It's easily shown using nodal analysis, or otherwise, [2] that the Colpitts oscillation conditions are

$$\text{Transconductance required, } g_m = \frac{C_1}{C_2 r_d} \quad (1)$$

$$\text{Oscillation frequency} = \frac{1}{2\pi\sqrt{LC_0}} \quad (2)$$

$$\text{where } C_0 = \frac{C_1 C_2}{C_1 + C_2} \quad (3)$$

Somewhat more complicated equations result when the inductor is lossy, or when a load is included, and these are given in the footnote reference above. However, they add nothing to understanding the principle.

[2] See again *"Theoretical and computer analysis of systems and networks"*, Bold and Tan, Chapter 4, page 67, where the oscillation conditions are derived by equating the circuit's transient response matrix to zero.

The Clapp Oscillator

The Series-tuned Colpitts or Clapp circuit was invented much later, in 1948, by James K. Clapp. There's a good description on VK2TIP's website at

`http://www.electronics-tutorials.com/oscillators/clapp-oscillators.htm`

The only difference is that an additional capacitor is placed in series with the inductor, as in figure 5 (a). However, the same oscillation condition as for the Colpitts is required, that is, this inductor-capacitor combination must still be *look like an inductor* for oscillation to occur.

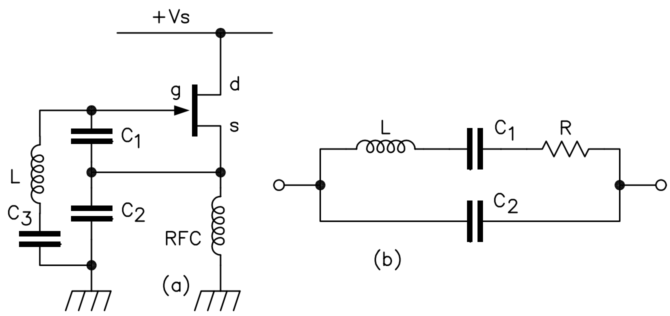

Figure 5: (a) The Clapp oscillator. (b) Equivalent circuit of a crystal.

A series LC circuit looks capacitive *below* the resonant frequency, and inductive *above* it. Hence this circuit will oscillate at a frequency which is *also above* the resonant frequency.

The crystal oscillator we are considering is more correctly a derivative of the Clapp circuit, since the crystal equivalent circuit also has a capacitor in series with the inductor. However, it has another capacitor too, so the equivalence is not exact.

Equivalent Circuit of a Quartz Crystal

Figure 5 (b) shows the standard electrical equivalent circuit of a quartz crystal near its fundamental resonant frequencies - there are two of them, see below.

- The inductor, L represents the vibrating mechanical mass of the crystal.
- The capacitor C_1 represents the compliance of the crystal when it's in oscillatory motion.
- The resistance R represents losses, frictional and molecular heating energy dissipation in the crystal. This, together with the inductance value, determines the Q of the crystal.
- The capacitor C_2 represents the capacitance between the mounting plates of the crystal.

The inductance is typically large, tens of milli-Henry to Henry. The series capacitance C_1 is very small, typically tens or hundreds of femto-Farad (fF). The loss resistance R ranges from a few ohm (20 MHz) to about 200 kΩ (1 kHz crystals).

This results in a *very* high L/C ratio, much higher than achievable with lumped components. Hence the resulting Q, the ratio of reactance to resistance, is *also* very large, typically some tens of thousands. This is one reason for the high stability of crystal oscillators compared to lumped-component oscillators.

Resonant Behaviour of the Crystal

Because the crystal equivalent circuit contains both series and parallel capacitance, it will also exhibit *both* series and parallel resonance. Their frequencies are readily shown to be

$$f_s = \frac{1}{2\pi\sqrt{LC_1}} \qquad \text{(series resonance)} \tag{4}$$

$$f_p = \frac{1}{2\pi\sqrt{LC_p}} \qquad \text{(parallel resonance)} \tag{5}$$

$$\text{where } C_p = \frac{C_1 C_2}{C_1 + C_2} \qquad \text{(Series combination of } C_1 \text{ and } C_2\text{)} \tag{6}$$

The parallel resonant frequency, f_s, is higher than the series resonant frequency, f_s, but typically by much less than 1%, since the ratio of the capacitors is very large.

The complex impedance of the crystal is given by

$$z_s = \frac{\left(sL + \frac{1}{sC_1}\right)\left(\frac{1}{sC_2}\right)}{sL + \frac{1}{sC_p}} \tag{7}$$

$$\text{where } s = j\omega \tag{8}$$

To evaluate this, we need typical parameter values. These are very hard to find. However, Adachi[3] gives these values for a nominal 10 MHz crystal.

$$L = 18.72\,\mathrm{mH} \tag{9}$$

$$C_1 = 13.54\,\mathrm{fF} \tag{10}$$

$$R = 11.9\,\Omega \tag{11}$$

$$C_2 = 2.75\,\mathrm{pF} \tag{12}$$

Note that C_2, the holder capacitance, is about 200 times larger than C_1, the compliance capacitance associated with the oscillating mass of the crystal! For these values, using equations 4 and 5 we find

$$f_s = 9.995728\,\mathrm{MHz} \tag{13}$$

$$f_p = 10.02130\,\mathrm{MHz} \qquad \text{(0.26\% higher)} \tag{14}$$

If such a crystal is used to replace the inductor of figure 3 (a) in a Colpitts/Clapp oscillator, its impedance must *also* look inductive at the oscillation frequency.

This is well understood. For example, the Wikipedia entry on "Colpitts Oscillator" (slightly re-worded for clarity) says that

"The crystal, in combination with the external Colpitts capacitors C_1 and C_2 forms a *pi* network band-pass filter, which provides a 180 degree phase shift from the output to input at approximately the resonant frequency of the crystal. At the frequency of oscillation, the crystal appears inductive; thus it can be considered a large inductor with a high Q. The combination of the phase shift from the pi network and the gain with 180^o phase-shift from the active element results in a positive loop gain, leading to oscillation."

Figure 6(a) shows the magnitude of the crystal impedance, plotted using Adachi's values, in the vicinity of its nominal frequency. Since the range is large, I've plotted this on a logarithmic scale. We

[3] *Analysis of Crystal Oscillator*, Takehiko Adachi, PDF available on the web.

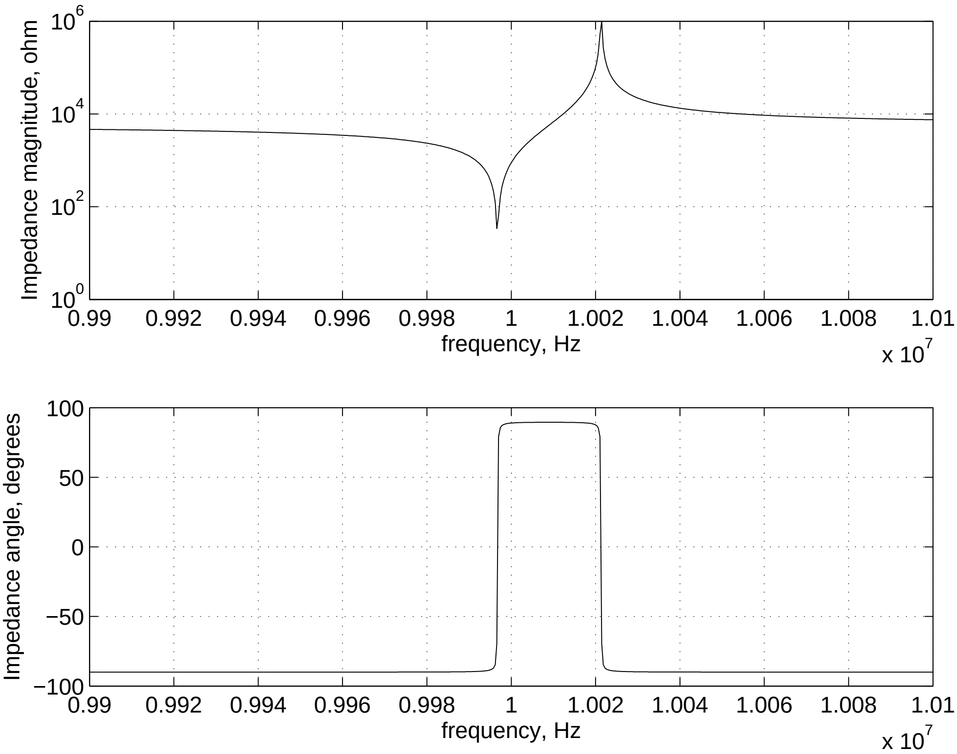

Figure 6: Impedance and phase angle of nominal 10 MHz crystal

see the impedance dip just below 10 MHz at the series resonance (f_s) and rise just above 10 MHz at the parallel resonance (f_p).

Figure 6 (b) shows the *phase angle* of this impedance over the same range. We see that *only between* the resonant frequencies is this positive, signifying inductive impedance! Thus, the oscillation frequency must lie somewhere in this small region, which is only about 0.25% of the nominal crystal frequency in width.

Aside: The reason for the high and low-frequency capacitive reactance is easy to see physically.

- At very low frequencies, the inductor tends to a short circuit, so the equivalent circuit looks like a lossy capacitor (C_1) in parallel with another (C_2). Hence it looks capacitive.
- At very high frequencies, the inductor tends to an open circuit, so no current flows in the upper branch and the circuit again looks capacitive (now C_2 alone)

Crystals used in this circuit are conventionally described as operating in “parallel mode”, which is unfortunate, since the impression is given that the oscillation frequency will be that of the *parallel resonance*. This is not true. The oscillation frequency will always be somewhere in the window *between* the series and parallel resonant frequencies, and always *lower* than the parallel resonant frequency.

Crystals operating in “series mode” require a non-inverting active element. There is no difference in the way crystals operating in these two modes are constructed, only in the circuits using them.

Oscillator SPICE Simulation

To demonstrate the characteristics of this oscillator, it’s instructive to perform a `SPICE` simulation, in “transient” mode. Here the circuit is defined and left to evolve in behaviour by itself. This also allows us to demonstrate several other difficulties involved in `SPICE` oscillator simulation.

Figure 7 shows the schematic of a Colpitts Crystal oscillator set up for simulation with `LTSpice`[4]

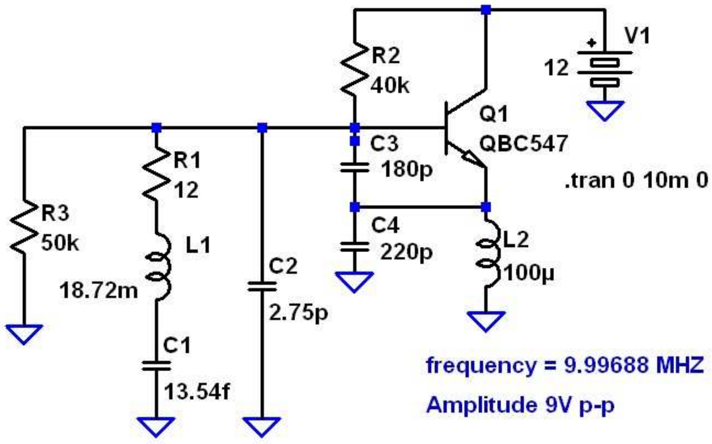

Figure 7: Colpitts Crystal Oscillator Circuit

The crystal is represented by components L_1, C_1, R_1 and C_2. C_3 and C_4 capacitance values are those given by Adachi in his sample circuit. A bi-polar transistor is used because that's what was easiest. This has finite input impedance, but that will only change the gain required, and the BC547 has plenty of that.

The first 5 ms of the transient response from turn-on of the waveform at the emitter is shown in figure 8. The oscillator takes about 0.6 ms to start. The trace shows as completely black because of the high frequency - there are about 10,000 cycles per millisecond.

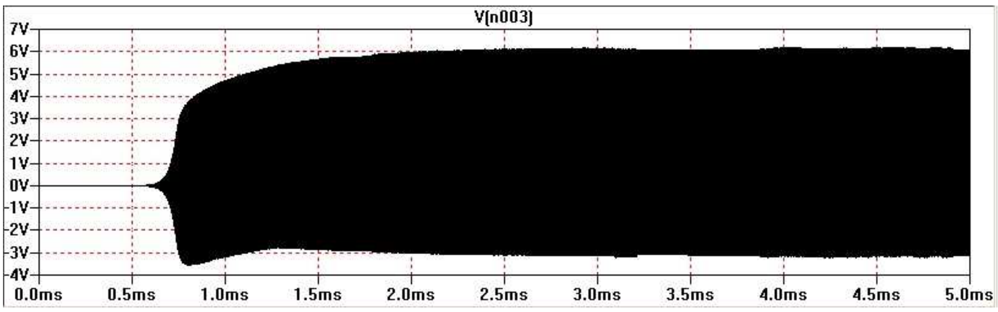

Figure 8: The first 5 ms of the crystal oscillator waveform at the emitter after simulation start.

Problems with Simulating Oscillators

`SPICE` Oscillator simulations are always tricky. Since the state when all currents are zero satisfies the circuit constraints, nothing may happen for quite a while. Real oscillators are triggered into leaving this state by thermal noise in resistors providing small transients which seed oscillation growth. These are supplied in `SPICE` simulations by accumulating round-off error in the differential equation solvers. Crystal oscillators are even more tricky since the Q of the crystal is so high. Sometimes injecting a transient voltage or current somewhere at start-up helps. A real oscillator may start much faster than a simulation predicts.

[4]A marvellous, free, unlimited SPICE simulator by *Linear Technology*, freely downloadable form the web. Google on "LTspice".

Figure 9 shows an expanded view of a portion of the same waveform, near the end of the trace. Measurements on this show that the oscillation fundamental frequency is about 9.9450 MHz (not the frequency displayed in the figure, see the explanation below). This is disturbing, since it is *below* the resonant frequency, and *not* in the "allowed frequencies" window!

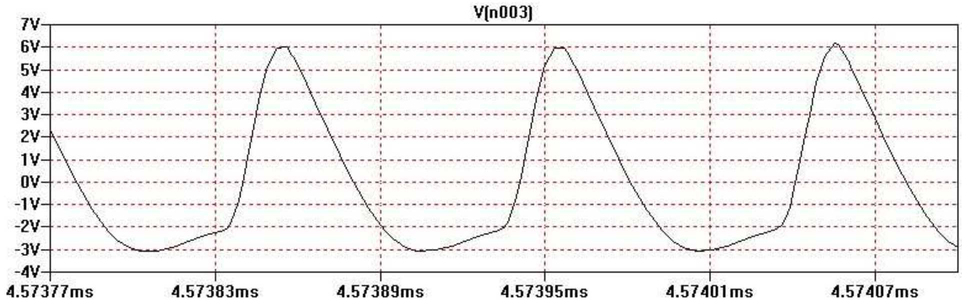

Figure 9: Crystal oscillator waveform

There are two reasons for this.

1. Even though the waveform appears stable, oscillations have *still* not reached their steady state, even after 5 ms, because of the very high Q (about 98,000) of the crystal. In this "oscillation growth" phase, the waveform amplitude grows exponentially, governed by Van der Pol's equation, with an instantaneous frequency which is *less* than its asymptotic steady-state value.

   Mathematically, in this phase, the complex conjugate poles of the transient response which form the oscillatory waveform have small but finite real parts, which slowly, and exponentially, decrease to zero as steady-state approaches. At the same time, their imaginary parts, which in the build-up phase reflect *slightly lower* frequencies than their steady-state values, are moving exponentially along the imaginary axes towards their asymptotic values. The exponential time constant increases with crystal Q, here very high, so the rate of increase is very slow.

   Thus, you have run the simulation for a long time before an accurate estimation of the final frequency is possible, and even then it will be slightly *lower* than its "infinite-time" value.

2. I have found that for accurate oscillator frequency estimation the maximum time-step of the simulation must be, in practice, small enough to give about 50 steps per waveform cycle.

The estimated frequency quoted above is below the allowed frequency window because the total simulation time was insufficient, and it was performed with no maximum time-step specified. By default, `SPICE` makes what it thinks is a reasonable choice of time-step, which it modifies as simulation proceeds. This works pretty well for source-driven simulations, where a network is energized with a sinusoid at a defined frequency. But this is inadequate here. For greatest accuracy in oscillator simulations, especially with crystals, you need to over-ride this default and specify a maximum time-step yourself. To illustrate the difference, another simulation was performed by

- specifying a maximum time-step of 2 ns (50 simulation steps per cycle),
- increasing the simulation time to 10 ms (about 980,000 simulation steps,)
- using only the *last millisecond* of the resulting waveform to estimate the frequency (using the supplied FFT function.)

An oscillation frequency of 9.99688 MHz now results, the value displayed in the figure. This is just above the crystal's series resonant frequency (9.996728 MHz) and thus within the window of possible oscillation frequencies. This simulation took `LTSPICE` several minutes. It grunted a lot and appeared to be swapping memory to and from disk.

The oscillation frequency is affected slightly by the values of C_3 and C_4 (the Colpitts feedback network components) since they define it, in combination with the inductive value of the crystal. If they change, a slightly different inductance will be required. But the inductive reactance of the crystal changes so rapidly in the "allowed frequencies" window that the frequency change will be small. Changing one, however, can be used to fine-tune the frequency. This will also change the loop gain, so the change can't be too great.

The Pierce Oscillator

This is another common circuit topology, where the crystal is connected directly between the drain (emitter, plate) of the active element and the gate (base, grid). See figure 10 (a), implemented using an *npn* transistor. Resistors R_1 and R_2 in combination with R_e form the *dc* biasing network, C_e bypasses R_e. This really *is* a bypass capacitor! C_1 and C_2 are the feedback network resistors needed to implement the 180^o in combination with the crystal.

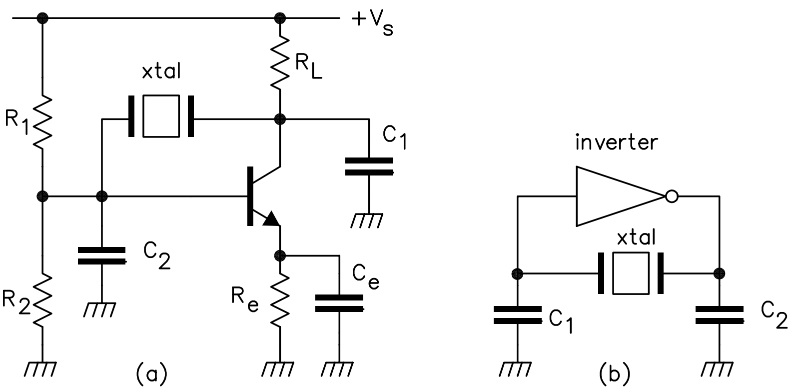

Figure 10: (a) Pierce Oscillator, as conventionally drawn. (b) Schematic of common microprocessor clock oscillator, also called a "Pierce".

It looks rather different when drawn like this, but in fact the small-signal equivalent circuit has exactly the same topology as that of the Colpitts, and is easier to see directly from the schematic. In fact, the only topological difference between this and the Colpitts circuit is that the ground (common) is now physically at a different point, the emitter of the transistor - as we have in fact re-drawn the Colpitts circuit in figure 3 (c).

However, the load is now in parallel with one end of the feedback network, and also the emitter-collector capacitance of the transistor. The effective load capacitance may change if the load varies, and hence the oscillation frequency may also change. For this reason, the Colpitts circuit is preferred for transmitter circuits.

Figure 10 (b) shows a circuit using a logic inverter often used in microprocessors. This is usually described as a Pierce type.