

Outdesigning Wozniak (Crystal Oscillators)

Jesús Arias

1 Introduction

Well... 45 years later, and I'm not really talking about logic design with TTLs, a field where Wozniak was a sheer genius (IMO his floppy controller should be displayed at Moma as a piece of modern art, deserving it a lot more than some paintings or sculptures there ;). I'm talking about the crystal oscillator of the Apple II computer.




Figure 1: Impedance of a 10MHz crystal and the same crystal in parallel with a load capacitance.

As with any piece of hardware designed by Woz, this crystal oscillator is a little peculiar. But let's first present the general problem of generating a periodic signal with a quartz crystal, beginning with the analysis of the crystal impedance. As we can see in figure 1 the impedance can change from a few ohms at the series resonant frequency to a value ranging the megaohm at its parallel resonance frequency. We must also remark that at these two frequencies the crystal impedance is purely resistive, and that for in between frequencies the crystal has an inductive reactance.

We can divide crystal oscillators into two categories depending on the impedance of the crystal at the oscillation frequency: In Series-Resonant oscillators the crystal impedance is very low, usually under 100Ω , while in Parallel-Resonant oscillators the crystal impedance is about a thousand times larger. Also, as we can see in figure 1, the capacitance of the circuit also affects the frequency of oscillation, with manufacturers

stating a loading capacitance for their crystals. Only when the loading capacitance in parallel with the crystal terminals is correct the oscillation frequency is the nominal one.

Historically, series-resonant oscillators were the dominant circuits during the TTL era (The Apple II oscillator is no exception), and became replaced by parallel-resonant oscillators in the modern CMOS era. The main reason for that is the high input impedance of CMOS circuits, a very convenient characteristic for parallel-resonant oscillators.




Figure 2: Examples of series-resonant oscillators from vintage computers

Several typical series-resonant oscillator schematics are shown in figure 2. All circuits are basically the same but every one includes some particular quirks as the recognition of these circuits being somehow tricky to run reliably. The simplest one is that of the Amstrad-CPC where the NAND gates are used as inverters and these, in turn, as linear amplifiers thanks to the biasing provided by the 330Ω resistors.




Figure 3: Inverter gate as amplifier

In figure 3 the transfer function of each inverter is shown along with the dependence of the DC input voltage with the output. As TTL gates are sourcing current at their inputs there is a voltage drop in R_f and the $V_i(V_o)$ dependence isn't a simple one (green curve). This bias current at the input of gates forces the use of small values for feedback resistors in order to achieve a convenient operating point. Otherwise the input would simply be high and the output low (for TTLs floating inputs remains at high level). If properly biased by R_f , the gate acts as an amplifier with high gain (the simulated inverter has a gain of -58.8).

But the biasing resistor also has the effect of lowering a lot the input impedance of the amplifier due to the Miller effect, where any impedance connected between the input and output of an inverting amplifier is seen as divided by the gain (actually the gain plus one) at the input terminal. So, in our Amstrad case the input impedance is $330\Omega/59.8 = 5.52\Omega$ in both amplifiers. The quartz crystal has a series resistance about

tens to hundreds of ohm, still low enough to provide the necessary feedback for the start of oscillations in spite of the severe attenuation introduced by the low input impedance.

2 The Apple II clock oscillator




Figure 4: Schematic of the Apple II crystal oscillator and a simplified equivalent circuit.

The schematic of the Apple II oscillator is shown in figure 4, where a TTL gate is used only for buffering while the actual oscillator is built using discrete transistors in a differential amplifier configuration. In this amplifier the input impedance is explicitly controlled by a base resistor of 47Ω , and it has a value similar to the series-resonance resistance of the crystal, limiting the feedback attenuation to about $1/3$.

A nice and clever circuit, it has still a couple of small inconveniences:

- It requires two different voltages to run. This wasn't a problem because the main memory also required the same voltages, but it would be better to run only on 5V if we want to use this circuit for other systems.
- It consumes a lot of power, almost $1/4$ of watt. Series-resonant oscillators are all power hungry and this certainly is, due mainly to the use of a 12V supply (current consumption is about 20mA)

So, here is still some room for improvement: The design of an equivalent but improved clock oscillator with a single 5V supply and way less power consumption while maintaining the capability of driving TTL logic.

3 Parallel-resonant oscillators and alternative Apple II clock

The typical circuit for a parallel-resonant crystal oscillator is shown in figure 5. Its main distinctive characteristics are the presence of explicit loading capacitors (33pF in the example), and the high input impedance of the amplifier. In this case the CMOS gate has almost no DC current at its input and a quite high feedback resistor, R_f , can be used, resulting in a high input impedance even when accounting for the Miller effect.

The capacitance at the crystal terminals is the series equivalent of the two 33pF capacitors, or 16.5pF, and we must also add the parasitic capacitance of the gate terminals, resulting in a value close to the usual nominal load capacitance of 18pF (some crystals may have other loading capacitances).




Figure 5: Typical CMOS, parallel-resonant, quartz crystal oscillator

In this circuit the only remaining device, the series resistor R_s , is used to further reduce the current consumption of the gate to a few milliamps.

The circuit of figure 5 is very ubiquitous today, but I want to build such an oscillator without resorting to fast CMOS gates. My idea here is to follow the Wozniak way and to use discrete transistors for the oscillator amplifier. The result is the circuit of figure 6.

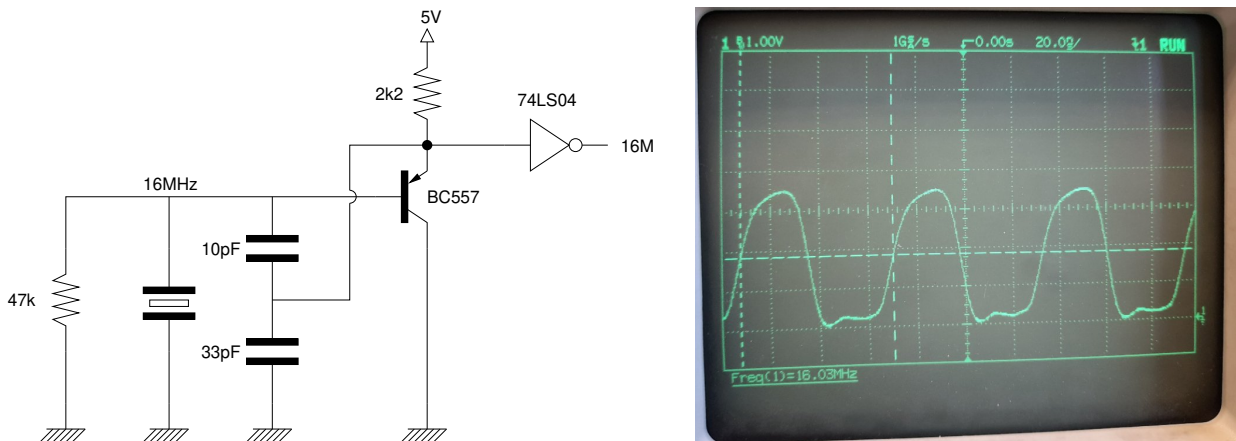


Figure 6: Proposed clock generator for TTL logic and experimental results

Here the amplifier uses a single emitter-follower PNP transistor with unity gain. The transistor is a general purpose one and can be substituted by many other models, (for example: 2N3906). The circuit consumes 1.8mA of current, and drives a 74LS gate that is fast enough for the intended frequency. When compared with the original Apple II oscillator this circuit uses less devices, requires a single 5V supply, and runs with 1/25 of the power.