CD Jitter

by
Paul Winser

Part 1, CD JITTER - The Problem

A lot has been written over the years on the subject of improving the analogue sections of CD players, and many DIY'ers can testify the benefits of better DAC chips, high speed op-amps, quality passive components and well designed power supplies. Quite a lot is also written on the subject of "jitter", with the implication that it effects the sound, and ought to be reduced wherever possible. And that's about it - stuff is written, and manufacturers do take steps to deal with the problem, but almost nothing has been done in the DIY scene as far as I know. I hope this offering will begin to put this right. In this issue I describe what the problem is and the reasons why it is as bad as it is. Part 2 gets more practical and describes an audiophile low-jitter clock subsystem for Philips style CD players.

What is Jitter?

Jitter is imprecision in the timing of the discrete samples when converting from analogue to digital or vice-versa. It should only be considered at the actual points of conversion. The inbetween storage of digital data and its recovery, transmission and processing are immune from the problem of jitter (assuming that the values of the samples and the order of their arrival are preserved). Only at the interface between the discrete time world of digital audio and the continuous analogue domain do we need to address the timing precision. Note that clock jitter is not related to the quantisation into digital code at all, only the quantisation into discrete time samples. Jitter would still be a problem if the samples were maintained as analogue values.

The effect of jitter is to introduce noise and distortion into the signal. A pure tone of one single frequency will be modulated by the timing jitter and cause a `spreading' of the signal frequency as shown in Figure 1.

If the timing jitter is random, the resulting effects will be similar to noise added to the signal. If the jitter is at all repetitive or correlated with the signal, bizarre modulation effects will occur. As we shall see, both of these are present in a CD playing system.

There is nothing we can do about jitter at the analogue to digital conversion end. This has happened in the recording process and the distortions introduced are there for good. We can only hope to lobby record labels into using the best equipment and techniques available to smarten up the sound of CD's in the future. Some of the bad reputation of CD sound rests firmly on the shocking quality of some masterings. Things seem to have become better, but if all CD's sounded as good as the few that are excellent, I for one would be much happier.

Why does it happen?

Left only with the possibility of improving the jitter performance at the D to A end, let's consider our typical CD system, with a transport section reading data from the disc, using the redundancy error correction data to deal with any errors, and passing the data stream on to the DAC, via a filter circuit which on the way typically multiplies the sample rate by 4 or 8 times to make the final analogue low-pass filter simpler. Whether this process happens in one box or two doesn't really matter, the point is that unless the transport is faulty or the disc has flaws, at the connection between the error correction and the digital filter we have a stream of correct 16 bit digital values. Yes the error correction does actually work, and the masking or interpolation function is only invoked for long bursts of uncorrectable errors such as happens when the disc is badly scratched. What this means is that any reasonable transport will deliver an identical data stream from any reasonable disc, regardless of how much green ink is around.

Given this, it is difficult to think of many ways in which transports can have differing sound. In fact I can think of only one: the timing precision with which the samples are delivered by the transport to the subsequent processing. In nearly all CD systems, the sample timing at the DAC is derived from a digital clock generator in the transport section, although in one-box designs the distinction is somewhat blurred. In the case of separate transport and DAC units, the situation is potentially much worse, since the sample clock is encoded into a single signal along with the serial 16 bit data stream (in "biphase-mark" format) and transmitted down a cable to the DAC unit where the clock has to be recovered and cleaned up by a phase locked loop (PLL) circuit. This involves keeping a "loose" crystal oscillator locked to the incoming clock over long durations but ignoring as much of the short term irregularities (jitter) as possible. Figures 2 & 3 show the case for a single box CD player and separate transport & DAC respectively.

Sources of Jitter

The jitter at the DAC device of a CD system has two causes. We shall consider the one-box case first because it's simpler. Firstly, the clock oscillator itself may have imprecise output timing. Just because the oscillator is driven by a crystal it doesn't mean it's perfect. There are good and bad oscillator circuits and in this context the ones that are bad have higher "phase noise" than the good ones. This phase noise (the same thing as jitter) is caused by the oscillator not being highly tuned to a single frequency, but wandering in frequency over very short durations. A good oscillator for low jitter has very low short term frequency variation, or low phase noise. It's the difference between the sound of hitting a bell or a frying pan. For low phase noise we want a highly tuned circuit with a large `Q' factor, which means low resistance in the resonant circuit. Other performance factors effecting phase noise include power supply rejection of the oscillator circuit. The crystal itself, being a peizo electric device, should also be isolated from physical vibration which could also upset the purity of the oscillator output.

The other source of jitter in CD players is the link between the clock source and the clock pin of the DAC device. Noisy digital circuitry in this path will introduce its own timing imprecisions and increase the jitter of the clock.

How bad is it?

If we look at the situation in an ordinary CD player, we find that the master clock oscillator circuit is part of a larger device, in Philips style players it's the SAA7220 digital filter chip. The oscillator circuit is almost certainly a standard cell of the type in Figure 4. Although performing quite adequately as the clock generator in most digital systems, its jitter performance is terrible. Its power supply rejection and temperature stability are also terrible, and remember this is just a tiny circuit stuck on the edge of a much larger chip, and will be subject to a large amount of power supply modulation noise, since it is fed off the main digital supply. This noise will be partly random digital rubbish and partly correlated to both the transport system and the data processing. Some of the noise will be due to the currents in the transport servo circuitry modulating the power supply, so hard-to-track discs requiring hard work from the servos might be worse in this respect than good ones. Whether this might explain things like the green pen phenomenon is open to speculation.

Having started with a poor quality clock source, the problems are compounded by the clock passing through a divider stage and output buffer, before leaving the filter chip. This circuitry is in CMOS switching logic and subject to the same effects of power supply noise introducing extra timing imprecision.

In the case of a two-box system, on top of these effects in the generation of the clock, it also has to survive being transmitted down a metre of coax (or optical fibre) while being mixed up with the data stream. The exact scheme how this is done is not important, all that matters is that the clock precision in the DAC now critically depends on the performance of the phase locked loop (PLL) used to extract and stabilise the clock signal from the biphase-mark serial data format. The PLL consists of a voltage controlled oscillator (VCO) whose control voltage is determined by comparing the incoming and local clock frequencies. The local clock is forced to track the incoming clock in order to remain in sync with the data, but the locally produced clock signal should be free from all the transport and cable induced short term jitter. In practise, this is not achieved for two reasons. The VCO is, although crystal based, by definition not highly tuned to a specific frequency, but can be varied by means of its control voltage. Its `Q' is low so its phase noise is intrinsically high. Secondly, the scheme is something of a fudge anyway, because we want to track long-term variations in the incoming clock but ignore short term ones. This is done by means of a low-pass filter in the control voltage of the VCO, and the parameters of this filter are something of a compromise. In practise, preserving synchronisation with the data stream takes first priority, and jitter rejection comes second.

I suggest that the fact that different digital interconnects (coax & optical) sound different tells us that the PLL scheme does not work. This interconnect is only carrying digital pulses, and assuming that the receiver circuitry can correctly distinguish a `high' from a `low', the only difference can be in the timing and shape of the pulse waveform that is received. The waveshape and restoration of its transition edges will depend strongly on the bandwidth and characteristic impedance of coax cable, or the transfer functions of optical transducers, and the analogue PLL circuit will not be able to fully reject these differences in the input signal timing precision.

Elektor magazine recently published a project for a PLL jitter reducer to be inserted in the link between transport and DAC unit [Ref. 1]. Their circuit still suffers from all the problems I have outlined, and their results show a slight

improvement with optical link, and none with coax. To really solve the problem requires a complete rethink of the clocking system.

Linn have done this with their Numerik 2-box CD system. The master clock lives in the DAC unit where it belongs and the transport clock is slaved to it by a special second connection between the two units. Has anyone had the opportunity to hear what it sounds like?

Well that's enough hand waving. Next time we shall put theory into practise and describe the design of a high quality clock source that is now being put together. I very much hope that the results will be worth reporting.

References

1. T. Giesberts, "Digital Audio Enhancer", Elektor Electronics, Feb. 1993.

Part 2, CD JITTER - A Solution

In part 1 (Newsletter 9) we saw how jitter in the digital to analogue conversion process introduces distortion in devices such as CD players. Jitter is imperfection in the regularity of the timing clock that drives the DAC, and results from compromises in the design of the player. Two-box CD systems, potentially superior in some ways, are actually worse in this respect because of the problems involved in transmitting the timing signal and the data itself down the same piece of wire, between transport and DAC.

What follows is my attempt at fitting a precision timing source into a one-box Philips style multi-bit CD player (with the SAA7210, SAA7220, TDA1541A chip set). The design is modular so that the sections can be used or adapted to suit other systems. It may also be used with 2-box players, without compromising the performance very much. Figure 1 shows a block diagram of the player before and after `augmentation' with the new clock source, and we shall now look at the different sections in detail.

The crystal oscillator in the player is disabled. This was simple a 11.2896MHz crystal attached to the filter chip SAA7220. Instead, we use an identical crystal in the series-resonant circuit in Figure 2. This circuit is lifted directly from Matthys[1] with component values adjusted for 11MHz. The active device is an ECL differential receiver 10216, which internally is simply a long tailed pair with emitter follower outputs (see the appendix for more about ECL circuits). This circuit was chosen because of its high performance and low parts count. Note the low value load resistor R1, combined with the low series resistance of the crystal (approx. 10R) give the circuit a high Q factor and good short-term stability, which is just what we want for low jitter.
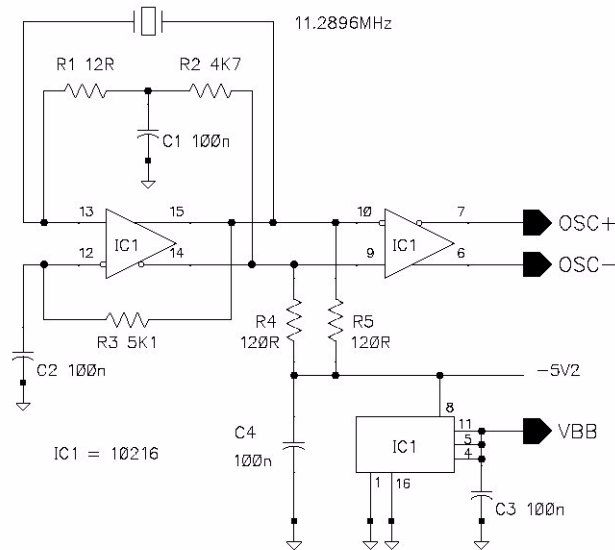

Figure 2. Crystal Oscillator

IC1 has three separate receivers on it, and the second is used as a buffer for the oscillator output. The third is unused. We do, however, use the bias voltage VBB later on. Because ECL circuits are single-ended, we need load resistors R4,R5 on the outputs. The differential outputs from the buffer OSC+ and OSC- pass to the circuit in Figure 3, the system clock interface. This section performs a number of tasks. Firstly it provides a clock output SYSCK back to the SAA7220 chip to replace the one we disabled. This is done by the differential pair T1,T2 driven by OSC- and VBB. Because we want fast precision timing throughout the circuit, we should minimise capacitive loading on the signals, which would otherwise slow down the edges. T1's collector is grounded, and T2's is cascoded by T3 to form the output, a configuration which prevents T1's Miller effect from presenting a large effective capacitance to IC1 output OSC-.
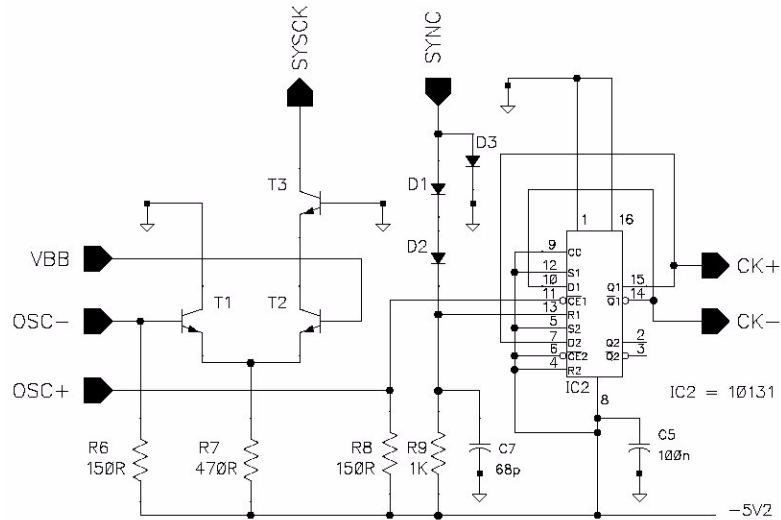
Figure 3. System Clock Interface

The other output OSC+ from the oscillator feeds a divide by two circuit 10131, IC2. This is required because the TDA1541A 4xoversampling DAC takes a bit clock of half the system clock frequency. The SAA7220 performs this function internally on its own oscillator but we have to do it ourselves for the new subsystem. IC2 is another ECL chip containing two flipflops. One of these is used to divide OSC+ by two in frequency to create differential outputs CK+ and CK- at 5.6448MHz. The SYNC input to the circuit serves to synchronise the dividers in SAA7220 and Figure 3 so that the 5.6448MHz clocks are always in phase and not 180 degrees out. This is important for when we recombine the filtered data output from the SAA7220 with the new bit clock that we have generated, CK+ & CK-.

This recombination should take place, together with clock buffering, as close to the DAC as possible. For this reason, we need special provision for transmitting the clock from the circuit board of Figures 2 & 3, to wherever the DAC is. This might be a few inches away in a one-box player, or a few feet with an outboard DAC, and we want to preserve the integrity of the signal as well as we can. Figure 4 shows the circuit of the clock line driver, a differential voltage to current convertor driven by CK+ and CK-. T4,T5 are a differential pair fed by current source T6, which keeps the combined collector currents of T4,T5 at a constant 2mA. T7 and T8 cascode the differential pair, eliminating their Miller capacitance effect and isolating the differential pair from the voltages on current outputs TCK+ and TCK-. These outputs can be twisted together with a ground wire in Litz formation to provide the clock connection to the DAC.

TCK−

TCK+

T7   T8

CK+

CK−

T4   T5

R10
430R

T6

D4

R11   R12   R23   R13
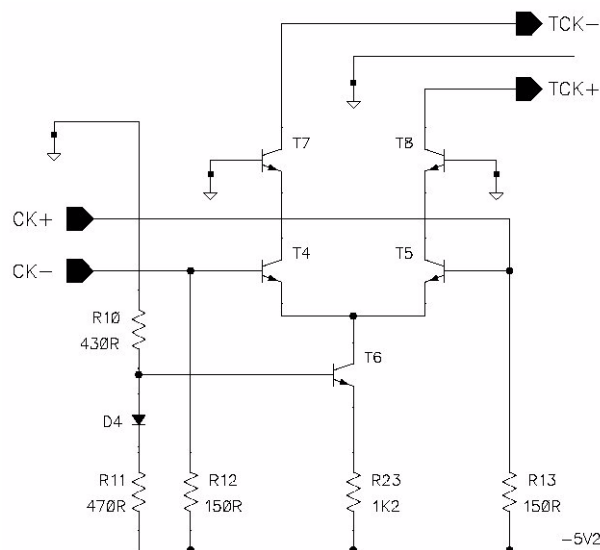470R   150R   1K2   150R

−5V2

Figure 4. Clock Line Driver

At the receiving end of TCK+,TCK- we find the circuit of Figure 5. The job of this section is to convert the differential current clock signals into a voltage drive for pin 2 of the DAC chip TDA1541A. I wanted to use a discrete transistor stage for this function, but the more I thought about it the less attractive it became. We need to produce a clock signal with 3V to 4V swing, with edge rise and fall times as fast as practical, certainly no more than a few nanoseconds. The problem is that these edge rates require a complementary push-pull stage, and a simple transistor arrangement is bound to operate with both transistors momentarily `on' during the switching, which will cause large current spikes to be drawn from the 5V supply. Preventing this action would take a fiddly arrangement of discrete components. We would much rather have a compact arrangement which we can retrofit into a small space very close to a DAC chip on an existing circuit board. Hence IC3, a fast comparator LT1016 with TTL logic level outputs. This chip provides the speed that we need and is specifically designed to minimise switching transients on the supply.

DABD

WSBD

+5V

R16   R14
220R   220R

C6
100n

4 PRE   5   3 DATA
3 CLK
2 D
1 CLR
IC4a   TDA1541A

TCK+

2 1   7

10 PRE   9   1 WS
11 CLK
12 D
13 CLR
IC4b

3   IC3
     8

TCK−

4 5,6

R17   R15
330R   330R

2   BCK

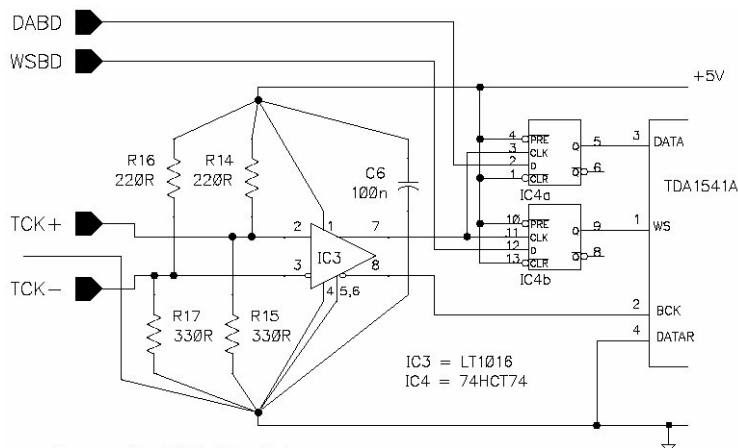4   DATAR

IC3 = LT1016
IC4 = 74HCT74

Figure 5. DAC Pin Driver

The inputs to IC3 are biased into their linear common-mode region by R14-R17, which also provide a 130R termination resistance to the transmitted clock signals TCK+,TCK-. Termination is essential on these lines to prevent reflections and waveshape deterioration of the fast signal edges. 130R presents a reasonable match to most cable types. We can take advantage of the complementary outputs of IC3 - one output goes straight to the bit clock input BCK of DAC TDA1541A, with as short a length of wire as possible, certainly no more than two inches. The other output can be used to reclock the two data inputs to the DAC, DABD and WSBD from the SAA7220. These are the serial 16 bit data and left/right word select lines. Reclocking them will ensure that they are noise free, and also that

when the BCK input switches, all other DAC inputs are at a steady voltage, preventing any cross-coupling from disturbing the timing of BCK.

Two flipflops in IC4 (an ordinary high speed CMOS device 74HCT74) perform the reclocking function. The relative timing of TCK+,TCK- and DABD,WSBD depend on the synchronising function of Figure 3 to be correct. Of course once the bit clock has entered the DAC, it may internally be subject to timing disturbances, but there's nothing we can do about that except to ensure that the power supplies to the DAC chip are as clean as possible. This means voltage regulation up to frequencies of 100MHz or so. Common three-terminal IC regulators are effective only up to about 100KHz, so at the frequencies relevant to jitter, we rely entirely on decoupling capacitors on the digital supplies to provide the smoothing, but I think that's another story.

As a bonus, switching between the complementary data outputs pins 5 & 6 of IC4a gives a signal polarity inversion function in the digital domain.

Lastly, Figure 6 shows the player interface circuitry around the SAA7220. T9 buffers signal SYSCK and applies it to pin 11, XIN. This pin was originally connected to the crystal and a loading capacitor, both of which should be disconnected. T10 is driven from bit clock output CLBD pin 16, which usually connects straight to the BCK input of the DAC. Here we use it to form the SYNC input to the divider of Figure 3. T10's emitter connects not to +5V but to the player's KILL signal (test point 93), which is at +2V for approximately a second after switch-on, to mute the outputs, and at -9V thereafter. This causes the SYNC signal to applied only for that second after switch-on. Once the dividers are synchronised they will stay that way, and we don't want unnecessary clock signals potentially disturbing the circuitry of the new clock system while we are playing discs.
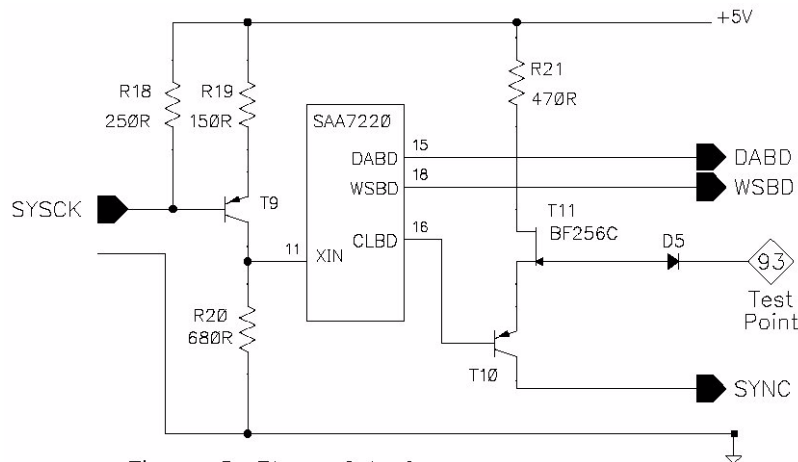

Figure 6. Player Interface

Power Supplies & Construction

Most of the new circuitry requires a -5V2 supply, at a total of around 250mA. This voltage is set by the ECL devices. The supply should be free from high frequency noise, and be of low impedance. An LT337 adjustable regulator will do the job, but since the current load will be pretty constant once the circuit has warmed up, we could use a local transistor shunt regulator that will have better high frequency performance. Do not daisy-chain the supply around the devices, use a star arrangement with local decoupling at each load (see below). We also steal about 25mA from the DAC +5V supply in Figure 5, for IC3 and its bias resistors. The diagram indicates how the supply and ground plus bias resistors and decoupling for IC3 should be starred.

Building these circuits requires a little care. IC's 1, 2 and 3 are high gain, high bandwidth devices, and given the slightest opportunity to oscillate at 100MHz they will do so. Things needn't be that scary as long as a few simple rules are applied. Firstly, use circuit board that has a ground plane. There are plenty of prototyping boards designed for digital logic that are suitable, and we only need two small pieces. I can supply suggested component layouts to anyone interested.

Keep all wires connected to IC1 & IC2 as short as possible. Mount the loading resistors of the outputs of these devices (R4,5,6,8,12,13) at the far ends of the wires and not at the outputs themselves. All of the IC's should be decoupled directly across their power pins with 100n multilayer ceramic (gasp!) capacitors. Do not be tempted to use physically large polypropylenes for this purpose - these are digital circuits and we don't care about linearity or temperature coefficients. What we must have is effective decoupling up to 100MHz, and the long leads of large caps will have sufficient inductance to prevent this. All wires that carry clocks between circuit boards must have a ground twisted with them. This applies to SYSCK, SYNC and TCK+,TCK-.

The crystal is a piezoelectric device, so its electrical performance may be effected by physical vibration. Best to mount it on a lump of Sorbothane, then glue a damping mass of metal to it and connect the leads with thin stranded wires about half an inch long. Ideally the board with the oscillator and associated circuits should be shielded to minimise radiated noise. My favourite way of doing this is to cut six pieces of copper-clad to make a box to house the board. Five sections can be joined by soldering along the seams with the copper on the inside surfaces. The sixth piece serves as a lid, attached by copper braid hinges that also serve as earth connection.

How does it sound?

Impending holidays mean that I'm writing this instead of doing a proper installation and listening tests! I have tried a quick lash-up and I think the new arrangement noticeably refines the top end. Bells and cymbals have a more lifelike quality and instruments are that much better resolved. A full evaluation will have to wait for later. Hopefully I can demonstrate the effect at January's DAC Day.

Adaptions

The sections of this design can be mix'n'matched to suit other CD systems. For example, an 8xoversampling DAC may require a bit clock of 11.2896MHz, so the divider of Figure 3 can be left out. Most people, though, would want to adapt this design to work with a two-box system. This can be done without too much trouble since the clock drive signals TCK+,TCK- should work well over connections of a couple of feet, as long as they are terminated as shown and the cable used has low loss at VHF frequencies. Teflon coated wire-wrap wire in Litz configuration should do the job, but it might be better to used screened dual core high speed data transmission cable.

The data signals DABD and WSBD shout not be run over distance as they stand. They should be buffered, and if transmitted at logic levels of 4V, they must be well screened. Buffer and terminate at the DAC end. Alternatively, they could be transmitted in the same way as the clock signals and recovered with their own LT1016 devices. This is the belt-and-braces solution. If anyone wants to try or experiment with this design, I'm very happy to discuss ideas with them. After all, at 379 pounds, the Audio Alchemy DTI is an expensive band-aid measure for a problem that only exists because of the single cable standard for connecting transport to DAC. There are good reasons for doing it, but they are all to do with convenience and standardisation, at the expense of sound quality. Even the DTI is limited by the performance of its phase-locked loop and the receiving circuitry in the DAC unit, and so is only making the best of a bad job.

Component Sources

The crystal is specific to CD players. Best to salvage one from an old unit. Don't use the one from the player being modified as it will make it hard to compare `before and after'. The crystals are also available from Philips as a spare part, no. 4822 242 71644. The ECL devices IC1 & IC2 are stocked by ESD (formerly STC), and IC3 & IC4 can be bought from RS.

All NPN transistors should be high frequency, fast switching types. I used BC548. Diodes are IN4148 or similar.

References

1. Crystal Oscillator Circuits, Robert J. Matthys, John Wiley & Sons, 1983.

Appendix: ECL Circuits

Emitter Coupled Logic differs from the usual bipolar and CMOS technology of common digital logic. Traditionally used for very high speed systems, it is now serving only niche applications because of the rise in speed of CMOS devices. Unlike other logic families, ECL employs a low voltage swing of around 1V, driven by non-saturating transistors. Logic elements work by current steering rather than voltage switching. It is suitable for this design because the switching speed is not only high, but as the transistors do not saturate, there is no charge storage delay effect so the switching is precise and regular. Because it is single ended, ECL circuits require load resistors, and these need to be of low value to maintain switching speeds, so the power consumption is relatively high. Noise immunity is highest when signals are referenced to the positive supply, so that is generally grounded, and a negative supply of - 5.2V provided.